Elon Musk says relaxing content rules on Twitter will boost free speech, but research shows otherwise

May 9 2022, by Filippo Menczer

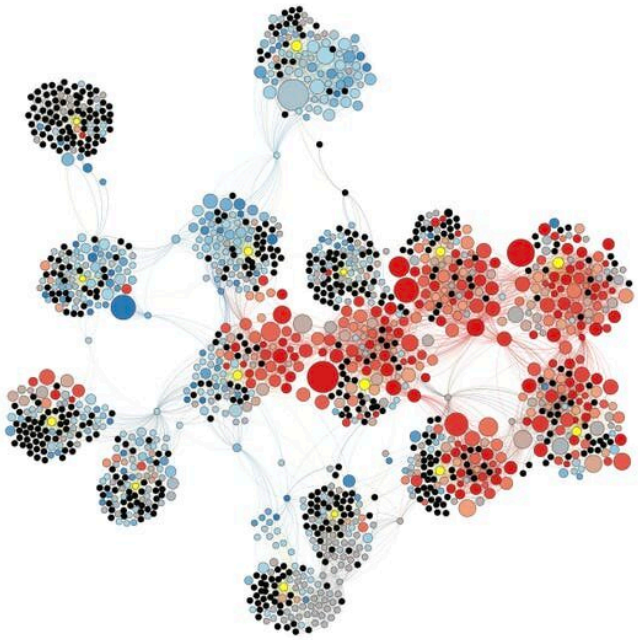

Neutral bots (yellow nodes) and a sample of their friends and followers in an experiment to measure partisan bias on Twitter. Node color indicates political alignment of an account: red for conservative, blue for liberal, black for unknown. Node size is proportional to share of links to low-credibility sources. The closely connected red clusters indicate conservative echo chambers. Credit: Filippo Menczer, CC BY-ND

Elon Musk's accepted bid to purchase Twitter has triggered a lot of debate about what it means for the future of the social media platform, which plays an important role in determining the news and information many people—especially Americans—are exposed to.

Musk has said he wants to make Twitter an arena for free speech. It's not clear what that will mean, and his statements have fueled speculation among both supporters and detractors. As a corporation, Twitter can regulate speech on its platform as it chooses. There are bills being considered in the U.S. Congress and by the European Union that address social media regulation, but these are about transparency, accountability, illegal harmful content and protecting users' rights, rather than regulating speech.

Musk's calls for free speech on Twitter focus on two allegations: political bias and excessive moderation. As researchers of online misinformation and manipulation, my colleagues and I at the Indiana University Observatory on Social Media study the dynamics and impact of Twitter and its abuse. To make sense of Musk's statements and the possible outcomes of his acquisition, let's look at what the research shows.

Political bias

Many conservative politicians and pundits have alleged for years that major social media platforms, including Twitter, have a liberal political bias amounting to censorship of conservative opinions. These claims are based on anecdotal evidence. For example, many partisans whose tweets were labeled as misleading and downranked, or whose accounts were suspended for violating the platform's terms of service, claim that Twitter targeted them because of their political views.

Unfortunately, Twitter and other platforms often inconsistently enforce their policies, so it is easy to find examples supporting one conspiracy

theory or another. A review by the Center for Business and Human Rights at New York University has found no reliable evidence in support of the claim of anti-conservative bias by social media companies, even labeling the claim itself a form of disinformation.

A more direct evaluation of political bias by Twitter is difficult because of the complex interactions between people and algorithms. People, of course, have political biases. For example, our experiments with political social bots revealed that Republican users are more likely to mistake conservative bots for humans, whereas Democratic users are more likely to mistake conservative human users for bots.

To remove human bias from the equation in our experiments, we deployed a bunch of benign social bots on Twitter. Each of these bots started by following one news source, with some bots following a liberal source and others a conservative one. After that initial friend, all bots were left alone to "drift" in the information ecosystem for a few months. They could gain followers. They acted according to an identical algorithmic behavior. This included following or following back random accounts, tweeting meaningless content and retweeting or copying random posts in their feed.

But this behavior was politically neutral, with no understanding of content seen or posted. We tracked the bots to probe political biases emerging from how Twitter works or how users interact.

Surprisingly, our research provided evidence that Twitter has a conservative, rather than a liberal bias. On average, accounts are drawn toward the conservative side. Liberal accounts were exposed to moderate content, which shifted their experience toward the political center, while the interactions of right-leaning accounts were skewed toward posting conservative content. Accounts that followed conservative news sources also received more politically aligned followers, becoming embedded in

denser echo chambers and gaining influence within those partisan communities.

These differences in experiences and actions can be attributed to interactions with users and information mediated by the social media platform. But we could not directly examine the possible bias in Twitter's news feed algorithm, because the actual ranking of posts in the "home timeline" is not available to outside researchers.

Researchers from Twitter, however, were able to audit the effects of their ranking algorithm on political content, unveiling that the political right enjoys higher amplification compared to the political left. Their experiment showed that in six out of seven countries studied, conservative politicians enjoy higher algorithmic amplification than liberal ones. They also found that algorithmic amplification favors right-leaning news sources in the U.S.

Our research and the research from Twitter show that Musk's apparent concern about bias on Twitter against conservatives is unfounded.

Referees or censors?

The other allegation that Musk seems to be making is that excessive moderation stifles free speech on Twitter. The concept of a free marketplace of ideas is rooted in John Milton's centuries-old reasoning that truth prevails in a free and open exchange of ideas. This view is often cited as the basis for arguments against moderation: accurate, relevant, timely information should emerge spontaneously from the interactions among users.

Unfortunately, several aspects of modern social media hinder the free marketplace of ideas. Limited attention and confirmation bias increase vulnerability to misinformation. Engagement-based ranking can amplify

noise and manipulation, and the structure of information networks can distort perceptions and be "gerrymandered" to favor one group.

As a result, social media users have in past years become victims of manipulation by "astroturf" causes, trolling and misinformation. Abuse is facilitated by social bots and coordinated networks that create the appearance of human crowds.

We and other researchers have observed these inauthentic accounts amplifying disinformation, influencing elections, committing financial fraud, infiltrating vulnerable communities and disrupting communication . Musk has tweeted that he wants to defeat spam bots and authenticate humans, but these are neither easy nor necessarily effective solutions.

Inauthentic accounts are used for malicious purposes beyond spam and are hard to detect, especially when they are operated by people in conjunction with software algorithms. And removing anonymity may harm vulnerable groups. In recent years, Twitter has enacted policies and systems to moderate abuses by aggressively suspending accounts and networks displaying inauthentic coordinated behaviors. A weakening of these moderation policies may make abuse rampant again.

Manipulating Twitter

Despite Twitter's recent progress, integrity is still a challenge on the platform. Our lab is finding new types of sophisticated manipulation, which we will present at the International AAAI Conference on Web and Social Media in June. Malicious users exploit so-called "follow trains "—groups of people who follow each other on Twitter—to rapidly boost their followers and create large, dense hyperpartisan echo chambers that amplify toxic content from low-credibility and conspiratorial sources.

Another effective malicious technique is to post and then strategically

delete content that violates platform terms after it has served its purpose. Even Twitter's high limit of 2,400 tweets per day can be circumvented through deletions: We identified many accounts that flood the network with tens of thousands of tweets per day.

We also found coordinated networks that engage in repetitive likes and unlikes of content that is eventually deleted, which can manipulate ranking algorithms. These techniques enable malicious users to inflate content popularity while evading detection.

Musk's plans for Twitter are unlikely to do anything about these manipulative behaviors.

Content moderation and free speech

Musk's likely acquisition of Twitter raises concerns that the social media platform could decrease its content moderation. This body of research shows that stronger, not weaker, moderation of the information ecosystem is called for to combat harmful misinformation.

It also shows that weaker moderation policies would ironically hurt free speech: The voices of real users would be drowned out by malicious users who manipulate Twitter through inauthentic accounts, bots and echo chambers.

This article is republished from The Conversation under a Creative Commons license. Read the original article.

Provided by The Conversation

Citation: Elon Musk says relaxing content rules on Twitter will boost free speech, but research shows otherwise (2022, May 9) retrieved 25 April 2024 from

https://techxplore.com/news/2022-05-elon-musk-content-twitter-boost.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.