

Antisemitic tweets soared on Twitter after Musk took over, study finds

By Cristiano Lima with research by David DiMolfetta

In the months after Elon Musk's takeover, antisemitic posts on Twitter skyrocketed, according to a report shared first with The Technology 202, which offers a new detailed look into the growing prevalence of hate speech on the site.

The study, which used machine-learning tools to identify likely antisemitic tweets, found that the average weekly number of such posts "more than doubled after Musk's acquisition" — a trend that has held in the months after Musk took over.

The analysis found an average of over 6,200 posts per week appearing to contain antisemitic language between June 1 and Oct. 27, the day Musk completed his \$44 billion deal to buy Twitter. But that figure rose to over 12,700 through early February — a 105 percent increase.

The report — conducted by the Institute for Strategic Dialogue (ISD), a nonpartisan think tank, and CASM Technology, a start-up that researches disinformation and hate speech online — also found a "surge" in the number of new accounts created immediately after Musk took over that posted at least some antisemitic content.

Researchers wrote that it represented a three-fold increase in the rate of "hateful account creation." But critically, the researchers behind the study said the uptick in hateful content extended well beyond that initial wave of new accounts.

"We're seeing a sustained volume of antisemitic hate speech on the platform following the takeover," said Jacob Davey, who leads research and policy on the far-right and hate movements at ISD.

The study marks one of the most extensive efforts to date to quantify how Musk's drastic makeover of the company has impacted the prevalence of hate speech on the platform.

According to the report, researchers trained a machine-learning tool to spot tweets that "plausibly" matched at least one interpretation of the International Holocaust Remembrance Alliance's definition of antisemitism. The organization lists making "dehumanizing, demonizing, or stereotypical allegations about Jews" and "Calling for, aiding, or justifying the killing or harming of Jews" as examples of antisemitic remarks.

Researchers then manually reviewed a smaller subset of the posts to compare it with their algorithmic sorting tool, finding that it matched with 76 percent accuracy.

"There are inherent challenges in training language models on as nuanced a topic as antisemitism," the researchers wrote.

Even with the caveats, researchers say the findings paint a clear picture: Antisemitic tweets have become far more prevalent under Musk.

“We’re pretty confident that this is the most sophisticated attempt to map antisemitism on Twitter in the pre- and post-Musk era,” said Tim Squirrel, ISD’s head of communications.

Twitter replied to a request for comment with an email containing a poop emoji. Musk tweeted Sunday that Twitter’s press email will automatically respond in that manner.



Musk has denied claims that hate speech has risen on Twitter under his leadership.

After a December report by advocacy groups found that hate speech was appearing on the site more often on average after his acquisition, Musk tweeted a graph that he said showed “hate speech impressions” were on the “decline” — without providing additional data to substantiate the claim. He suggested that one spike in engagement was due to a small number of accounts.

Musk has also claimed that Twitter will limit the circulation of “hate tweets” so that they are “max deboosted,” part of a policy he described as “freedom of speech, but not freedom of reach.”

But according to the new findings, any impact from those changes appears marginal, with researchers identifying “only a very small decrease in the average levels of engagement” with antisemitic content and stating that it did not amount to an “appreciable change.”

Milo Comerford, who leads policy and research on counter-extremism for ISD, said the slight drop in engagement with antisemitic posts could be explained by the surge in overall volume.

“When you have a substantially higher volume of content that is propagating particular narratives, you can't expect all of that content to continue to have the same level of engagement,” he said.

Researchers and Democratic lawmakers have expressed concern that it could be significantly harder to track the prevalence and reach of hate speech on Twitter under Musk, given his plans to charge outside groups a significant amount to access data about the platform.

