

SCIENTIFIC COMMUNITY

Selective and deceptive citation in the construction of dueling consensus

Andrew Beers^{1*}, Sarah Nguyễn², Kate Starbird¹, Jevin D. West², Emma S. Spiro²

The COVID-19 pandemic provides a unique opportunity to study science communication and, in particular, the transmission of consensus. In this study, we show how “science communicators,” writ large to include both mainstream science journalists and practiced conspiracy theorists, transform scientific evidence into two dueling consensus using the effectiveness of masks as a case study. We do this by compiling one of the largest, hand-coded citation datasets of cross-medium science communication, derived from 5 million Twitter posts of people discussing masks. We find that science communicators selectively uplift certain published works while denigrating others to create bodies of evidence that support and oppose masks, respectively. Anti-mask communicators in particular often use selective and deceptive quotation of scientific work and criticize opposing science more than pro-mask communicators. Our findings have implications for scientists, science communicators, and scientific publishers, whose systems of sharing (and correcting) knowledge are highly vulnerable to what we term adversarial science communication.

Copyright © 2023
The Authors, some rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
License 4.0 (CC BY).

INTRODUCTION

Public consensus often diverges from scientific consensus (1, 2). Something is lost in translation from the scientist to the public, as “science communicators,” writ large to include both mainstream science journalists and practiced conspiracy theorists, reshape and transform scientific evidence in ways that may deviate from its authors’ intents (3). For example, in 2015, a paper suggested that cloth masks increase the risk of respiratory illness in hospital contexts (4). It became one of the most shared papers about masks on social media during the coronavirus disease 2019 (COVID-19) pandemic, was quoted in one of the most popular television news programs in the United States, and was cited in legal challenges to mask mandates (5, 6). However, the paper’s first author wrote more than 10 papers over a 2-year period during COVID-19 affirming the benefits of masks and speaking out against the selective and misleading understanding of the 2015 paper, before and after its popularity on anti-mask media (7–16).

How does this reimagining of scientific consensus outside of the hands of scientists themselves happen in today’s media-rich environment? What are the implications, and what can be done about it?

Here, we hand-curate one of the largest corpora for investigating these questions. Leveraging the high demand for research findings about masks during the early stages of the COVID-19 pandemic, we show via citation analysis how science communicators and practiced conspiracy theorists alike manufacture the appearance of consensus using selective and critical citation to create dueling, opposite perceptions of scientific knowledge. We find that scientific publishers and authors are unprepared for how their work is being misused by science communicators to advance scientific misunderstanding.

Science communication during a pandemic

Rarely have so many people been so personally invested in emerging scientific discoveries, and never before has the process of science communication been so thoroughly documented on social media as it has during the COVID-19 pandemic. These circumstances have shown vividly how varied the sources of scientific knowledge have become: not only scientists but also journalists (17), government officials (18), celebrities (19), viral videos (20), and social media users themselves (21, 22). We present and analyze here two of the largest hand-coded datasets of scientific citation to understand the development of a pandemic controversy around masks in the United States. The first dataset shows how scientists built and shared among themselves a varied evidence base on the effectiveness of masks for COVID-19. The second dataset shows how science communicators, credentialed and not, transformed this evidence base into two different, opposing perceived consensus: one in favor of masks and one against. In documenting this process of consensus transformation, we have three primary findings:

1) While scientists largely agreed that masks were effective at reducing the spread of infectious disease and tended to cite works with similar methodologies, science communicators sharply disagreed and tended to cite work with similar stances.

2) While scientists rarely engage in negative citation, science communicators and particularly anti-mask science communicators more frequently disparage and criticize scientific work.

3) Science communicators cite much of the same literature as scientists. However, consensus diverge via misleading contextualizations that often contradict the scientists’ publicly stated views.

While popular rhetoric has held that those who disbelieve mainstream science are “anti-science” (23, 24), our dataset describes a communicative world where both pro- and anti-mask social media users are attuned to scientific norms and invest heavily in citing the primary literature (25). In what is akin to context collapse (26), anti-mask science communicators gladly repurpose the knowledge infrastructure of mainstream science to spread misinformation in the language of science itself (27). These science communicators’ frequent use of negative citation (28) can be seen as

¹Department of Human Centered Design and Engineering, University of Washington, WA 98195, USA. ²Information School, University of Washington Seattle, WA 98195, USA.

*Corresponding author. Email: albeers@uw.edu

inoculating their audiences against the views of mainstream science, the mirror of a strategy of “prebunking” and “inoculation” advocated by science communication researchers (29–31). Anti-mask science communicators’ deceptive but likely effective citational practices reveal how scientific authors and publishers are currently woefully unprepared for the ways that their articles are misrepresented to the public (32).

Conceptually, this work bridges consensus research in science communication, which focuses on the public communication of science (1, 33–35), science and technology studies, which focuses on how consensus is constructed and then understood by both scientists and the public (36–38), and the science of science, which focuses on consensus formation through citational networks (39–42). We focus on the role of formal (government health bulletins) and informal (anonymous viral Twitter posts) science communication. The work is inspired by the concept of knowledge infrastructures: the “robust internetworks of people, artifacts, and institutions which generate, share, and maintain specific knowledge about the human and natural worlds” (43). By centering both formal and informal knowledge creators as objects of analysis, we create opportunities for quantitative and qualitative inspections of knowledge infrastructures and the way they shape public health outcomes.

RESULTS

The creation of evidence bases by scientists

In these findings, we make a distinction between scientists, who publish their work for consideration by the scientific community, and science communicators, who publish their work for consideration by the lay public and may or may not include formally trained scientists. We assume that the goal of science communicators in crises is to translate the (sometimes developing) consensus of scientists into public-facing information about probable risks and recommended responses (44). In practice, we find that popular science communicators in the United States are sharply divided on the benefits of masks for preventing COVID-19. Thus, our first task is to understand whether these science communicators are faithfully transmitting a true dissensus among scientists themselves or are generating it from whole cloth.

To assess the existing evidence base on mask wearing at the start of the pandemic, we use the Web of Science (45) to identify papers about masks that were highly cited in the first year of the COVID-19 pandemic. We annotate all citations in these papers and mark whether they are positive/neutral citations or negative citations. We define negative citations as citations that go beyond discussing a work’s limitations or conflicting results to actively recommend that readers at least partly disregard the cited work. We construct a signed (i.e., positive and negative citations specified through hand coding) cocitation graph of these highly cited papers to identify subsets of papers that tend to be cited together (Fig. 1). We proceed under the assumption that papers frequently cocited positively together, within a topical area, represent endorsements of shared evidence on the topic of masks in COVID-19 (46).

We find three evidence bases that correspond to different methodologies and fields that highly cited masks papers tend to reference. Specifically, these papers cite (i) studies grounded in particle physics that evaluate the material properties of masks, (ii) studies grounded in epidemiology, which evaluate masks effectiveness on different metrics during the first year of COVID-19, and

(iii) studies grounded in public health, which evaluate the effectiveness of masks in different contexts, such as hospitals or households, and evaluate practical considerations such as the disinfection of masks. Each evidence base is composed of papers that mostly suggest that masks are effective for preventing respiratory illness and especially COVID-19.

We find that, in general, scientists use negative citation sparingly, with only 1.3% of citations across 80 papers manually coded as negative (32 of 2426 citations). This percentage is notably lower than previous estimates, which have ranged from 2 to 15% (28, 47, 48). Affirming Bruggeman *et al.* (39), we find that negative ties, while rare, are not randomly distributed throughout the citation corpus. Rather, they are disproportionately directed at a few papers, including both papers with positive and negative results toward masks. The most negatively cited paper, with four negative citations, is MacIntyre *et al.*’s 2015 paper in *BMJ Open* that “cautions against the use of cloth masks” for preventing respiratory illness in hospitals and suggests that they may increase the risk of infection to the wearer (4). This paper is particularly notable because its lead author has subsequently published no fewer than 10 publications over a 2-year period arguing for the use of masks of all types to lessen the spread of COVID-19 (7–16). Among these publications is a letter titled “Masks in the community are an effective strategy” and a follow-up to the criticized 2015 study offering a possible explanation for their earlier, negative findings around cloth masks (11, 13). Chu *et al.*’s 2020 paper claiming that “face mask use could result in a large reduction in risk of [respiratory] infection” was also negatively cited three times primarily on methodological grounds, some of which are described by Jüni *et al.* in a letter response published in the same journal (49, 50). Overall, scientists negatively cite papers that have perceived methodological weaknesses rather than papers that share a certain stance on the effectiveness of masks.

In this first year of the pandemic, the consensus on the effectiveness of masks was still developing, and new studies continue to be published since our analysis period. We show here not a proof of consensus that mask wearing is an optimal choice for combating COVID-19 but rather the creation of a shared literature base from which scientists, within their own fields, collectively endorsed and used to move scientific knowledge and work forward. In the following section, we describe the different way that science communicators select literature bases, namely, by their ability to support a pro- or anti-mask stance.

The creation of perceived consensus by science communicators

Science communicators, by contrast, come to remarkably diverging conclusions about the effectiveness of masks within their own community. To surface the published works of science communicators, we analyze the URLs shared in posts by Twitter users discussing masks. These URLs sometimes link directly to the published scientific literature, but more often, these URLs link to the works of science communicators, broadly defined: science journalism articles, government- and hospital-sponsored websites, self-published personal blogs, viral social media posts, and videos hosted on a variety of platforms. Some examples of frequently shared URLs that support mask wearing include, with screenshots in the top row of Fig. 2, the following: (i) “Still Confused About Masks? Here’s the Science Behind How Face Masks Prevent Coronavirus,” hosted by the University of California San Francisco (UCSF)




Fig. 1. A visualization of the cited literature of scientists and science communicators with respect to mask wearing and COVID-19. Papers that tend to be cited together by (A) scientists and (B) science communicators discussing masks and COVID-19. Scientists tend to cite groups of papers by topic or field. In (A), the green cluster represents papers concerned with the physical properties of mask materials; the orange cluster represents epidemiological papers on masks effectiveness against COVID-19, and the purple cluster represents older public health papers on the effectiveness of masks in hospitals and households for COVID-19 and other respiratory illness. In (B), meanwhile, science communicators cite largely the same papers as scientists but segregate their citations by stance: supporting the use of masks in blue and opposing them in red. Three papers (nodes) are labeled in both clusters: (1) (62), the most shared paper in both datasets, (2) (89), a fluid physics paper from 2010 shared in the anti-mask perceived consensus, and (3) (4), a 2015 paper with a negative finding on masks whose lead author has since publicly supported mask wearing for COVID-19. Larger nodes in the scientist network (A) are cited by more highly cited papers on masks, while larger nodes in the science communication network (B) are cited by more frequently shared Twitter URLs in a dataset of people arguing about masks.

Medical School (51), (ii) "Guidance for Wearing Masks" hosted by the U.S. Centers for Disease Control and Prevention (CDC) (52), (iii) a video hosted on YouTube titled "How Well Do Masks Work? (Schlieren Imaging In Slow Motion!)" published by the U.S. Public Broadcasting Service (PBS) (53), and (iv) a long, threaded Twitter post that aggregates links to scientific studies purportedly supporting masks, created by a pseudonymous Twitter user (54).

Conversely, examples of highly shared URLs that oppose mask wearing, again with screenshots in the bottom row of Fig. 2, include the following:

1) "Mask Facts," a website hosted by the Association of American Physicians and Surgeons (AAPS), a medical conspiracy group aligned with the U.S. libertarian movement, which opposes the scientific consensus on vaccines, HIV/AIDS, climate change, and abortion among other topics (55).

2) "Masks Don't Work: A Review of Science Relevant to COVID-19 Social Policy," a self-published literature review authored by a scientist who also contested the scientific consensus on climate change and vaccines. This review has been rehosted on many different websites, including the science aggregator ResearchGate (56).

3) A video hosted on YouTube titled "Why Face Masks DON'T Work, According To SCIENCE" created by a journalist then employed by RT, a Russian state-owned news outlet known to publish false and misleading content. The author advocates for counterconsensus views on vaccines and a variety of conspiracy theories related to mass shootings in the United States, the Syrian Civil War, and other topics (57).

4) "The Science of Masks," a website and companion video that links to more than 60 studies and news articles about masks. It is hosted on the website of an alternative medicine podcast host that sells diet planning services, books about wellness, and speaking opportunities (58).

5) A long, threaded Twitter post that aggregates links to scientific studies and other materials purportedly opposing masks, created by a pseudonymous Twitter user (59).

The science communicators behind each of these websites, whether pro- or anti-mask, explicitly cite peer-reviewed scientific literature to explain their positions on the effectiveness of masks. Anti-mask science communicators cite the scientific literature in ways similar to pro-mask communicators, for example, by using explicit "References" sections with formatted citations. They cite non-scientific sources, such as personal blogs or anecdotal news stories, only rarely. Most anti-mask communicators use sophisticated, scientific language to contextualize their findings, for example, referring to "randomized controlled trials" and "statistical significance," and likely give the impression of a knowledgeable expert explaining the current state of scientific consensus.

To understand how science communicators reference science differently from scientists, we code their citations just as we did with the Web of Science data. We find that the proportion of negative citations is notably higher from science communicators than scientists themselves at 6% (89 of 1496). This proportion rises to 9% for websites opposing or reporting negative findings on masks compared to 4% for websites with neutral or positive findings. Overall, anti-mask websites averaged 1.5 negative citations per website, while other websites average 0.7 citations. Several anti-mask websites have sections specifically devoted to refuting research that supports masks. For example, one such section is titled "Popular Mask Effectiveness Studies That Ignore Real World Conditions," and lists the titles of pro-mask papers accompanied with short explanations of their supposed flaws in scientific language.

As with the Web of Science data, we use this dataset to create a signed cocitation network (Fig. 1). We see several small groups of two to seven papers, each likely the result of different science

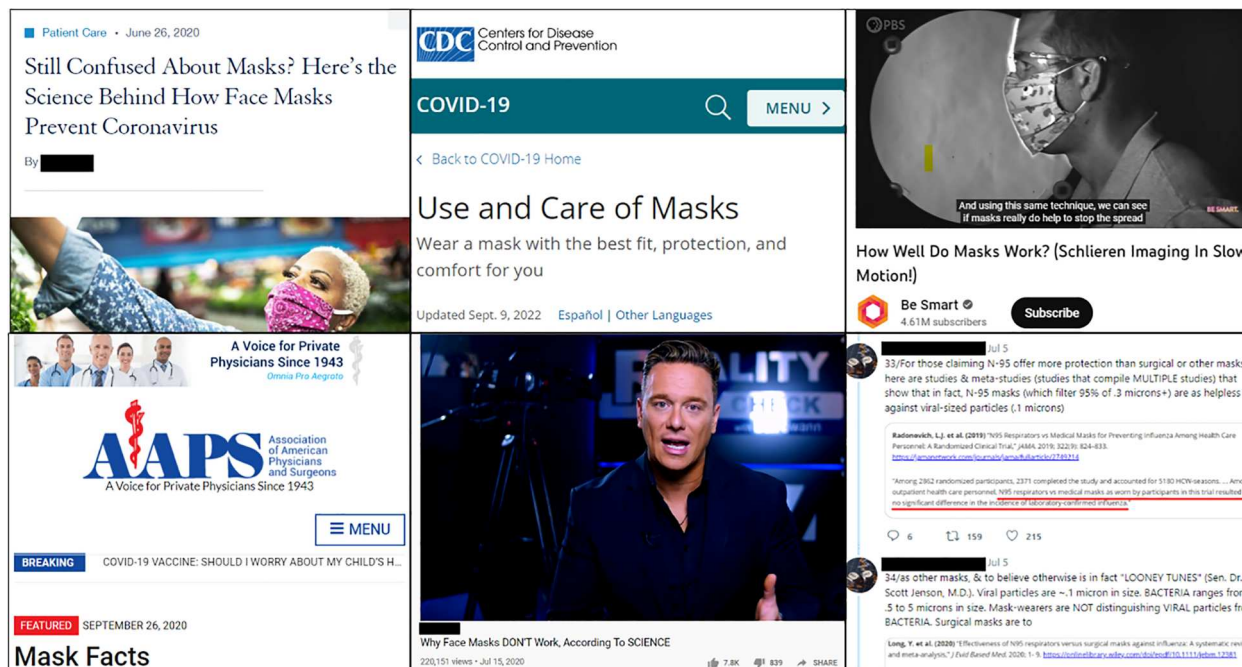


Fig. 2. Screenshots of six popular science communicators in our dataset. The top row are examples of generally pro-mask science communicators, including, from left to right, an article published by the University of California San Francisco (UCSF), a help page published by the CDC, and a YouTube video published by the PBS. The bottom row contains generally anti-mask science communicators, including, from left to right, a page published by the medical conspiracy group, the Association of American Physicians and Surgeons (AAPS), a YouTube video published by a conspiracy theorist formerly employed by Russian state media RT, and a viral threaded Twitter post by an anonymous user containing URLs and screenshots of scientific articles. We have appended black boxes to the first image in the top row, as well as the second and third images in the bottom row, to remove names of individual authors.

communicators using extremely similar citation lists or, in some cases, directly copying and excerpting each others' pieces. The rest of the graph is dominated by two large groups of papers that, unlike citations in the scientific literature network, group together by stance. One group is pro-mask and consists entirely of literature from mostly high-impact scientific journals whose results support the claim that masks are at least somewhat effective for preventing the spread of COVID-19 or other respiratory illnesses. A typical example of a paper in the pro-mask evidence base is that of Leffler *et al.* (60), which states in its abstract that "societal norms and government policies supporting the wearing of masks by the public, as well as international travel controls, are independently associated with lower per-capita mortality from COVID-19."

The other group is anti-mask and consists of the scientific literature and one university-affiliated website with positive, qualified, or negative findings about masks, almost all of which are construed to appear as negative findings in the context of their citations. A typical example in this cluster is the work of Jacobs *et al.* (61), which argues in its abstract that "face mask use in health care workers has not been demonstrated to provide benefit in terms of cold symptoms or getting colds." Both groups of papers contain almost exclusively papers published in academic journals, and neither cluster displays a clear preference for papers from the fields or journals identified in the scientific literature network.

Both the pro- and anti-mask paper groups leveraged papers from the three different fields found in the scientific literature, extracting those that can be made to appear in agreement with their chosen stance. Some papers are cited positively by both pro- and anti-

mask science communicators: Fifteen percent of incoming citations to the pro-mask cluster come from websites with anti-mask orientations, and 23% of incoming citations to the anti-mask cluster come from websites with pro-mask orientations. However, rather than take these cross-stance citations as evidence of shared scientific knowledge between pro- and anti-mask perspectives, we find that papers are often contextualized differently between pro- and anti-mask websites (see Table 1 for excerpts). Accordingly, the separation between the evidence bases that readers perceive may be even wider than these data can illustrate. We discuss two such papers at length to illustrate two distinct tactics by which ostensibly pro-mask papers can be contextualized to appear to advise against masks: selective quotation and deceptive quotation.

We first examine the work of Leung *et al.* (62), the paper that was most often cited by science communicators in our dataset. Inspecting how this paper was cited in text, we find that pro- and anti-mask websites selectively quote entirely different parts of the paper. Pro-mask science communicators tend to quote or paraphrase the top-level finding from the abstract: "surgical face masks could prevent transmission of human coronaviruses and influenza viruses from symptomatic individuals." Anti-mask science communicators instead quote a sentence in the Discussion section, "our findings indicate that surgical masks can efficaciously reduce the emission of influenza virus particles into the environment in respiratory droplets, *but not in aerosols*" [emphasis added, for additional discussion of aerosols and COVID-19, see (63)]. In practice, the work of Leung *et al.* exists as two papers with two, seemingly contradictory conclusions: one that claims that surgical masks can prevent transmission

Table 1. A list of the top five most-shared citations (measured by number of shares of sources that cited them) in the Twitter dataset. We provide examples of positive and negative framings for each of the papers' results/conclusions.		
Paper citation	Pro-mask context	Anti-mask context
Leung <i>et al.</i> (62)	"It found that even loose-fitting surgical masks blocked almost all the contagious droplets the wearers breathed out and even also some infectious aerosols – tiny particles that can linger in the air." (90)	"Surgical masks can efficaciously reduce the emission of influenza virus particles into the environment in respiratory droplets, but not in aerosols." (55)
MacIntyre <i>et al.</i> (4)	"One of the most frequently mentioned papers evaluating the benefits and harms of cloth masks have been by MacIntyre <i>et al.</i> Findings have been misinterpreted, and therefore justify detailed discussion here... Its implementation does not inform the effect of using cloth masks versus not using masks in a community setting for source control of SARS-CoV-2, which is of the same genus as seasonal coronavirus, which has been found to be effectively filtered by cloth masks in a source control setting." (79)	[Transcribed from video] "Now here's what's even more shocking about the study. Not only were the rates of infection higher in the cloth mask group than they were in the medical mask group. The rates of infection in the cloth mask group were significantly higher than the people in the control group, who are allowed to occasionally wear a mask or not even wear a mask at all." (58)
Long <i>et al.</i> (64)	"Further, the evidence is mixed from randomized studies on types of masks and risk for influenza-like illness transmission to mask wearers; for example, a recent systematic review and meta-analysis comparing N-95 respirators versus surgical masks found a statistically insignificant decline in influenza risk with N-95 respirators." (91)	"There were no statistically significant differences in preventing laboratory-confirmed influenza, laboratory-confirmed respiratory viral infections, laboratory-confirmed respiratory infection, and influenza-like illness using N95 respirators and surgical masks." (56)
Radonovich <i>et al.</i> (92)	"Radonovich <i>et al.</i> found in an outpatient setting that use of N95 respirators, compared with medical masks... resulted in no significant difference in the rates of laboratory-confirmed influenza." (79)	[Transcribed from video] "In the 2019 study, there was no significant difference by wearing an N95 respirator mask." (57)
Smith <i>et al.</i> (93)	"If someone asks: What's the evidence for mask wearing? Here is a list of *SEVENTY* papers, including reviews/meta-analysis and individual studies, in reverse chronological order. Includes 31 from 2020 alone (!!). META = meta-analysis or systematic review." (54)	"This 2016 meta-analysis found that both randomized controlled trials and observational studies of N95 respirators and surgical masks used by healthcare workers did not show benefit against transmission of acute respiratory infections." (94)

of COVID-19 and another that claims that they cannot, given that COVID-19 is an aerosolized disease. We observe that science communicators tend to selectively quote or cite the paper to support their stance on masks.

Next, we look at the work of Long *et al.* (64), a paper that is cited positively by both popular pro- and anti-mask websites. Long *et al.* concluded from their literature review on mask wearing that “the use of N95 respirators compared with surgical masks is not associated with a lower risk of laboratory-confirmed influenza” and recommended that surgical masks may thus be a suitable substitute for N95 respirators. In three separate anti-mask websites, their results were quoted out of context as follows: “There were no statistically significant differences in preventing laboratory-confirmed influenza, laboratory-confirmed respiratory viral infections, laboratory-confirmed respiratory infection and influenza-like illness using N95 respirators and surgical masks.” Taken alone, this quotation can make it appear that neither surgical masks nor N95 respirators had any effectiveness for preventing disease, a finding that directly contradicts the main result of Long *et al.*, which is about the effectiveness of N95 masks compared to surgical masks. In this case, quotation is not so much selective as deceptive, implying that surgical masks are ineffective through the removal of additional context rather than that surgical masks are as effective as N95 masks and both of which are associated with lower risk of illness.

We conclude our findings by noting that the anti-mask paper group is not a marginal phenomenon in the realm of Twitter debates, and the websites that cite them are shared more often than those that share the pro-mask paper group. To quantify the difference in how often Twitter users are exposed to these two paper groups via shared URLs, we sum the shares for each cited

publication across all the science communicators citing them positively. According to this metric, nodes in the anti-mask cluster are shared positively more often than nodes in the pro-mask cluster (124,000 shares versus 76,000 shares). By this metric, some nodes in the pro-mask cluster have more negative exposure than positive exposure, such as that of Leffler *et al.* (60), which is not true for any node in the anti-mask cluster. We echo other researchers’ caution of “digital bias” (36), namely, that arguments offered in online forums are not likely to reflect public opinion writ large. Our chosen forum in this analysis, Twitter users arguing with each other in the wake of public health announcements by U.S. governors, likely attracts a unique population subset of politically engaged, internet-savvy commenters. Anti-mask viewpoints may have been specifically elicited by the choice of a highly politicized discussion forum. However, to the extent that U.S. governors’ official announcements on Twitter are meant to be a public and visible forum for discussing the pandemic, responses in the aggregate support the appearance of a misleading anti-mask consensus and oppose the mainstream pro-mask consensus.

DISCUSSION

We describe how science communicators cite the collected works of scientists differently to create the appearance of two radically different scientific consensuses: one supporting mask use during COVID-19 and one opposing. We find that science communicators perform a transformation of the original evidence bases used by scientists, which are mostly organized by fields and methodologies, into separate evidence bases segregated by stance toward masks. Science communicators, who, on one side, are from mostly

Downloaded from https://www.science.org on October 07, 2023

mainstream scientific and journalistic institutions and, on the other, are mostly a coalition of conspiracy theorists and creators on alternative media platforms, collect different scientific papers and through a process of aggregation (collection of positive citation) and denigration (negative citation) construct a perceived scientific consensus for their audiences. Some papers are referenced by both pro- and anti-mask science communicators, but rather than indicating shared scientific understanding, we see that these papers are either selectively or deceptively quoted to imply the opposite of their conclusions by anti-mask communicators. Overall, we portray the development of an anti-mask community of belief, which, far from being anti-science as their detractors may claim (23), puts concerted (and deceptive) effort into the boundary-work of separating favored “good science” from disfavored “bad science” (65). We also contribute a method for mapping the transformation, which occurs from the work published by scientists, so often the stopping point of scientometric analysis, into the consensus that the public perceives in online settings.

We find that negative citation is an important practice in the development of these two separate perceived consensus on Twitter and that anti-mask science communicators in our data use negative citation more frequently than other communicators. Science communication scholars have argued that prebunking, i.e., telling audiences that they will encounter misinformation, may be an effective way to communicate science and inoculate them against false beliefs (29–31). We find that, unfortunately, charismatic disinformers also inoculate their readers with critical citations of otherwise widely trusted research. We reflect that mainstream science communicators’ apparent reluctance to cite critically may stem from not wanting to attract attention, even critically, to counterconsensus findings. Although some research supports the benefits of prebunking (66), other findings have claimed that the “oxygen of amplification” may ultimately be a negative factor and that counterconsensus sources are best left uncovered by mainstream media (67, 68). Those outside the scientific mainstream, however, need not worry about amplifying their opponents, as they may assume that most people will inevitably encounter the mainstream position. We recommend further research on the circumstances in which it may or may not be advantageous to negatively cite information that audience members may not have previously encountered.

Although we study Twitter here, it is likely that the pattern that we find persists across other forms of media. A quotation from one of the most-shared papers in the anti-mask consensus that we identified, by MacIntyre *et al.* (4), was cited in 2021 on what was at the time the most popular broadcast television program in the United States, Tucker Carlson Tonight, and in a 2022 legal brief successfully advocating for the abolition of federal mask mandates on public transportation (5, 6). This is particularly unfortunate because Dr. MacIntyre herself has worked to support mask wearing in many subsequent letters and peer-reviewed publications and even specifically qualified the results of the 2015 paper in question (13). While experts in public health have the ability to assess MacIntyre’s 2015 paper holistically in light of her more recent work, non-experts, who encounter the 2015 paper in isolation on the journal’s website or aggregators such as PubMed, likely do not. This phenomenon has previously been called context collapse, where one group of users (non-experts) encounters information meant for a different group (experts) and consequently lacks the resources to correctly reason about it (26). A key qualitative result from this work is how

unprepared scientific publishers are for the implications of context collapse on their published archives.

To overcome context collapse, publishers could allow experts to append contextual information to papers that have been or are likely to be misused and do so in formats that are difficult for non-experts to ignore, i.e., appending a large, conspicuous warning box rather than the small ones that we find for these papers [e.g., in (69)]. Sites that rehost scientific content, such as PubMed and ResearchGate, must show a willingness to propagate such contextualizations, as the content in our study was often linked to these sites. Scientists themselves must show a willingness to create such contextualizations, and scientific institutions must create credit mechanisms that reward the creation of these contextualizations. Scientists’ peers must understand and acknowledge that rather than a show of weakness on the part of the original authors’ work, these additional comments represent valuable public service. Scientists, reviewers, and editors should also take care to craft their papers’ abstracts with specific statements that are not so easily misquoted, as paper abstracts are sometimes the only information available for lay readers without subscriptions to closed-access research. Such interventions can make the knowledge infrastructure of science more resilient to misuse by bad actors while preserving the value of that infrastructure to expert scientists. We encourage additional research into what forms such interventions may take and which are effective in defanging the work of counterconsensus science communicators (70).

The implication of this work that consensus are socially constructed by non-expert mediators may cause familiar anxiety among scientists, as it can suggest that mainstream scientific consensus is non-objective and thus potentially just as untrustworthy (71) as alternative consensus. We take the perspective of Oreskes (37) in her work *Why Trust Science* and believe that the selection of a certain consensus should depend on the diversity and strength of the vetting system that a consensus’s authors practice. The architects of the anti-mask consensus found here are largely an informal coalition of conspiracy theorists, creators on alternative media platforms, and others who have previously shared false theories about climate change, the benefits of vaccination, the existence of mass shootings, HIV/AIDS, abortion, the reality of the moon landing, and a variety of other topics. Few have any scientific background in mask effectiveness or even public health more generally, and similarly, few have any formal accountability to the public. By contrast, the architects of the pro-mask consensus are a coalition of mostly journalists, scientists, and government agencies, all of whom have established different formal systems of peer review and accountability. While it is widely known that these institutions do not always produce reliable knowledge, especially for matters concerning marginalized people (72–75), their systems for generating and vetting information are still stronger than those of a loose movement of unaffiliated conspiracy theorists and creators on alternative media platforms. As a community, scientists and publishers must work to build better, more trustworthy consensus by strengthening both the popular perception and the actual reality of a science informed by rigorous peer review that is performed by and accountable to a diverse array of knowledgeable stakeholders (76), and they must understand that their work is filtered through the lens of science communicators, both friendly and unfriendly, and participate in the process of correction and communication that leads to their work being interpreted as intended.

We conclude with some comments on the construction of the datasets used in this study, which has been undoubtedly time-consuming, and how this process may be automated in the future. While automated positive/negative citation coding would allow analyses of larger controversies over longer periods of time, our experience of manually coding citations suggests that automated processes will face difficulties. We find that scientists in archival work do not criticize papers the way people criticize, for example, restaurants or movies, and will often use nongeneralizable science- and field-specific language. Criticisms can be circumspect because of admirable comity norms in science, and their meaning may be diffused over several sentences. Outside of automatic coding, simply gathering the citations in a consistent manner is likely to be quite difficult. Websites communicating important scientific information can change their content almost daily, which we observed particularly with the CDC's websites, and this complicates attempts to understand the state of perceived scientific consensus at any given moment. Automatically collecting citations from websites with very different layouts and citation formats will also inevitably provide challenges, especially given how many of these websites do not hyperlink (or even consistently format) their citations. The popularity of video-based science communication, which combines audio content with images and text screenshots of scientific papers sometimes without titles, further increases the difficulty of completing an automatic census of communicated scientific literature as actually experienced by the media-consuming public. While we encourage research into how such data collection may be automated, we also encourage researchers to continue to engage in similar hand coding experiments in scientific communication to surface such challenges.

MATERIALS AND METHODS

Web of Science dataset

To understand how scientists were developing consensus around masks in the wake of the COVID-19 pandemic, we used the Web of Science paper database to collect mask-related papers in approximately the first year of the pandemic. Specifically, we collect 3593 articles from the Web of Science published between 28 March 2020 and 4 April 2021 containing in their title the following: "mask," "personal protective equipment," "ppe", and "n95." These choices were informed by familiarizing ourselves with popular papers referencing masks and the terms most frequently used to describe masks. We filter this dataset to the 82 papers that have at least 50 citations as of 18 November 2021, representing the top 2% of papers retrieved. We choose this threshold to limit the labor needed for subsequent manual coding. We exclude two papers unrelated to masking during the COVID-19 pandemic and three papers with no collectable citations, leaving a dataset of 77 papers. These papers are listed in the Supplementary Materials.

Most papers are peer-reviewed experiments and literature reviews; others are non-peer-reviewed editorials. Most directly pertain to measuring the effectiveness of masks for preventing the spread of COVID-19, but some pertain to adjacent issues such as the effectiveness of reusing masks or mitigating the waste products created from COVID-19-related mask production.

For the 77 papers, we retrieve all listed outgoing citations within the papers that are archived by the Web of Science using the digital object identifier (DOI). While this does not include all citations, our

objective is to replicate those data curation practices typically performed in large-scale bibliometric studies of science. Even including citations without DOIs, the Web of Science does not include every citation that these articles make, especially when citations are made to work published outside of academic journals. There are rare instances in which the Web of Science consistently misrouted a given paper's citation to the wrong paper, and we attempted to correct these when identified. This process results in 2426 citations, with the most cited paper as Leung *et al.* (62) referenced in 29 papers.

Twitter Science Dataset

Most social media posts do not necessarily constitute scientific debate, which poses a challenge from a data collection perspective. A naive collection of Twitter posts containing the word "mask" during our study period surfaces mostly content related to sensational stories about masks. Such links rarely contain scientific content and, instead, document, for example, an altercation in a gas station started over a request to wear a mask. To discover users' content about the effectiveness of masks, we instead specifically sample reply threads between users discussing masks, often in disagreement with one other. These arguments stimulate the trading of evidence between users in the form of external URLs, which then become the object of analysis for the Twitter Science Dataset used here.

We choose as our forum for arguments posts made in reply to the set of all Twitter accounts for U.S. governors (and the mayor of the District of Columbia) as of 28 March 2020. These political leaders are often tasked with announcing and explaining public health regulations that require wearing masks during the COVID-19 pandemic, and their posts about masks tend to generate fierce arguments between pro- and anti-mask users. From 28 March 2020 to 4 April 2021, we collect all hierarchical reply trees to these politicians that contain the word "mask" at any point in the tree. The number of posts in all reply trees totaled 37.8 million, and the subset of trees containing an instance of "mask" totaled 5.1 million posts (13% of the total). From these posts, we extracted all posts that shared a URL for a total of 158,000 unique URLs shared across 351,000 posts. We collected Twitter posts using the Twitter streaming API and independently unwound redirected links to retrieve the final destination of original URLs. Some URLs could not be retrieved because of redirect URLs no longer being operational.

From this dataset of URLs, we first remove marketing tags and other character strings appended to the end of URLs that do not affect the final landing page and then deduplicate the results. This reduces the URL dataset size by 10% to 142,000 unique URLs. To find a tractable sample for manual coding, we subsample the 194 URLs that were shared more than 100 times, which collectively are responsible for 15% of all URLs shared. Two researchers then perform a qualitative coding process on the linked content to identify which are explicitly concerned with the effectiveness (or ineffectiveness) of mask wearing for COVID-19. This restricted set of links comprises 98 unique URLs. We combined unique URLs that rehosted identical content on different websites, for a final set of 81 unique pieces of science communication. Many URLs had either been modified or, in the case of social media posts, deleted by the time of coding. To address these gaps, we use the Internet Archive (77) to access the most recent nonredirected version as of 4 April

2021. In only one case could a URL, linking to a removed Twitter post, not be retrieved via this method. We note that 8 of 10 of the most shared unique URLs specifically concerned themselves with a scientific analysis of the effectiveness of masks, with the remaining nonscientific links corresponding to a post by former U.S. president Donald Trump and a petition to recall California Governor Gavin Newsom.

From this dataset of 81 URLs, two researchers manually recorded every reference in the main text of these URLs. These references could be represented as embedded URLs, appended as formal academic citations, screenshots of text, or explicitly referenced by article name. In the case of URLs leading to videos, researchers watched the videos and recorded verbal citations and citations displayed in text on the video itself, a frequent behavior. All citations to external sources were recorded, including material published in scientific journals and sources such as videos, scientific journalism, and other websites. We coded the most recent version of the website as of 4 April 2021 and only content on the site that directly pertain to masks and the article in question. This resulted in 1496 citations, with the most-cited article (as in the Web of Science dataset) being that of Leung *et al.* (62), cited by 19 unique URLs.

Annotating citation and website valence

In both the Web of Science and the Twitter citations datasets, the lead researcher manually annotated citations for valence, specifically noting negative citations. We define two criteria for a negative citation: (i) criticizing the methods, conclusions, or statistical power of a cited study [e.g. "Many studies in these reviews were underpowered, and most failed to measure adherence..." (78)] and (ii) arguing that a cited study has been popularly misinterpreted by peers or the public, either in general or for the specific case in context [e.g., "One of the most frequently mentioned, but misinterpreted, papers evaluating cloth masks..." (79)].

We do not label citations negatively when authors raise limitations or areas for improvement in otherwise trustworthy work or in the case where cited findings conflict with the findings of the present paper but are not themselves criticized. Our intended definition for negative citations is that, in effect, the citing article is asking readers to at least partially disregard the conclusions of a cited paper.

For our analysis of negative citation rates in Twitter science communicators, we also labeled websites in our dataset as either having a stance generally critical of masks as intervention for COVID-19 or generally positive or neutral. We considered websites to have an anti-mask orientation if they either explicitly advocated against the wearing of masks or primarily reported scientific findings supporting the ineffectiveness of masks. Anti-mask websites that had been publicly corrected and modified by the time of our analysis were considered to have neutral stance. We labeled 40 websites as critical of masks and 41 as positive or neutral.

Network construction

Cocitation graphs are essentially undirected unipartite projections of directed bipartite graphs, where an edge between two nodes signifies that two papers are cited together in at least one citing article (80). Projecting a bipartite network into a unipartite network is never trivial, and choices about how to represent edges and edge weights can significantly affect resulting analyses, particularly clusterings (81). These choices are made even more complex in the

signed case that we introduce here. We perceive three principal researcher choices with respect to these projections: (i) the choice to analyze a unipartite projection rather than the original bipartite network, (ii) the scheme for generating edges and edge weights for the unipartite projection, and (iii) the scheme for filtering edges such that communities of interest are preserved.

The first is the decision to project a bipartite graph, especially when the original graph could be considered unipartite given that some sources cite other sources in the citing set. Our choice to treat "citing" and "cited" entities as separate and, thus, bipartite stems from our understanding of this data as representing the curatorial choices of a source's authors rather than a deep engagement with a source's content. We find that many anti-mask communicators misleadingly cite research that actually endorses masks and itself cites extensively pro-mask literature. In directed unipartite form, this would lead to anti-mask researchers seemingly endorsing a wide variety of pro-mask literature via one transitive connection.

The second choice is how to define edges in a signed projected graph. We follow a method described in (82), where positive edges are considered to be (+1) and negative edges are considered to be (−1), and the weight contribution to an edge cited by two nodes being the product of their cited edges ($1 \times 1 = 1$, $-1 \times -1 = 1$, $1 \times -1 = -1$). Nodes are either together within trustworthy science (1×1), together within untrustworthy pseudoscience (-1×-1), or apart from each other on either side of the boundary (1×-1). In subsequent analysis, we analyze two such projections: one that contains all of the summed positive edge contributions (nodes cited with the same valence) and one that contains all of the negative edge contributions (nodes cited with opposite valence) over the same set of nodes. Splitting these projections, rather than summing them into one graph, aids in cluster detection via methods that optimize partitionings between multigraphs, as described in the following section.

The third choice is how to weight and filter the edges to only those relevant to the clusters at hand. As described in (81), projections tend to generate local structures that are highly amenable to clustering via modularity-based methods because each citing node creates a densely connected subgraph of all the nodes that it cites in the projection. Accordingly, modularity-based methods can easily return clusters that simply represent "all the papers cited by a single source." While weighting methods have been proposed that mitigate the cluster-making ability of single sources, such as hyperbolic weight projection, we find their application to signed citation networks nontrivial and not necessarily well suited to the relatively normal degree distribution found in our data. We instead opt for a simple filtering mechanism: Connections between two papers are not considered unless they have been cited together in at least three sources in our datasets. This eliminates the cluster-making ability of single authors and single authors who create two sources and use similar citation lists. Our data does not feature authors with more than two sources, making stricter filters unnecessary. We note that this is not a simple filter on edge weights less than 3 on both subgraphs, as two sources could be cited positively together in two sources (+2) and negatively together in one (−1), and these edges would still be preserved in their respective projections because they represent cocitation in three sources.

Clustering and visualization

Our primary clustering algorithm is based on the modularity statistic (83) and optimizes modularity using a multigraph method derived from the Leiden algorithm (84) as implemented in the Python package (85). Specifically, this optimization method attempts to maximize modularity on the graph representing positive ties while simultaneously minimizing modularity on the graph representing negative ties.

Historically, cocitation analyses use large, automatically collected datasets with hundreds to millions of nodes and edges. Our case is unique in that, because of hand coding, our network contains only hundreds of nodes and thousands of connections. The small size of our dataset increases the influence of noise in our data and, correspondingly, the level of uncertainty in clusterings and network representations. To address this challenge, we use a method akin to soft clustering, where individual nodes are assigned cluster probabilities rather than strict assignments, and inspired by the clustering bootstrapping procedure proposed in (86). Specifically, we resample each weight in the projected graphs according to a normal distribution with mean of the given weight and variance equal to the variance of all weights in the graph. We perform such replications and cluster the results 10,000 times for each graph, permuting the order of nodes provided to algorithm each time. We then create a derived representation of these graphs where each node represents a cited paper, and links between nodes are weighted by the percentage (0 to 1.0) of times these nodes were clustered in the same graph. We then apply clustering using an unsigned modularity method to this derived graph. We visualize these representations in Fig. 1, with a filter such that nodes that are clustered less than 50% of the time are removed.

We visualize graphs using the ForceAtlas2 algorithm in the network visualization software Gephi (87, 88). Nodes in Fig. 1 are sized by the summed number of citations that the papers citing them have received to illustrate their importance in scientific discourse around masks in the first year of the pandemic. Nodes in Fig. 1 are similarly sized by the summed number of times the websites citing them have been shared on Twitter, for a parallel measure of importance.

Supplementary Materials

This PDF file includes:

Table S1

REFERENCES AND NOTES

1. S. Lewandowsky, G. E. Gignac, S. Vaughan, The pivotal role of perceived scientific consensus in acceptance of science. *Nat. Clim. Change* **3**, 399–404 (2013).
2. D. A. Scheufele, N. M. Krause, Science audiences, misinformation, and fake news. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 7662–7669 (2019).
3. N. Oreskes, E. M. Conway, *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming* (Bloomsbury, 2012).
4. C. R. MacIntyre, H. Seale, T. C. Dung, N. T. Hien, P. T. Nga, A. A. Chughtai, B. Rahman, D. E. Dwyer, Q. Wang, A cluster randomised trial of cloth masks compared with medical masks in healthcare workers. *BMJ Open* **5**, e006577 (2015).
5. Tucker Carlson Tonight : FOXNEWSW : July 1, 2021 10:00pm-11:00pm PDT (2021). Num Pages: 3657.
6. P. Hitlin, Health issues topped the list of scientific studies reaching wide audiences in 2016 (Pew Research Center, 2016).
7. C. Raina MacIntyre, V. Costantino, A. Chanmugam, The use of face masks during vaccine roll-out in New YorkCity and impact on epidemic control. *Vaccine* **39**, 6296–6301 (2021).
8. C. R. MacIntyre, P. Y. Nguyen, A. A. Chughtai, M. Trent, B. Gerber, K. Steinhofel, H. Seale, Mask use, risk-mitigation behaviours and pandemic fatigue during the COVID-19 pandemic in five cities in Australia, the UK and USA: A cross-sectional survey. *Int. J. Infect. Dis.* **106**, 199–207 (2021).
9. C. R. MacIntyre, N. Binkin, In the room where it happens: The consequences of the lack of public health expertise during the COVID-19 pandemic. *Glob. Biosecurity* **3**, (2021).
10. C. R. MacIntyre, V. Costantino, L. Bian, C. Bethel, Effectiveness of facemasks for opening a university campus in Mississippi, United States - a modelling study. *J. Am. Coll. Health* **70**, 2505–2510 (2022).
11. C. R. MacIntyre, A. A. Chughtai, Masks in the community are an effective strategy: Author's response to Haslam et al (2020). *Int. J. Nurs. Stud.* **111**, 103751 (2020).
12. A. A. Chughtai, H. Seale, C. R. MacIntyre, Effectiveness of cloth masks for protection against severe acute respiratory syndrome coronavirus 2. *Emerg. Infect. Dis.* **26**, (2020).
13. C. R. MacIntyre, T. C. Dung, A. A. Chughtai, H. Seale, B. Rahman, Contamination and washing of cloth masks and risk of infection among hospital health workers in Vietnam: A post hoc analysis of a randomised controlled trial. *BMJ Open* **10**, e042045 (2020).
14. C. R. MacIntyre, A. A. Chughtai, A rapid systematic review of the efficacy of face masks and respirators against coronaviruses and other respiratory transmissible viruses for the community, healthcare workers and sick patients. *Int. J. Nurs. Stud.* **108**, 103629 (2020).
15. Y. Wang, H. Tian, L. Zhang, M. Zhang, D. Guo, W. Wu, X. Zhang, G. L. Kan, L. Jia, D. Huo, B. Liu, X. Wang, Y. Sun, Q. Wang, P. Yang, C. R. MacIntyre, Reduction of secondary transmission of SARS-CoV-2 in households by face mask use, disinfection and social distancing: A cohort study in Beijing, China. *BMJ Glob. Health* **5**, e002794 (2020).
16. A. A. Chughtai, H. Seale, M. S. Islam, M. Owais, C. R. MacIntyre, Policies on the use of respiratory protection for hospital health workers to protect from coronavirus disease (COVID-19). *Int. J. Nurs. Stud.* **105**, 103567 (2020).
17. M. Brüggemann, S. Engesser, Beyond false balance: How interpretive journalism shapes media coverage of climate change. *Glob. Environ. Chang.* **42**, 58–67 (2017).
18. N. M. Lee, M. S. VanDyke, Set it and forget it. *Sci. Commun.* **37**, 533–541 (2015).
19. V. Alatas, A. G. Chandrasekhar, M. Mobius, B. A. Olken, C. Paladines, "When celebrities speak: A nationwide twitter experiment promoting vaccination in indonesia, Working Paper, National Bureau of Economic Research, 2019.
20. J. Allgaier, Science and environmental communication on youtube: Strategically distorted communications in online videos on climate change and climate engineering. *Front. Commun.* **4**, 36 (2019).
21. C. Lee, T. Yang, G. D. Inchoco, G. M. Jones, A. Satyanarayan, Viral visualizations: How coronavirus skeptics use orthodox data practices to promote unorthodox science online, paper presented at the Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Yokohama Japan, 8 to 13 May 2021, pp. 1–18.
22. S. Lewandowsky, J. Cook, N. Fay, G. E. Gignac, Science by social media: Attitudes towards climate change are mediated by perceived social consensus. *Mem. Cognit.* **47**, 1445–1456 (2019).
23. A. Philipp-Muller, S. W. S. Lee, R. E. Petty, Why are people antisience, and what can we do about it? *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2120755119 (2022).
24. N. d. Tyson, "Neil deGrasse Tyson on the pandemic year: Science needs better marketing," *Wall Street Journal*, 2021.
25. K. S. Koltai, K. R. Fleischmann, Questioning science with science: The evolution of the vaccine safety movement. *Proc. Assoc. Inf. Sci. Technol.* **54**, 232–240 (2017).
26. A. E. Marwick, d. boyd, I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media Soc.* **13**, 114–133 (2011).
27. I. V. Pasquetto, A. F. Olivieri, L. Tacchetti, G. Riotta, A. Spada, Disinformation as infrastructure: Making and maintaining the QAnon conspiracy on Italian digital media. *Proc. ACM Hum.-Comput. Interact.* **6**, 1–31 (2022).
28. C. Catalini, N. Lacetera, A. Oettl, The incidence and role of negative citations in science. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 13823–13826 (2015).
29. S. Lewandowsky, S. van der Linden, Countering misinformation and fake news through inoculation and prebunking. *Eur. Rev. Soc. Psychol.* **32**, 348–384 (2021).
30. M. Basol, J. Roozenbeek, M. Berriche, F. Uenal, W. P. McClanahan, S. Linden, Towards psychological herd immunity: Cross-cultural evidence for two prebunking interventions against COVID-19 misinformation. *Big Data Soc.* **8**, 205395172110138 (2021).
31. S. van der Linden, J. Roozenbeek, J. Compton, Inoculating against fake news about COVID-19. *Front. Psychol.* **11**, 2928 (2020).
32. I. Boutron, P. Ravaud, Misrepresentation and distortion of research in biomedical literature. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 2613–2619 (2018).
33. A. van Stekelenburg, G. Schaap, H. Veling, J. v.'t. Riet, M. Buijzen, Scientific-consensus communication about contested science: A preregistered meta-analysis. *Psychol. Sci.* **33**, 1989–2008 (2022).
34. S. Chinn, D. S. Lane, P. S. Hart, In consensus we trust? Persuasive effects of scientific consensus communication. *Public Underst. Sci.* **27**, 807–823 (2018).

35. D. M. Kahan, H. Jenkins-Smith, D. Braman, Cultural cognition of scientific consensus. *J. Risk Res.* **14**, 147–174 (2011).
36. N. Marres, Why map issues? On controversy analysis as a digital method. *Sci. Technol. Hum. Values* **40**, 655–686 (2015).
37. N. Oreskes, *Why Trust Science?* (Princeton Univ. Press, 2021).
38. B. Miller, *The Routledge Handbook of Social Epistemology* (Routledge, 2019), p. 10.
39. J. Bruggeman, V. A. Traag, J. Uitermark, Detecting communities through network data. *Am. Sociol. Rev.* **77**, 1050–1063 (2012).
40. D. Edge, Quantitative measures of communication in science: A critical review. *Hist. Sci.* **17**, 102 (1979).
41. J. Mingers, L. Leydesdorff, A review of theory and practice in scientometrics. *Eur. J. Oper. Res.* **246**, 1–19 (2015).
42. S. Fortunato, C. T. Bergstrom, K. Börner, J. A. Evans, D. Helbing, S. Milojević, A. M. Petersen, F. Radicchi, R. Sinatra, B. Uzzi, A. Vespignani, L. Waltman, D. Wang, A. L. Barabási, Science of science. *Science* **359**, eaao185 (2018).
43. P. N. Edwards, S. J. Jackson, M. K. Chalmers, G. C. Bowker, C. L. Borgman, D. Ribes, M. Burton, S. Calvert, “Knowledge infrastructures: Intellectual frameworks and research challenges,” Working paper, Deep Blue, 2013; <http://hdl.handle.net/2027.42/97552>.
44. B. Fischhoff, D. A. Scheufele, The science of science communication. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 14031–14032 (2013).
45. Clarivate Analytics, Web of science core collection. *Citation database. Web of Science* (2017).
46. C. M. Trujillo, C. M. Long, Document co-citation analysis to enhance transdisciplinary research. *Sci. Adv.* **4**, e1701130 (2018).
47. D. O. Case, G. M. Higgins, How can we investigate citation behavior? A study of reasons for citing literature in communication. *J. Am. Soc. Inf. Sci.* **51**, 635–645 (2000).
48. M. J. Moravcsik, P. Murugesan, Some results on the function and quality of citations. *Soc. Stud. Sci.* **5**, 86–92 (1975).
49. D. K. Chu, E. A. Akl, S. Duda, K. Solo, S. Yaacoub, H. J. Schünemann, D. K. Chu, E. A. Akl, A. el-harakeh, A. Bognanni, T. Lotfi, M. Loeb, A. Hajizadeh, A. Bak, A. Izcovich, C. A. Cuello-Garcia, C. Chen, D. J. Harris, E. Borowiack, F. Chamseddine, F. Schünemann, G. P. Morgano, G. E. U. Muti Schünemann, G. Chen, H. Zhao, I. Neumann, J. Chan, J. Khabisa, L. Hneiny, L. Harrison, M. Smith, N. Rizk, P. Giorgi Rossi, P. AbiHanna, R. el-khoury, R. Stalteri, T. Baldeh, T. Piggott, Y. Zhang, Z. Saad, A. Khamis, M. Reinap, S. Duda, K. Solo, S. Yaacoub, H. J. Schünemann, Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: A systematic review and meta-analysis. *The Lancet* **395**, 1973–1987 (2020).
50. P. Jüni, B. R. d. Costa, P. Bobos, N. S. Bodmer, A. McGeer, Revisiting the evidence for physical distancing, face masks, and eye protection. *Lancet* **398**, 663 (2021).
51. N. Bai, *Still Confused About Masks? Here's the Science Behind How Face Masks Prevent Coronavirus* (UCSF, 2020).
52. Center for Disease Control Prevention, Masks (2022).
53. J. Hanson, How Well Do Masks Work? Schlieren Imaging In Slow Motion! (2020).
54. @nuanceORDEATH, #MASKUP THREAD (2020).
55. M. M. Singleton, Mask Facts (2020).
56. D. G. Rancourt, Masks Don't Work: A Review of Science Relevant to COVID-19 Social Policy (2020).
57. B. Swann, Why Face Masks DON'T Work, According To SCIENCE (2020).
58. S. Stevenson, Mask Facts (2021).
59. @jcho710, “An N-95 mask filters...” (2020).
60. C. T. Leffler, E. Ing, J. D. Lykins, M. C. Hogan, C. A. M. Keown, A. Grzybowski, Association of country-wide coronavirus mortality with demographics, testing, lockdowns, and public wearing of masks. *Am. J. Trop.* **103**, 2400–2411 (2020).
61. J. L. Jacobs, S. Ohde, O. Takahashi, Y. Tokuda, F. Omata, T. Fukui, Use of surgical face masks to reduce the incidence of the common cold among health care workers in Japan: A randomized controlled trial. *Am. J. Infect. Control* **37**, 417–419 (2009).
62. N. H. L. Leung, D. K. W. Chu, E. Y. C. Shiu, K. H. Chan, J. J. McDevitt, B. J. P. Hau, H. L. Yen, Y. Li, D. K. M. Ip, J. S. M. Peiris, W. H. Seto, G. M. Leung, D. K. Milton, B. J. Cowling, Respiratory virus shedding in exhaled breath and efficacy of face masks. *Nat. Med.* **26**, 676–680 (2020).
63. K. Randall, E. T. Ewing, L. C. Marr, J. L. Jimenez, L. Bourouiba, How did we get here: What are droplets and aerosols and how far do they go? A historical perspective on the transmission of respiratory infectious diseases. *Royal Soc.* **11**, 20210049 (2021).
64. Y. Long, T. Hu, L. Liu, R. Chen, Q. Guo, L. Yang, Y. Cheng, J. Huang, L. du, Effectiveness of N95 respirators versus surgical masks against influenza: A systematic review and meta-analysis. *J. Evid. Based Med.* **13**, 93–101 (2020).
65. T. F. Gieryn, Boundary-work and the demarcation of science from non-science: Strains and interests in professional ideologies of scientists. *Am. Sociol. Rev.* **48**, 781 (1983).
66. U. K. H. Ecker, S. Lewandowsky, M. Chadwick, Can corrections spread misinformation to new audiences? Testing for the elusive familiarity backfire effect. *Cogn. Res.* **5**, 41 (2020).
67. W. Phillips, The oxygen of amplification. *Data Soc.* **22**, 1 (2018).
68. I. Skurnik, C. Yoon, D. C. Park, N. Schwarz, How warnings about false claims become recommendations. *J. Consum. Res.* **31**, 713–724 (2005).
69. M. Klompas, C. A. Morris, J. Sinclair, M. Pearson, E. S. Shenoy, Universal masking in hospitals in the covid-19 era. *N. Engl. J. Med.* **382**, e63 (2020).
70. C. E. Smith, E. Nevarez, H. Zhu, in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20* (Association for Computing Machinery, 2020), pp. 1–13.
71. A. Kofman, Bruno Latour, “The post-truth philosopher, mounts a defense of science,” *The New York Times*, 2018.
72. K. TallBear, *Native American DNA: Tribal Belonging and the False Promise of Genetic Science* (University of Minnesota Press, 2013).
73. S. M. Shuster, *Trans Medicine: The Emergence and Practice of Treating Gender* (NYU Press, 2021).
74. M. Dusenbery, *Doing Harm: The Truth About How Bad Medicine and Lazy Science Leave Women Dismissed, Misdiagnosed, and Sick* (HarperCollins, 2018).
75. R. A. Hogarth, *Medicalizing Blackness: Making Racial Difference in the Atlantic World, 1780–1840* (UNC Press Books, 2017).
76. S. T. Fiske, C. Dupree, Gaining trust as well as respect in communicating to motivated audiences about science topics. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 13593–13597 (2014).
77. B. Kahle, Preserving the internet. *Sci. Am.* **276**, 82–83 (1997).
78. B. Javid, M. P. Weekes, N. J. Matheson, Covid-19: Should the public wear face masks? *BMJ* **369**, m1442 (2020).
79. J. Howard, A. Huang, Z. Li, Z. Tufekci, V. Dzimal, H.-M. van der Westhuizen, A. von Delft, A. Price, L. Fridman, L.-H. Tang, V. Tang, G. L. Watson, C. E. Bax, R. Shaikh, F. Questier, D. Hernandez, L. F. Chu, C. M. Ramirez, A. W. Rimoin, Face masks against COVID-19: An evidence review (Preprints.org, 2020).
80. H. Small, Co-citation in the scientific literature: A new measure of the relationship between two documents. *J. Am. Soc. Inf. Sci.* **24**, 265–269 (1973).
81. T. J. B. Cann, I. S. Weaver, H. T. P. Williams, Is it correct to project and detect? How weighting unipartite projections influences community detection. *Netw. Sci.* **8**, S145–S163 (2020).
82. T. Derr, C. Johnson, Y. Chang, J. Tang, Balance in signed bipartite networks, in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (ACM, 2019)*, pp. 1221–1230.
83. M. E. J. Newman, M. Girvan, Finding and evaluating community structure in networks. *Phys. Rev. E* **69**, 026113 (2004).
84. V. A. Traag, L. Waltman, N. J. van Eck, From Louvain to Leiden: Guaranteeing well-connected communities. *Sci. Rep.* **9**, 5233 (2019).
85. V. A. Traag, J. Bruggeman, Community detection in networks with positive and negative links. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **80**, 036115 (2009).
86. M. Rosvall, C. T. Bergstrom, Mapping change in large networks. *PLOS ONE* **5**, e8694 (2010).
87. M. Bastian, S. Heymann, M. Jacomy, Gephi: An open source software for exploring and manipulating networks. *Proceedings of the International AAAI Conference on Web and Social Media* **3**, 361–362 (2009).
88. M. Jacomy, T. Venturini, S. Heymann, M. Bastian, ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software. *PLOS ONE* **9**, e98679 (2014).
89. S. Rengasamy, B. Eimer, R. E. Shaffer, Simple respiratory protection—evaluation of the filtration performance of cloth masks and common fabric materials against 20–1000 nm size particles. *Ann. Occup. Hyg.* **54**, 789–798 (2010).
90. M. Godoy, Yes, Wearing Masks Helps. Here's Why (2020).
91. W. Lyu, G. L. Wehby, Community use of face masks and COVID-19: Evidence from a natural experiment of state mandates in the US. *Health Aff. (Millwood)* **39**, 1419–1425 (2020).
92. L. J. Radonovich Jr., M. S. Simberloff, M. T. Bessesen, A. C. Brown, D. A. T. Cummings, C. A. Gaydos, J. G. Los, A. E. Krosche, C. L. Gilbert, G. J. Gorse, A.-C. Nyquist, N. G. Reich, M. C. Rodriguez-Barradas, C. S. Price, T. M. Perl, ResPECT investigators, N95 respirators vs medical masks for preventing influenza among health care personnel: A randomized clinical trial. *JAMA* **322**, 824–833 (2019).
93. J. D. Smith, C. C. MacDougall, J. Johnstone, R. A. Copes, B. Schwartz, G. E. Garber, Effectiveness of N95 respirators versus surgical masks in protecting health care workers from acute respiratory infection: A systematic review and meta-analysis. *CMAJ* **188**, 567–574 (2016).
94. C. Huber, Masks are neither effective nor safe (2020).

Acknowledgments: We thank J. S. Schafer, J. Liu, M. Mayanja, M. Ionescu, M. Sioson, and P. Lockaby for advice and support on this work. **Funding:** 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), the William and Flora Hewlett Foundation (2019-9221), and the University of Washington Population Health Initiative Pilot Grants. We also are grateful for support from the NSF (grants 2027792 and 1749815). **Author contributions:**

Conceptualization: A.B. and S.N. Methodology: A.B., J.D.W., and E.S.S. Investigation: A.B. and S.N. Visualization: A.B. Supervision: K.S., J.D.W., and E.S.S. Writing (original draft): A.B. Writing (review and editing): A.B., S.N., K.S., J.D.W., and E.S.S. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** The code for reproducing this paper, Twitter URL dataset, Web of Science citation dataset, annotation datasets, and network visualization data can be accessed online via Zenodo: <https://zenodo.org/record/>

8034125. All data needed to evaluate the conclusions of the paper are present in the paper and/or the Supplementary Materials.

Submitted 1 March 2023

Accepted 18 August 2023

Published 22 September 2023

10.1126/sciadv.adh1933

Selective and deceptive citation in the construction of dueling consensus

Andrew Beers, Sarah Nguyen, Kate Starbird, Jevin D. West, and Emma S. Spiro

Sci. Adv. **9** (38), eadh1933. DOI: 10.1126/sciadv.adh1933

View the article online

<https://www.science.org/doi/10.1126/sciadv.adh1933>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2023 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution License 4.0 (CC BY).