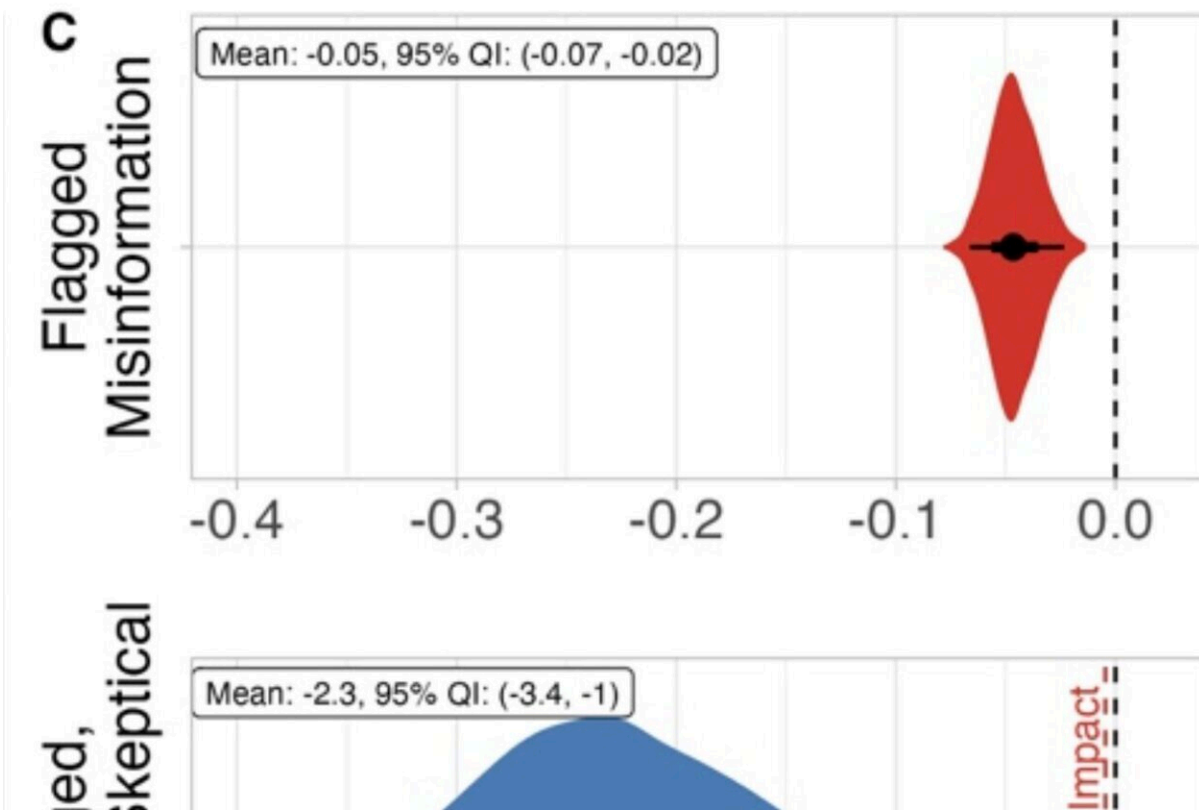


Misleading COVID-19 headlines from mainstream sources did more harm on Facebook than fake news, study finds

May 30 2024



Despite the greater potency of "fake news" on Facebook to discourage Americans from taking the COVID-19 vaccine, users' greater exposure to unflagged, vaccine-skeptical content meant the latter had a much greater

negative effect on vaccine uptake. Credit: Jennifer Allen, Duncan Watts, David G. Rand

Since the rollout of the COVID-19 vaccine in 2021, fake news on [social media has been widely blamed for low vaccine uptake](#) in the United States—but research by MIT Sloan School of Management Ph.D. candidate Jennifer Allen and Professor David Rand finds that the blame lies elsewhere.

In a new [paper](#) published in *Science* and co-authored by Duncan J. Watts of the University of Pennsylvania, the researchers introduce a [new methodology](#) for measuring [social media](#) content's causal impact at scale. They show that misleading content from mainstream news sources—rather than outright misinformation or "[fake news](#)"—was the primary driver of vaccine hesitancy on Facebook.

A new approach to estimating impact

"Misinformation has been correlated with many societal challenges, but there's not a lot of research showing that exposure to misinformation actually causes harm," explained Allen.

During the COVID-19 pandemic, for example, the spread of misinformation related to the virus and vaccine received significant public attention. However, existing research has, for the most part, only established correlations between vaccine refusal and factors such as sharing misinformation online—and largely overlooked the role of "vaccine-skeptical" content, which was potentially misleading but not flagged as misinformation by Facebook fact-checkers.

To address that gap, the researchers first asked a key question: What would be necessary for misinformation or any other type of content to have far-reaching impacts?

"To change behavior at scale, content has to not only be persuasive enough to convince people not to get the vaccine, but also widely seen," Allen said. "Potential harm results from the combination of persuasion and exposure."

To quantify content's persuasive ability, the researchers conducted randomized experiments in which they showed thousands of survey participants the headlines from 130 vaccine-related stories—including both mainstream content and known misinformation—and tested how those headlines impacted their intentions to get vaccinated against COVID-19.

Researchers also asked a separate group of respondents to rate the headlines across various attributes, including plausibility and political leaning. One factor reliably predicted impacts on vaccination intentions: the extent to which a [headline](#) suggested that the vaccine was harmful to a person's health.

Using the "wisdom of crowds" and [natural language](#) processing AI tools, Allen and her co-authors extrapolated those survey results to predict the persuasive power of all 13,206 vaccine-related URLs that were widely viewed on Facebook in the first three months of the vaccine rollout.

By combining these predictions with data from Facebook showing the number of users who viewed each URL, the researchers could predict each headline's overall impact—the number of people it might have persuaded not to get the vaccine. The results were surprising.

The underestimated power of exposure

Contrary to popular perceptions, the researchers estimated that vaccine-skeptical content reduced vaccination intentions 46 times more than misinformation flagged by fact-checkers.

The reason? Even though flagged misinformation was more harmful when seen, it had relatively low reach. In total, the vaccine-related headlines in the Facebook data set received 2.7 billion views—but content flagged as misinformation received just 0.3% of those views, and content from domains rated as low-credibility received 5.1%.

"Even though the outright false content reduced vaccination intentions the most when viewed, comparatively few people saw it," explained Rand. "Essentially, that means there's this class of gray-area content that is less harmful per exposure but is seen far more often—and thus more impactful overall—that has been largely overlooked by both academics and social media companies."

Notably, several of the most impactful URLs within the data set were articles from mainstream sources that cast doubt on the vaccine's safety. For instance, the most-viewed was an article—from a well-regarded mainstream news source—suggesting that a [medical doctor](#) died two weeks after receiving the COVID-19 vaccine. This single headline received 54.9 million views—more than six times the combined views of all flagged [misinformation](#).

While the body of this article did acknowledge the uncertainty of the doctor's cause of death, its "clickbait" headline was highly suggestive and implied that the vaccine was likely responsible. That's significant since the vast majority of viewers on social media likely never click out to read past the headline.

How journalists and social media platforms can help

According to Rand, one implication of this work is that [media outlets](#) need to take more care with their headlines, even if that means they aren't as attention-grabbing.

"When you are writing a headline, you should not just be asking yourself if it's false or not," he said. "You should be asking yourself if the headline is likely to cause inaccurate perceptions."

For platforms, added Allen, the research also points to the need for more nuanced moderation—across all subjects, not just public health.

"Content moderation focuses on identifying the most egregiously false information—but that may not be an effective way of identifying the most overall harmful content," she says. "Platforms should also prioritize reviewing content from the people or organizations with the largest numbers of followers while balancing freedom of expression. We need to invest in more research and creative solutions in this space—for example, crowdsourced moderation tools like X's Community Notes."

"Content moderation decisions can be really difficult because of the inherent tension between wanting to mitigate harm and allowing people to express themselves," Rand said. "Our paper introduces a framework to help balance that trade-off by allowing tech companies to actually quantify potential harm."

And the trade-offs could be large. An exploratory analysis by the authors found that if Facebook users hadn't been exposed to this [vaccine](#)-skeptical content, as many as 3 million more Americans could have been vaccinated.

"We can't just ignore this gray area-content," Allen concluded. "Lives could have been saved."

More information: Jennifer Allen, Quantifying the impact of misinformation and vaccine-skeptical content on Facebook, *Science* (2024). [DOI: 10.1126/science.adk3451](https://doi.org/10.1126/science.adk3451).
www.science.org/doi/10.1126/science.adk3451

Provided by MIT Sloan School of Management

Citation: Misleading COVID-19 headlines from mainstream sources did more harm on Facebook than fake news, study finds (2024, May 30) retrieved 18 June 2024 from <https://phys.org/news/2024-05-covid-headlines-mainstream-sources-facebook.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.