


# **Facebook has a misinformation problem, and is blocking access to data about how much there is and who is affected**

November 3 2021, by Ethan Zuckerman

---



Credit: Pixabay/CC0 Public Domain

Leaked internal documents suggest Facebook—which recently renamed itself Meta—is [doing far worse than it claims](#) at minimizing COVID-19 vaccine misinformation on the Facebook social media platform.

Online [misinformation](#) about the virus and vaccines is a major concern. In one study, survey respondents who got some or all of their news from Facebook were significantly [more likely to resist the COVID-19 vaccine](#) than those who got their news from mainstream media sources.

As a researcher who [studies social and civic media](#), I believe it's critically important to understand how misinformation spreads online. But this is easier said than done. Simply counting instances of misinformation found on a social media platform leaves two key questions unanswered: How likely are users to encounter misinformation, and are certain users especially likely to be affected by misinformation? These questions are the denominator problem and the distribution problem.

The COVID-19 misinformation study, "[Facebook's Algorithm: a Major Threat to Public Health](#)", published by public interest advocacy group Avaaz in August 2020, reported that sources that frequently shared health misinformation—82 websites and 42 Facebook pages—had an estimated total reach of 3.8 billion views in a year.

At first glance, that's a stunningly large number. But it's important to remember that this is the [numerator](#). To understand what 3.8 billion views in a year means, you also have to calculate the [denominator](#). The

numerator is the part of a fraction above the line, which is divided by the part of the fraction below line, the denominator.

## Getting some perspective

One possible denominator is [2.9 billion monthly active Facebook users](#), in which case, on average, every Facebook user has been exposed to at least one piece of information from these health misinformation sources. But these are 3.8 billion content views, not discrete users. How many pieces of information does the average Facebook user encounter in a year? Facebook does not disclose that information.

Market researchers estimate that Facebook users spend from [19 minutes a day](#) to [38 minutes a day](#) on the platform. If the [1.93 billion daily active users](#) of Facebook see an average of 10 posts in their daily sessions—a very conservative estimate—the denominator for that 3.8 billion pieces of information per year is 7.044 trillion (1.93 billion daily users times 10 daily posts times 365 days in a year). This means roughly 0.05% of content on Facebook is posts by these suspect Facebook pages.

The 3.8 billion views figure encompasses all content published on these pages, including innocuous health content, so the proportion of Facebook posts that are health misinformation is smaller than one-twentieth of a percent.

Is it worrying that there's enough misinformation on Facebook that everyone has likely encountered at least one instance? Or is it reassuring that 99.95% of what's shared on Facebook is not from the sites Avaaz warns about? Neither.

## Misinformation distribution

In addition to estimating a denominator, it's also important to consider the distribution of this information. Is everyone on Facebook equally likely to encounter health misinformation? Or are people who identify as anti-vaccine or who seek out "alternative health" information more likely to encounter this type of misinformation?

Another social media study focusing on extremist content on YouTube offers a method for understanding the distribution of misinformation. [Using browser data from 915 web users](#), an Anti-Defamation League team recruited a large, demographically diverse sample of U.S. web users and oversampled two groups: heavy users of YouTube, and individuals who showed strong negative racial or gender biases in a set of questions asked by the investigators. [Oversampling](#) is surveying a small subset of a population more than its proportion of the population to better record data about the subset.

The researchers found that 9.2% of participants viewed at least one video from an extremist channel, and 22.1% viewed at least one video from an alternative channel, during the months covered by the study. An important piece of context to note: A small group of people were responsible for most views of these videos. And more than 90% of views of extremist or "alternative" videos were by people who reported a high level of racial or gender resentment on the pre-study survey.

While roughly 1 in 10 people found extremist content on YouTube and 2 in 10 found content from right-wing provocateurs, most people who encountered such content "bounced off" it and went elsewhere. The group that found extremist content and sought more of it were people who presumably had an interest: people with strong racist and sexist attitudes.

The authors concluded that "consumption of this potentially harmful content is instead concentrated among Americans who are already high

in racial resentment," and that YouTube's algorithms may reinforce this pattern. In other words, just knowing the fraction of users who encounter extreme content doesn't tell you how many people are consuming it. For that, you need to know the distribution as well.

## **Superspreaders or whack-a-mole?**

A widely publicized study from the anti-hate speech advocacy group Center for Countering Digital Hate titled [Pandemic Profiteers](#) showed that of 30 anti-vaccine Facebook groups examined, 12 anti-vaccine celebrities were responsible for 70% of the content circulated in these groups, and the three most prominent were responsible for nearly half. But again, it's critical to ask about denominators: How many anti-vaccine groups are hosted on Facebook? And what percent of Facebook users encounter the sort of information shared in these groups?

Without information about denominators and distribution, the study reveals something interesting about these 30 anti-vaccine Facebook groups, but nothing about medical misinformation on Facebook as a whole.

These types of studies raise the question, "If researchers can find this content, why can't the social media platforms identify it and remove it?" The Pandemic Profiteers study, which implies that Facebook could solve 70% of the medical misinformation problem by deleting only a dozen accounts, explicitly advocates for the deplatforming of these dealers of disinformation. However, I found that 10 of the 12 anti-vaccine influencers featured in the study have already been removed by Facebook.

Consider Del Bigtree, one of the three most prominent spreaders of vaccination disinformation on Facebook. The problem is not that Bigtree is recruiting new anti-vaccine followers on Facebook; it's that Facebook

users follow Bigtree on other websites and bring his content into their Facebook communities. It's not 12 individuals and groups posting health misinformation online—it's likely thousands of individual Facebook users sharing misinformation found elsewhere on the web, featuring these dozen people. It's much harder to ban thousands of Facebook users than it is to ban 12 anti-vaccine celebrities.

This is why questions of denominator and distribution are critical to understanding misinformation online. Denominator and distribution allow researchers to ask how common or rare behaviors are online, and who engages in those behaviors. If millions of users are each encountering occasional bits of medical misinformation, [warning labels](#) might be an effective intervention. But if medical misinformation is consumed mostly by a smaller group that's actively seeking out and sharing this content, those warning labels are most likely useless.

## Getting the right data

Trying to understand misinformation by counting it, without considering denominators or distribution, is what happens when good intentions collide with poor tools. No social media platform makes it possible for researchers to accurately calculate how prominent a particular piece of content is across its platform.

Facebook restricts most researchers to its [Crowdtangle](#) tool, which shares information about content engagement, but this is not the same as content views. [Twitter explicitly prohibits researchers](#) from calculating a denominator, either the number of Twitter users or the number of tweets shared in a day. YouTube makes it so difficult to find out how many videos are hosted on their service that Google routinely asks interview candidates to [estimate the number of YouTube videos hosted](#) to evaluate their quantitative skills.

The leaders of social media platforms have argued that their tools, despite their problems, [are good for society](#), but this argument would be more convincing [if researchers could independently verify that claim](#).

As the societal impacts of social media become more prominent, pressure on the big tech platforms to release more data about their users and their content is likely to increase. If those companies respond by increasing the amount of information that researchers can access, look very closely: Will they let researchers study the denominator and the distribution of content online? And if not, are they afraid of what researchers will find?

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: Facebook has a misinformation problem, and is blocking access to data about how much there is and who is affected (2021, November 3) retrieved 9 April 2024 from <https://techxplore.com/news/2021-11-facebook-misinformation-problem-blocking-access.html>

<p>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.</p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------