


JACK STILGOE

Viral Suppression



When Facebook’s fact-checkers slapped a “missing context” label on a venerable medical journal’s article about breached vaccine trial protocols, they set off a very twenty-first-century fight about who should play what role in scientific communication.

In November 2021, an investigative story in the *British Medical Journal (BMJ)* with the headline “Covid-19: Researcher blows the whistle on data integrity issues in Pfizer’s vaccine trial” received more than 3.5 million hits. Altmetric, a company that measures scientific papers’ online attention, scored the article as the third most-shared research output of all time.

It also soon fell foul of mechanisms Facebook had put in place to stop misinformation. Those attempting to share the article received a message saying that independent fact-checkers had found it was “missing context.” Those running Facebook groups whose members shared the article were told there was “partly false information in your group.”

BMJ responded with fury. It posted an open letter to Mark Zuckerberg, cofounder of Facebook and CEO of its parent company, Meta, calling the fact check “inaccurate, incompetent, and irresponsible.” The 182-year-old journal objected to being characterized as a news blog, being included on a URL address that contained the phrase “hoax-alert,” and having a stamp marked “flaws reviewed” on a screenshot of the article. It emphasized that the article contained no outright errors and accused Facebook’s fact-checker, a third-party company called Lead Stories, of unjustified censorship.

Lead Stories responded that *BMJ*’s article had a “scare headline” with analysis that “oversells the whistleblower and overstates the jeopardy,” which resulted in “hundreds of Facebook posts and tweets, many by anti-vaccine activists using it as ‘proof’ the entire clinical trial was fraudulent and the vaccine unsafe.” The response also noted that Facebook had not restricted traffic or visibility but “merely warned of missing context.” That label, according to Facebook’s guidelines, is applied to content that may mislead or “content that implies a false claim without directly stating it.”

Accusations and name-calling pinged and poned over

Twitter. Lead Stories tweeted at the journalist who wrote the article, Paul Thacker, asking whether he minded being republished by the “Disinformation Dozen,” a small group identified by the nonprofit Center for Countering Digital Hate as spreading about two-thirds of the COVID-19 misinformation on social media. Thacker addressed Lead Stories as “You moron” and accused its team of trolling.

As a professor of science and technology policy at University College London, part of my job is studying public debates about science. I was intrigued that an esteemed journal like *BMJ* had been called out this way. I posted on Twitter that it would be a fascinating case study. Lead Stories responded: “When a reputable medical journal is so good at science communication that its article is republished verbatim by anti-vaccine activists because it can be easily misinterpreted to mean something it doesn’t, yes, that is quite the case study.”

The deeper I dug into this case, the more complicated it became. But ultimately, it boils down to competing assumptions about healthy information ecosystems and the role organizations such as *BMJ* and Facebook (and the fact-checkers they deploy) should play within these ecosystems. By focusing on the *BMJ* story’s heavy traffic and its use in narratives that went beyond the article’s explicitly stated claims, Lead Stories was, first, trying to control how the story would be used by people they saw as bad actors and, second, trying to control the spread of the story. By contrast, *BMJ* sought to ask important questions about shoddy oversight of clinical trials. *BMJ* did not address the issue of readers’ overinterpretation or promotion by anti-vax websites. Rather, it pointed to its own august reputation, the story’s lack of outright factual errors, and an (anonymous) external peer review as exempting it from Facebook’s labeling.

Although *BMJ*’s articles generally come from academics and clinicians, the November piece was a feature story written

by an investigative reporter. But conventional research articles have also been caught in social media platforms' fact-checking apparatus. One example of this involves Cochrane, a well-respected British nongovernmental organization that conducts rigorous systematic reviews of clinical treatments by collecting published articles, evaluating their quality, and publishing an overarching assessment of the state of the evidence. Around the same time *BMJ* published its article, Cochrane complained that Instagram users mentioning Cochrane's work received warnings that the organization had "repeatedly posted content that goes against our Community Guidelines on false content about COVID-19 or vaccines." After several weeks, prohibitions were lifted, though without any explanation I could find. Like Facebook, Instagram is owned by Meta, but even for those inside the process, the lines between what is deemed factually acceptable and what is not appear to be blurry: in its response to *BMJ*'s open letter, which mentioned Instagram's treatment of Cochrane, Lead Stories said that it had no role in this decision and added that Cochrane's content should "probably not be blocked."

What these tangles with reputation, misinformation, social media, and polarization show is that the players and the power in science communication are shifting. For most of science's history there has been a presumption that communication should be one-way (from the experts to the public) and tightly controlled (by the experts). More recently, there has been a democratic shift toward open dialogue, something that social media platforms can sometimes facilitate. But in practice, social media can also boost misinformation, contributing to political polarization and threatening public health. The pandemic has raised the stakes of information quality and distribution while also increasing the number and types of people paying attention to scientific publishing. The roles that media, the public, government, and scientific publishers play in disseminating, validating, and debating content are all in flux.

Disorganized skepticism

Robert Merton, a leading light of twentieth-century sociology, argued in the 1940s that science derives its authority in part from its "organized skepticism." Science, he contended, progressed through constant scrutiny, but this scrutiny happened within its own citadel. Mechanisms of peer review and "invisible colleges" such as the United Kingdom's Royal Society were seen as guarantees. Asking questions was encouraged, but only for those with institutionalized authority. The number of times an article was cited by other scientists was considered a proxy for the importance of a piece of research (and the researchers who conducted it).

Online access has changed this radically, providing quick, open forums for critique. The amount of attention scientific articles attract is being measured and valued in new ways—ones that many journals and researchers actively seek to maximize. Despite the fact that social media prioritizes

engagement over quality, tools like Altmetrics are increasingly used to measure and celebrate such attention. Although *BMJ*'s target audience is practicing medical doctors, the journal decided to provide its reporting on COVID-19 to the public for free, which doubled their website traffic during the pandemic. Of the 100 scientific articles with the highest Altmetric score in 2021, 98 were about COVID-19, including papers on ivermectin and hydroxychloroquine, two drugs that repeatedly failed to demonstrate efficacy yet received cult status among vaccine sceptics. Like *BMJ*, many journals have come to value high click rates and shares by nonspecialists, although I have seen little reflection on the changing role of these publications in a social media world.

The twenty-first century brought recognition that past notions of one-way science communication—also known as the "knowledge deficit model"—were both simplistic and paternalistic. Effective science communication requires dialogue with the public, not heavy-handed preaching. Clinicians now increasingly ask patients' opinions when recommending treatments, and policymakers acknowledge the need to listen to citizens' attitudes on new technologies such as synthetic biology or artificial intelligence.

Meanwhile, the internet has radically democratized access to information and the tools of research, completely transforming the news landscape. By 2016, more than half of internet users said they used social media for news at least weekly. What information is most visible in people's feeds is increasingly shaped by social media algorithms. People who would previously not have explored scientific journals share articles on social media. Some leverage the authority of accredited scientists and peer-reviewed journals to support their views.

The COVID-19 pandemic accelerated these trends. Conventional science journalists have been joined by a host of other communicators whose stories only reflect the uncertainties of science when it suits their narrative. All through the height of the pandemic, *The Joe Rogan Experience* was one of the world's most popular podcasts, with more than 10 million subscribers. It features a controversial host who provides a platform for issues such as the efficacy of unorthodox COVID-19 treatments with the catchphrase "I'm just asking questions."

The early panic of the pandemic saw a retreat to some paternalistic approaches to science communication in the United States, United Kingdom, and elsewhere. Scientists, public authorities, and others were alarmed by the narratives spreading through the internet, convinced that public trust in science, so crucial to vaccination campaigns and efforts to forestall infection and disease, was being sabotaged. Some scientific organizations fell back on what communication scholars Matthew Nisbet and Dietram Scheufele describe as the "still dominant assumption that science literacy is both the problem and solution to societal conflicts."

Anything that was seen as undermining or even scrutinising the COVID-19 vaccines' safety record attracted disproportionate concern. As early as April 2020, Facebook created a special COVID-19 policy to combat misinformation: it produced a list of claims that it would prohibit in posts and ads and announced that it would boost funding for both algorithmic fact-checkers and human ones. The company boasted that when users saw posts that had been flagged, "95% of the time they did not go on to view the original content."

Several months before its whistleblower article was flagged, *BMJ* had aired concerns of overreach. In May 2021, a feature story asked, "Who fact checks the fact checkers?" The piece explained that Facebook had "removed 16 million pieces of its content and added warnings to around 167 million." It pointed to the "the difficulty of defining scientific truth" and questioned whether social media platforms should be charged with the task.

Blunt and sharp tools

In November 2021, *BMJ* published Thacker's reported story, which focused on Ventavia Research Group, a company running three sites in Texas as part of Pfizer's 153-site COVID-19 vaccine clinical trial. A whistleblower at the company provided evidence that patients were inadequately monitored after receiving vaccines, their treatment status had potentially (and inappropriately) been disclosed to clinicians, biohazard waste had been improperly disposed of, and data had been recorded irresponsibly.

In my interview with Rebecca Coombes, *BMJ*'s head of journalism, she told me there was a clear public interest in reporting on the whistleblower. "We were presented with very hard evidence of serious problems that had occurred in one of the world's most important, valuable pharmaceutical products, the Pfizer vaccine." Coombes was well aware that COVID-19 vaccines were being watched closely by those looking to exaggerate concerns about their safety or efficacy. "It is very important that the *BMJ* doesn't lose its courage at times like this," she told me. "It's still important to ask the big questions." She criticized Lead Stories for noting that a Twitter account associated with the whistleblower had also promoted anti-vaccine content as "guilt by association."

Three days after *BMJ* published its account, a false story appeared on a now-defunct Canadian website, the *Conservative Beaver*, claiming that the Pfizer CEO had been arrested and that this news was being suppressed by US outlets. The story was quickly debunked, but the rumor continued to circulate within social media. The *Conservative Beaver* page included a tweet from *BMJ* about the Ventavia whistleblower, providing a grain of truth around which people could weave the conspiracy.

Alan Duke, the editor-in-chief at Lead Stories, told me he saw "a big problem that it was *BMJ*, and that made it become more viral, more effective in the hands of the anti-vaxxers."

Lead Stories said their concern was that the piece was being "wildly misinterpreted by many people." Rather than focusing on the details of *BMJ*'s reporting, its statement focused on missing perspectives and the broader message: "Did the British Medical Association's news blog reveal flaws that disqualify the results of a contractor's field testing of Pfizer's COVID-19 vaccine, and were the problems ignored by the Food & Drug Administration and by Pfizer? No, that's not true ... The benefits of the Pfizer vaccine far outweigh rare side effects and the clinical trial data are solid."

But a close reading of *BMJ*'s article shows that the article didn't explicitly say that results of the Pfizer's trial were invalidated or explore implications for the risks and benefits of Pfizer's vaccine. Nor did the original article include any wording to discourage this interpretation. And while it stated that blinding was not maintained properly and that patients were not monitored appropriately following vaccination, it did not explore whether and how results might have been compromised or whether patients were indeed harmed. The text also did not provide a response from Ventavia or Pfizer about the accusations presented in the article.

Other outlets responded with their own context-setting articles: one, titled "Experts Blow Whistle on Alleged COVID Vaccine Whistleblower Claims," from MedPage Today, quoted a physician who described the allegations against Ventavia as a "vague kind of handwaving." A story from a CBS affiliate in North Carolina carried the headline "Fact check: Report questioning Pfizer trial shouldn't undermine confidence in vaccines." (Although I didn't ask her about these specific articles, Coombes told me that her team at *BMJ* was "very careful not to overplay the findings.")

From Lead Stories' perspective, *BMJ*'s authority and its rigorous internal fact-checking added to the problem. The story was dangerous in the hands of anti-vaxxers not because it was false, but because it was true—and so easily overinterpreted. Lead Stories editor Dean Miller thought *BMJ* was being naïve in the face of coronavirus conspiracies, telling me, "If you're going to handle sharp tools, you should use them well."

BMJ did not engage with Lead Stories' concerns about misinterpretation. Announcing plans to appeal to Facebook's Oversight Board, *BMJ* editor-in-chief Kamran Abbasi raised his own concerns about the platform's motivations: "The real question is: why is Facebook acting in this way? What is driving its world view? Is it ideology? Is it commercial interests? Is it incompetence? Users should be worried that, despite presenting itself as a neutral social media platform, Facebook is trying to control how people think under the guise of 'fact checking.'"

When I asked editor Dean Miller whether Lead Stories thought the public could be trusted to deal with uncertainty, he responded, "Don't condescend to the public. They can figure it out. They can muddle through all this and decide what to believe."

The fact-checkers' actions, however, suggest a different ethos. In conversations, I felt the Lead Stories editors revealed an old-fashioned model of paternalism, deployed through new means. They took a Manichean view, presuming that views could be separated into pro- or anti-science and pro- or anti-vaccine and that people could be divided into those with good intentions and those with bad intentions.

There are many within science who believe that, like laws and sausages, the public should not look too carefully at the processes by which scientific knowledge is made. Instead, we should just appreciate the end product. *BMJ* invited its readers backstage, into a debate about the processes of science—but stopped short of examining or intervening in what narrative played out in the social sphere.

The COVID-19 pandemic, with its high stakes, large uncertainties, and urgent need for decisions, is a case of what philosophers of science Silvio Functowicz and Jerry Ravetz call “postnormal science,” in which the old assumptions about scientific reliability look shaky. According to Functowicz and Ravetz, the answer in such cases is not to defend the barricades of laboratory science, but instead to seek “extended peer review” in the form of more dialogue with the public, building

Fact-checking is a naïve response to a complex problem. Facebook's fact-checkers face an impossible task, unable to control the tide of information promulgated by Facebook's algorithms, which are tuned to emphasize controversy over veracity. As scholars such as sociologist Tarleton Gillespie have explained, Facebook, though claiming to be merely a platform, plays a powerful but chaotic editorial role.

Overall, social media invites conversation that is contentious rather than constructive. In my opinion, Lead Stories' belligerence triggered an overly defensive reaction from *BMJ*. It stepped on *BMJ*'s editorial responsibilities, adding a label without first offering suggestions or improvements. And while *BMJ*'s Coombes complained to me that the institution suffered “reputational damage” because of labels that discouraged online sharing, the publication does not seem to have yet worked out what its new roles and responsibilities should be in a world of online sharing. Tellingly, *BMJ* never issued a statement to clarify that its article was about clinical trial oversight, not about the vaccine's effectiveness.

Social media companies' importance to science communication will continue to grow, but they have resisted engaging with their responsibilities as publishers. And as

There are many within science who believe that, like laws and sausages, the public should not look too carefully at the processes by which scientific knowledge is made.

what sociologist Helga Nowotny and colleagues call “socially-robust knowledge.” For public health issues, particularly those involving vaccination, this would mean a greater emphasis on local conversations. Dialogue between community doctors and citizens, for instance, could complement and inform top-down campaigns.

For most of its history, science communication has depended on good-faith actors: benevolent scientists inform a receptive public of their discoveries. The presumption of good faith was based on Robert Merton's old idea of organized skepticism—that science was a self-governing entity. The question with today's decidedly disorganized skepticism is how new models of trust can be built.

In this environment, fact-checking might help interrupt the spread of obvious falsehoods, but it won't resolve emerging situations in which evidence cannot be separated from its context. Tedros Adhanom Ghebreyesus, director-general of the World Health Organization (WHO), told a security conference in February 2020, “Fake news spreads faster and more easily than this virus, and is just as dangerous.” But his call for “facts, not fear” was undermined by the uncertainties that inevitably accompany a new disease. As evidence accumulated, WHO itself had to change its messaging about whether COVID-19 was airborne and whether wearing masks curbed its spread.

science publishers are pressured to become more transparent, so should platforms. To start, they should allow probes of their own algorithms. Meta's Oversight Board, a body of journalists, scholars, lawyers, and others, is a case in point. A recent story in *WIRED* reported that the board's remit is to cover posts and content, not the algorithms that influence users' attention, and that relations between Meta and the board have been fractious. If there is no way to scrutinize the wellspring for how so many people now receive information about science, it will be harder to fix problems downstream. In an ecosystem where algorithms encourage engagement, controversy, and trolling, information presented to suit charged narratives will spread.

Simple fact-checking cannot solve this problem. Social media companies should seek healthy relationships with conventional science publishers, and science publishers should return the favor. Solutions will require dialogue with platforms such as Facebook, transparency about social media algorithms, new institutions capable of protecting the public interest, and resources to pursue solutions. It is indeed time to bring the public behind the scenes of science, but this brings with it new questions and editorial responsibilities.

Jack Stilgoe is a professor of science and technology policy at University College London.