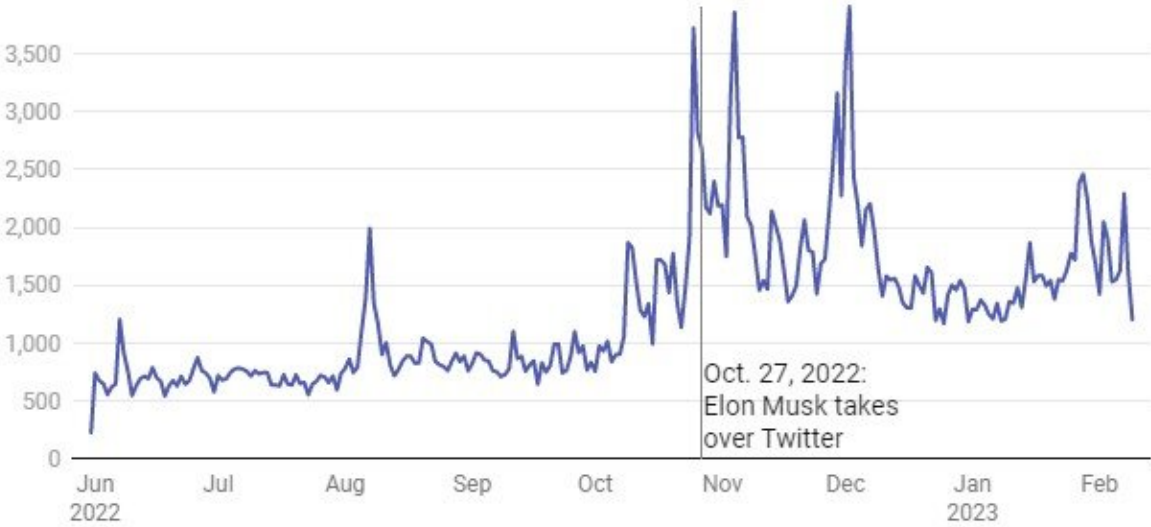


Antisemitism on Twitter has more than doubled since Elon Musk took over the platform, says researcher

March 21 2023, by Carl Miller

Antisemitic tweets reached new heights as Musk took helm

A relatively flat level of antisemitic tweets more than tripled around the time Elon Musk took control of Twitter on Oct. 27, 2022.



Credit: The Conversation

In the days after Elon Musk [took over Twitter](#) in October 2022, the social media platform saw a "surge in hateful conduct," which its [then safety chief put down to](#) a "focused, short-term trolling campaign." New research suggests that when it comes to antisemitism, it was anything

but.

Rather, antisemitic [tweets](#) have [more than doubled](#) over the months since Musk took charge, according to [research that I](#) and colleagues at tech firm [CASM Technology](#) and the [Institute for Strategic Dialogue](#) think tank conducted. Between June and Oct. 26, 2022, the day before Twitter's acquisition by Musk, there was a weekly average of 6,204 tweets deemed "plausibly antisemitic"—that is, where at least one reasonable interpretation of the tweet falls within the International Holocaust Remembrance Alliance's [definition of the term](#) as "a certain perception of Jews, which may be expressed as hatred towards Jews."

But from Oct. 27 until Feb 9, 2023, the average was 12,762—an increase of 105%. In all, a total of 325,739 tweets from 146,516 accounts were labeled as "plausibly antisemitic" over the course of our study, stretching from June 1, 2022 to Feb. 9, 2023.

Finding antisemitism with AI

To identify plausibly antisemitic tweets, my co-authors and I combined 22 published hate speech-identifying algorithms into a single mechanism and used even more machine learning to see which combinations of decisions led to the correct result. We then passed through all tweets—over a million in total—that contained any one of 119 words, phrases, slurs and epithets related to antisemitism.

No such process is perfect. We estimate our model to make a correct decision about 75% of the time. We also no doubt missed some antisemitic tweets not containing any of those 119 key words, as well as those taken down before early December when we collected the data.

We then used an algorithm to draw out 10 different themes of antisemitism seen in the tweets. Some centered around the use of

specific antisemitic derogatory epithets. Others alluded to [conspiracy theories](#) concerning hidden Jewish influence and control.

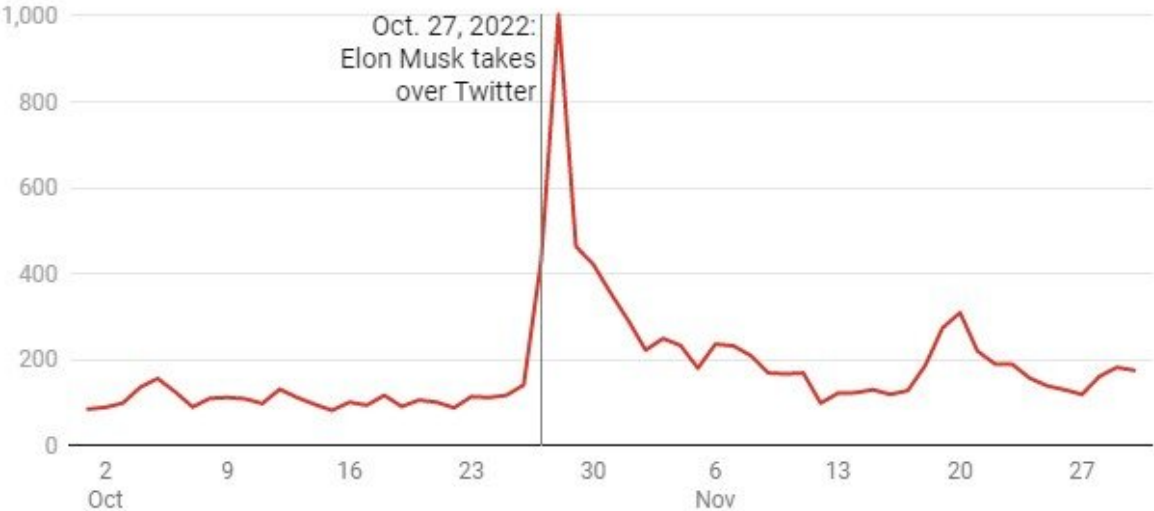
Antisemitic tweets directed at Jewish investor and philanthropist George Soros warranted its own category. He was mentioned more than any other person in our data, over 19,000 times, with tweets [claiming he was a member](#) of a hidden globalist, Jewish or "Nazi" world order.

Another theme were tweets defending the rapper Ye, formerly Kanye West, who had made a [number of antisemitic remarks](#) after he had his account briefly reinstated by Musk.

Our research, which has not yet been peer-reviewed, also found around 4,000 of the antisemitic tweets were [focused on the Russian invasion of Ukraine](#). These variously claimed that the conflict was caused by Jews, or that Jews secretly caused the U.S. to support Ukraine. They also contained direct antisemitism directed against the Ukrainian president, Volodymyr Zelenskyy, who is Jewish.

New antisemitic Twitter accounts spiked as Musk took over

The number of antisemitic Twitter accounts created in late 2022 saw a massive increase around Oct. 27, the day Elon Musk took over the company.



Credit: The Conversation

Musk rolls back content moderation

Musk's acquisition of Twitter came on the back of what I have observed as a decadelong trend among tech giants to take more responsibility for hate speech, harassment, incitement, disinformation and other harms lurking in the information flowing through their platforms. Over that period, companies such as Facebook and Twitter gradually enacted policies to respond to extremism, hate speech and harassment, or increase "civility," as Twitter itself [described it in 2018](#), and built out the teams and tools to enforce them.

Musk, a self-professed "[free speech](#) absolutist," pointed the platform in a different direction after taking control. In short order, Twitter's independent Trust and Safety Council was [dissolved](#), previously [banned](#)

[accounts were reinstated](#) and over half of [Twitter's staff was laid off](#) or simply left—including many of those responsible for enforcing the company's hate speech policies.

As someone who has tracked hate speech on places like Twitter for around 10 years, I believe the changes to Twitter's moderation practices are only partly to blame for the jump in antisemitism on the platform.

The media spectacle surrounding Musk's takeover, along with his very vocal views on free speech, likely also encouraged exactly those people to join or rejoin the platform who had fallen foul of its previous attempts to confront hate. Our research gives some backing to this theory. Some 3,855 accounts we identified as posting at least one plausibly antisemitic tweet joined Twitter in the 10 days after Musk took over. This is, however, only a small proportion of the 146,516 accounts that sent at least one antisemitic tweet over the course of the entire study.

Little effect on curbing hate speech

A surge in [hate speech](#) on Twitter was [flagged](#) by researchers in the weeks after Musk took over, concerns the billionaire [dismissed](#) as "utterly false," having earlier vowed to "[max deboosted & demonetized](#)" hateful tweets.

If Twitter has been de-amplifying antisemitism, our research shows almost no evidence of it. Before Oct. 27, antisemitic tweets received an average of 6.4 "favorites" and 1.2 retweets. Since then, they have averaged six "favorites" and 1 retweet. Although such engagement isn't a perfect measure for visibility, tweets made much less visible to users would generally receive less engagement.

We also attempted to measure takedowns of antisemitic tweets. On Feb.

15, 45 days after we initially collected the data, we tried to re-collect all the tweets we identified as antisemitic. Tweets can be unavailable for lots of reasons, and Twitter's enforcement is only one of them. Imperfect though this is, it does give us a tentative glimpse of what might be happening in regard to the removal of antisemitic posts. And across those dates, 17,589 antisemitic tweets were taken down—8.5% of the total.

Rising tide of antisemitism

Our findings come at a time when many fear [growing threats to Jewish communities](#). In 2021, the Anti-Defamation League tracked the highest number of antisemitic incidents—including harassment, vandalism and assaults—in the U.S. since they started tracking numbers in 1979. And this is not just a U.S. phenomenon; in the U.K., the Community Security Trust has recorded [a similar spike in anti-Jewish activity](#), while in Germany, anti-Jewish crimes [surged](#) by 29% over the pandemic.

Studying [social media](#) has shown me again and again just how powerfully it helps to form the cultures and ideas that underlie its users' behavior. Ultimately, the proliferation of tweets that hold Jews responsible for all the world's ills, that circulate dark conspiracies of control and cover-up, or that fire derogatory attacks directed toward Jews, can only support antisemitism online—and in the real world.

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

Provided by The Conversation

Citation: Antisemitism on Twitter has more than doubled since Elon Musk took over the platform, says researcher (2023, March 21) retrieved 27 April 2024 from

<https://techxplore.com/news/2023-03-antisemitism-twitter-elon-musk-platform.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.