

A HOME FOR HATE

How Antisemitism is Reaching Millions on X





The Center for Countering Digital Hate works to stop the spread of online hate and disinformation through innovative research, public campaigns and policy advocacy.

Our mission is to protect human rights and civil liberties online.

Social media platforms have changed the way we communicate, build and maintain relationships, set social standards, and negotiate and assert our society's values. In the process, they have become safe spaces for the spread of hate, conspiracy theories and disinformation.

Social media companies erode basic human rights and civil liberties by enabling the spread of online hate and disinformation.

At CCDH, we have developed a deep understanding of the online harm landscape, showing how easily hate actors and disinformation spreaders exploit the digital platforms and search engines that promote and profit from their content.

We are fighting for better online spaces that promote truth, democracy, and are safe for all. Our goal is to increase the economic and reputational costs for the platforms that facilitate the spread of hate and disinformation.

If you appreciate this report, you can donate to CCDH at counterhate.com/donate. In the United States, Center for Countering Digital Hate Inc is a 501(c)(3) charity. In the United Kingdom, Center for Countering Digital Hate Ltd is a nonprofit company limited by guarantee.



Since its founding 81 years ago, the Jewish Council for Public Affairs (JCPA) has served as the national convener of Jewish coalitions working to build a just and inclusive democracy—bringing together national and local partners to address the issues that most deeply affect our community and advance our most essential values. Throughout its history, JCPA has been at the forefront of the fight for civil rights, justice, and equality in the United States.

Today, JCPA is mobilizing the robust community relations network and its partners in pursuit of a more just and democratic society for all – based on an abiding belief that Jewish safety is inextricably linked with the safety of other communities and a strong, pluralistic democracy.

This is even more crucial at a moment when rising attempts to normalize antisemitism and other forms of hate and extremism seek to drive wedges between the Jewish community and other targeted groups. Growing extremism, disinformation, and attacks on inclusive democracy are fueling rollbacks of rights and a cycle of hate and violence that profoundly endangers Jews and so many others.

JCPA's work is rooted in the understanding that only by overcoming lines of difference and fostering solidarity across communities can we advance an inclusive future in which Jews – and all people – are safe and free.

You can learn more about the JCPA's work and how to support it at www.jewishpublicaffairs.org.

CONTENTS

1 Introduction	5
2 Executive Summary	6
3 Key Findings	9
4 Methodology	10
5 Volume And Views	13
6 Antisemitic Conspiracies	16
7 Effectiveness Of Community Notes	18
8 Antisemitism Influencers	20
Appendix: Methodology	26
Appendix: Limitations	35
Appendix: Examples Of Antisemitism From List Of Top 10 “Antisemitism Influencers”	36
Appendix: Examples Of Ads By Monetized “Antisemitism Influencers”	41
References	46

Note: This report’s AI analysis of antisemitic content is 85% accurate

Researchers used OpenAI’s GPT-4o to analyze 1.1 million posts, identifying and categorizing antisemitic content. After rigorous manual checks, researchers concluded that GPT-4o has an F1 score of 0.85, on a scale between 0 and 1. An F1 score is used to measure a model’s accuracy and is a weighted average of the following two metrics:

- Recall: the number of posts for a category that were correctly identified by the model as belonging to that category.
- Precision: the number of posts labelled as belonging to a category by the model that were actually correct.

An F1 score of 0.85 indicates high but not perfect accuracy, meaning that not all the 679,584 posts that GPT-4o identified as antisemitic are actually so. More details can be found in Appendix: Methodology at the end of this report.

For readability, we sometimes use the phrase “antisemitic posts” to refer to all posts GPT-4o identified as antisemitic, accepting that the analysis is not perfectly accurate.

1 INTRODUCTION

For thousands of years, antisemitism has adapted to every new form of mass communication, spreading hate through the printing press, radio, television, and now, social media. Each wave has brought real-world harm, and today's digital platforms are no exception.

From the horrors of the Holocaust to the recent murders at Jewish events in D.C. and Colorado, we know that antisemitism and the conspiracy theories and lies at its core result in violence, intimidation, harassment, and discrimination in the real world. What has changed is the scale and speed at which it spreads and becomes normalized. Platforms like X enable antisemitic content to reach millions for zero cost to the producer, with little to no accountability.

The old conspiracy theories that simmered in the margins of society now thrive in plain sight, amplified by X's ineffective content moderation policies. In many cases, the platform not only tolerates this content but allows users to monetize it, giving antisemitic influencers both reach and revenue.

At a time when polarization, extremism, and violence are rising at home and abroad, the unchecked spread of antisemitism online is a direct threat to public safety. The platforms turn a blind eye to harmful content being spread on their sites, and they leave society to suffer the consequences.

Unless platforms change course, live up to their terms of service, and stop the spread of antisemitism and broader hate and extremism, it will likely, and sadly, lead to further violent incidents targeting our communities and our democracy. With Congress still debating whether to sunset Section 230's special immunity shield, it is even more crucial to advance a whole of society response to make investments that can protect communities, including advancing digital and media literacy and countering extremism.

Imran Ahmed, Founder and CEO, Center for Countering Digital Hate

Amy Spitalnick, CEO, Jewish Council for Public Affairs

2 EXECUTIVE SUMMARY

Antisemitism is one of the oldest forms of hate. The language, symbols, imagery, and arguments used to inculcate hate against Jewish people have proliferated for millennia, and persist today, even though the channels exploited to spread hate against Jewish communities have changed considerably in our increasingly digitized world. Throughout our work looking at identity-based hate, chaotic information ecosystems, and conspiracism at CCDH, in diverse areas such as the migration debate, climate denial and public health disinformation, we see clear echoes of anti-Jewish tropes such as the blood libel, control of our institutions, greed, intent to harm children, and Satanic conspiracies. Social media has become an incredibly potent vector for hatred against Jews and the lies that underpin hatred against Jews, allowing antisemitism to flourish in a way unseen for over a generation. X has repeatedly been brought up by Jewish communities as a space in which they see a lot of antisemitism, in which they feel unwelcome and unsafe.

Over the course of a year-long investigation into antisemitic content on X using AI tools in the form of OpenAI's GPT-4o, research uncovered systemic failures in content moderation, as well as platform amplification of antisemitic conspiracy theories, and the rise of influential, often monetized, accounts spreading anti-Jewish hate to a large audience.

Volume and Reach

Between February 1, 2024 and January 31, 2025, our study identified 679,584 posts as containing antisemitism likely to violate X's own policies on hateful conduct. These posts were viewed 193 million times, despite public commitments from X to reduce the visibility of hateful content. The platform's content moderation approach, which CCDH previously identified as insufficient, has not effectively limited the spread of this hateful material, suggesting that X is failing to enforce its own policies on antisemitism.ⁱ

Antisemitic Conspiracies

Antisemitic conspiracies are rampant on X. These conspiracies are not new and can be separated into three categories: Jewish control or power conspiracies, Jewish satanic conspiracies, and Holocaust denial. Of the total antisemitic posts, 59% were identified as being conspiracy theories. Despite this being just over half of the posts, they accounted for 73% of all likes. There is a clear pattern that posts promoting antisemitic conspiracies seem to be more likely to generate likes and engagement than other forms of antisemitic content.

Posts relating to Jewish control or power represented 30% of the yearly sample. Yet with 44% of the total likes and views, this category of conspiracy theory proved the most common and successful at securing engagement. These online conspiracies cannot be taken in isolation. They are linked to real-world harm. The FBI and extensive other research has warned that antisemitism is a “persistent driver” of violent extremism, with many attackers referencing the tropes in manifestos or online interactions.ⁱⁱ Alarming, polling indicates that teens who heavily use social media are more likely to support the Jewish power conspiracy.ⁱⁱⁱ

Ineffectiveness of Community Notes

X Community Notes system is a decentralized method for combating false information. This research shows that the system is failing in practice when it comes to antisemitism. By identifying the top 100 most-viewed posts from each antisemitic conspiracy theory category, reaching a total of 300 posts, we found that only 4 were given a publicly visible Community Note, representing just over 1%. This includes only 2 of the 100 top Holocaust denial posts being given a publicly-displayed Community Note. In total, the posts in this set of 300 without Community Notes achieved 86 million views.

Of the four posts with Community Notes, X data shows that they received 2.5 million views, while their accompanying notes received just 553,600 views. This indicates that Community Notes were visible for just 22% of views on these posts, due to delays arising from users writing and voting on notes before they get displayed to all users. Even when Community Notes were applied, they were not visible for over three quarters of views. Despite Elon Musk’s assurances that “a falsehood like Holocaust denial... can immediately be corrected” by Community Notes, this report finds negligible evidence of effective correction.^{iv} The system is not functioning at the speed or level required to address the reach of antisemitic content.

Antisemitism Influencers Profit on X

X has become the go-to platform for antisemitic posters. The platform makes it easy to monetize antisemitic content and reach millions of viewers. Of all the posts identified as antisemitic, ten individual “antisemitism influencers” account for 32% of total likes on posts in our sample. The other 68% of likes were from 159,055 users, displaying the disproportionate levels of influence that the top ten antisemitism influencers have.

For nine of these ten antisemitism influencers, X serves as their biggest platform in terms of the number of followers they have compared to other platforms. This indicates that X may offer these influencers access to a much larger audience than they are able to achieve on other social media platforms due to its features or its approach to content

moderation. Despite putting out content that violates X's own policies on antisemitism, 6 of the 10 antisemitism influencers are verified, meaning they are subscribed to X Premium. These subscribers receive a blue checkmark, a boost in their engagement and visibility across the platform, and crucially gain more opportunities to monetize content.^v The Jewish Council for Public Affairs found ads displayed near content posted by 5 of the 10 antisemitism influencers, whose content is estimated to earn X up to \$141,239 a year in ad revenue.

A Home for Hate

X has become a hotbed for online antisemitism to thrive. The findings in this report expose how X's levels of reach, failure of moderation, and its monetization model create a hospitable environment for antisemitic content. Antisemitic conspiracy theories perform well on the platform, reaching millions of views, and empower a class of antisemitism influencers who are usually bigger on X than anywhere else. Community Notes, X's primary mechanism for combating false information, is failing to stem the tide. These failures have profound implications – not just for online discourse, but for the safety of Jewish and all communities around the world.

3 KEY FINDINGS

Researchers used AI to identify 679,000 posts that violate X's policies on antisemitism

- Using OpenAI's GPT-4o model, researchers identified 679,584 English-language X posts containing antisemitism likely to violate X's rules with an accuracy of 85%.^{vi}
- The study analyzes X posts made from February 1, 2024, to January 31, 2025.^{vii}

Posts identified as antisemitic got 193 million views despite X promises to limit visibility^{viii}

- Posts identified as antisemitic by GPT-4o were viewed 193 million times and were placed into six categories of antisemitic conspiracy and abuse. Including:
 - "Jewish satanic blood letting ritual. They drink gentile blood."^{ix}
 - "The Holocaust is the biggest lie told in history."^x
 - "Money actually isn't the root of all evil. It's jews"^{xi}
 - "He is a fucking Jew He is a rat by nature"^{xii}
 - "he's not lying. jewish people own everything and got all the money."^{xiii}
 - "We need to fckng 🗡️ them they must not feel safe those 🗡️🗡️ must be eradicated"^{xiv}
- X has previously promised to remove hateful content or limit its visibility.^{xv}

X is failing to address the popularity of antisemitic conspiracies

- Antisemitic conspiracies appear to perform disproportionately well on X, constituting 59% of posts in our sample but 73% of likes.^{xvi}
- Elon Musk has claimed Community Notes ensure that "if somebody tries to push a falsehood like Holocaust denial, they can immediately be corrected."^{xvii}
- However, we found 4 in 300 posts promoting antisemitic conspiracies had a publicly visible Community Note, including just 2 in 100 Holocaust denial posts.^{xviii}

X has enabled the rise of "antisemitism influencers" not present on other platforms

- Ten "antisemitism influencers" accounted for 32% of likes on posts in our study.^{xix}
- These antisemitism influencers appear to enjoy unique benefits on X:
 - 9 in 10 have more followers on X than any other platform
 - 5 in 10 have had ads displayed their content
 - 6 in 10 have a verified 'blue checkmark'
 - 3 in 10 offer paid subscriptions on X
- X is estimated to earn up to \$141,239 a year from ads displayed near antisemitism influencers' content.
- While X has stripped perks such as verification from some users promoting antisemitism, this approach has not been applied consistently.

4 METHODOLOGY

Researchers used AI to label 679,000 posts as containing antisemitism violating X rules

To assess antisemitism on X, researchers studied 1.1 million English-language posts from February 1, 2024 to January 31, 2025, containing keywords relating to antisemitism.^{xx}

Researchers used OpenAI's GPT-4o model to identify 679,584 X posts from this broader sample that contain antisemitism likely to violate X's platform policies. Manual checks of a sample of 5,000 posts from our study found that GPT-4o had an accuracy score of 85%.^{xxi}

Our research focused on six types of antisemitism that violate X rules

Antisemitism comes in many forms and is always evolving. To produce an effective and accurate analysis, our research focuses on categories of antisemitism that appear to be common on X, that are likely to violate X policies and that pose real harm to Jewish people.

Researchers drew up a longlist of categories of antisemitic content, taking direction from resources like the International Holocaust Remembrance Alliance (IHRA) working definition of antisemitism, as well as research by scholars and organizations studying antisemitism.^{xxii}

They narrowed this longlist to just six categories listed in the table below, half of which they grouped as "antisemitic conspiracies" and the other half as "antisemitic abuse".

Keyword searches indicated that these categories have significant post volume on X and are likely to violate X's platform policies. Although the category of anti-Jewish violent speech had lower post volume, it was included due to the threat such content represents.

	Relevant X Policy		
Category of antisemitic content	Hateful Conduct ^{xxiii}	Abusive Behavior ^{xxiv}	Violent Content ^{xxv}
<i>Antisemitic conspiracies</i>			
Jewish control conspiracy theories	x		
Jewish satanic conspiracy theories	x		
Holocaust denial or distortion	x	x	
<i>Antisemitic abuse</i>			
Anti-Jewish character attacks	x		
Dehumanization of Jewish people	x		
Anti-Jewish violent speech	x	x	x

These categories represent a subset of the antisemitic content on X, therefore the findings of this report should not be taken as a comprehensive analysis of all antisemitic content on the platform. Importantly, our analysis does not specifically identify neo-Nazi or Israel-Palestine content that could be antisemitic, unless it falls into one of our six categories.

Researchers analyzed a year of English-language posts for analysis

For each of the six categories used in our analysis, researchers created a search query containing related terms. Researchers used these queries to study posts from February 1, 2024, to January 31, 2025, using a third-party tool with licensed access to X data. The full search queries can be viewed at the end of this report in Appendix: Methodology.

Our research studied only English-language posts with a location recorded as the US, UK, Canada, Australia, New Zealand, Ireland or without a stated location. All duplicates, reposts and quote posts were removed for this study.

Researchers used GPT-4o to classify six categories of antisemitic content in our sample

Researchers devised a prompt to guide OpenAI's GPT-4o model in evaluating posts. Performance was enhanced by using chain-of-thought prompting to improve the AI's reasoning by breaking down the categorization task into a series of intermediate steps.^{xxvi}

The prompt directed the large language model (LLM) through a series of steps to break down each post and discard posts with features such as sarcasm, quotes or ambiguous language. The LLM was directed to sort antisemitic posts into the six categories of antisemitism used in our report, or to exclude posts that did not fit into those categories.

To assess the LLM's performance, researchers manually labeled a sample of 5,000 posts and found that the model matched their labels 85% of the time, indicating a high level of reliability.^{xxvii} A breakdown of its performance can be found in Appendix: Methodology.



Our analysis is limited to text content and has limited data on post views

Our analysis is limited to the text content of posts on X, and therefore cannot identify posts containing antisemitic images, videos or audio. Another key limitation is that the third-party tool used to study posts does not provide view data for posts without engagement, making the total number of views for posts in our study a low estimate.

5 VOLUME AND VIEWS

Posts identified as antisemitism were viewed 193 million times

Posts identified as antisemitism likely to violate X's policies were viewed 193 million times, despite the platform's promises to remove hate or limit its visibility.^{xxviii}

Researchers used OpenAI's GPT-4o model to sort English-language X posts into six key categories of antisemitic content likely to violate X's policies. The model labeled a total of 679,584 posts as antisemitic, with an overall accuracy score of 85%.^{xxix}


X previously stated its "ongoing commitment to combat antisemitism", including automated content moderation and a "zero tolerance approach" to violent speech.^{xxx} The platform has also described its broader "enforcement philosophy" as "Freedom of Speech, not Freedom of Reach", which it describes "restricting the reach of [posts] that violate our policies by making the content less discoverable."^{xxxi}

Our analysis suggests that X is failing to limit the reach of antisemitic content, driven in part by a failure to tackle antisemitic conspiracies and platform decisions that have enabled the rise of "antisemitism influencers" who are analyzed later in this report.

Antisemitic conspiracies make up 59% of posts in our sample

Over 400,000 posts in our dataset were classified as antisemitic conspiracies, including claims about Jewish control or power, claims that Judaism is Satanic, and Holocaust distortion or denial.

■ Control Conspiracies ■ Satanic Conspiracies ■ Holocaust Denial ■ Character Attacks ■ Dehumanization ■ Violence



Antisemitic conspiracies, including claims about Jewish control or power, that Jews are Satanic, and Holocaust denial, make up 59% of posts in our study. Conspiracies about Jewish control were the single largest category representing 30% of posts.

Later sections of this report examine how conspiracies drive most likes of antisemitic content in our sample, and how X's Community Notes system of user-generated fact-checks may be failing to adequately address the spread of conspiracies.

ANTISEMITIC CONSPIRACIES
59% of posts labeled as antisemitic with 146.7 million views
Over half of our sample was classified as antisemitic conspiracies that are either conspiracies about Jewish people or blame them for negative world events. ^{xxxii}
Jewish control conspiracy theories
30% of posts with 84.7 million views
This category includes conspiracies that Jewish people possess excessive power or control over the world, media, finance or politics.^{xxxiii} Examples in our study include: • "he's not lying. jewish people own everything and got all the money."^{xxxiv} • "You know what the jewish-funded and controlled, lying corporate media and 'education' (mind programming system) taught you."^{xxxv} • "It doesn't matter which party comes to power. The puppet masters are always  and they want to mix whites out of existence..."^{xxxvi}
Jewish satanic conspiracy theories
23% of posts with 41.1 million views
This category includes false claims that Judaism is satanic, that Jewish people are children of the devil or that Jews worship the devil. It also includes claims that Jews participate in Satanic practices or that Jews drink the blood of non-Jews (known as "blood libel").^{xxxvii} Examples in our study include: • "Jewish satanic blood letting ritual. They drink gentile blood."^{xxxviii} • "It's not the left. It's not the right. It's not the globalists. It's not the communists. IT IS THE SYNAGOGUE OF SATAN."^{xxxix} • "The only terrorists I see evidence of on this earth are the ones referred to as the devils children  absolutely disgusting creatures"^{xl}
Holocaust denial or distortion
6% of posts with 20.9 million views
This category includes claims that deny, minimize or distort the historical truth of the Holocaust, including the number of Jewish people who were killed and how the Holocaust was carried out.^{xli} Examples in our study include: • "The Holocaust is the biggest lie told in history."^{xlii} • "...I saw a document on X stating after WW2  asked the usgov't to fudge theus  population numbers to go along with the holohoax fairytale."^{xliii} • "6 million was a made up number for the holohoax, real estimates are around 220k, most of which were from famine, also the whole gas chambers thing was fake."^{xliv}

ANTISEMITIC ABUSE
41% of posts labeled as antisemitic with 46.6 million views
Antisemitic abuse refers to hateful rhetoric that directly attacks Jewish people based on their identity through explicit degradation, dehumanization, and threats.
Anti-Jewish character attacks
28% of posts with 34.1 million views
This category includes claims that Jewish people or Judaism are evil, greedy, untrustworthy, liars, cowardly or have other character defects.^{xlv} Example posts in our study include: • “Jews all lie, it’s all they know how to do”^{xlvi} • “Money actually isn’t the root of all evil. It’s jews”^{xlvii} • “It is not that Jews are traitors, but rather that traitors are Jews.”^{xlviii}
Dehumanization of Jewish people
10% of posts with 9.2 million views
This category includes comparisons of Jewish people to animals, insects or suggestions that they are subhuman.^{xlix} Example posts in our study include: • “He is a fucking Jew He is a rat by nature”^l • “jews are the ones who push this on Europe. Vile parasites. jews are the problem.”^{li} • “It’s the Jews... 🦹‍♂️ That’s your snake 🐍”^{lii}
Anti-Jewish violent speech
3% of posts with 3.3 million views
This category consists of incitement, threats, or glorification of violence against Jewish people, including glorification of Hitler’s actions against the Jewish population in WWII.^{liii} Example posts in our study include: • “Hitler was Right about the filthy jews.”^{liv} • “I want to know who runs this account so we can get rid of him in person ..all Jews must be eradicated from the earth ...the most evil people in the world are Jews”^{lv} • “We need to fckng 🗡️ them they must not feel safe those 🦹‍♂️🦹‍♂️ must be eradicated”^{lvi}

6 ANTISEMITIC CONSPIRACIES

X is failing to address the popularity of antisemitic conspiracy theories

Antisemitic conspiracies appear to perform disproportionately well on X, receiving 73% of all likes in our sample despite representing just 59% of posts identified as antisemitic.

The disproportionate popularity of antisemitic conspiracies was particularly pronounced for conspiracies about Jewish control, which amassed 44% of likes while comprising 30% of all posts labeled as antisemitic. This pattern suggests that posts promoting antisemitic conspiracies are more likely to generate likes than other forms of antisemitic content.

Category of antisemitic content	Likes	Likes %	Posts	Posts %
Antisemitic conspiracies	3,087,068	73%	400,315	59%
<i>Jewish control conspiracy theories</i>	<i>1,855,046</i>	<i>44%</i>	<i>202,667</i>	<i>30%</i>
<i>Jewish satanic conspiracy theories</i>	<i>813,518</i>	<i>19%</i>	<i>154,101</i>	<i>23%</i>
<i>Holocaust denial or distortion</i>	<i>418,504</i>	<i>10%</i>	<i>43,547</i>	<i>6%</i>
Antisemitic abuse	1,156,357	27%	279,269	41%
<i>Anti-Jewish character attacks</i>	<i>818,390</i>	<i>19%</i>	<i>187,912</i>	<i>28%</i>
<i>Dehumanization of Jewish people</i>	<i>260,434</i>	<i>6%</i>	<i>70,599</i>	<i>10%</i>
<i>Anti-Jewish violent speech</i>	<i>77,533</i>	<i>2%</i>	<i>20,758</i>	<i>3%</i>
TOTAL OF BOTH CATEGORIES	4,243,425	100%	679,584	100%

Without access to X data, it is difficult to distinguish what causes antisemitic conspiracies to outperform some other forms of antisemitic content in receiving likes. It is possible that X users find antisemitic conspiracies more appealing, that X's algorithm is more likely to promote this content in users' feeds, or a combination of causes, with further research required to establish which causes are present and most influential.

However, the overall effect is that X is failing to address the popularity of antisemitic conspiracies despite claiming to prohibit "tropes or other content that intends to degrade or reinforce negative or harmful stereotypes about a protected category."^{lvii}

The following section of this report outlines how X's proposed solution to addressing false or misleading content – its Community Notes system of user-generated fact-checks – appears to be having little impact on antisemitic conspiracies.^{lviii}

Conspiracy theories are linked to real-world harm

Previous polling from CCDH indicates that US teens who use social media more heavily are more likely to express support for the conspiratorial claim that “Jewish people have a disproportionate amount of control over the media, politics and the economy.”^{lix}

These narratives have fueled violence against Jewish communities for centuries, from medieval pogroms to contemporary hate crimes.^{lx} Several studies find evidence that “stronger conspiracy beliefs are associated with heightened levels of violent support and intentions.”^{lxi} The FBI warned that antisemitism is a “persistent driver” of violent extremist attacks and went on to explain how several violent domestic extremists cited anti-Jewish conspiracy theories in their manifestos or online interactions.^{lxii}

7 EFFECTIVENESS OF COMMUNITY NOTES

Only 4 out of 300 posts promoting antisemitic conspiracies had a Community Note


Researchers found that just four out of 300 posts promoting antisemitic conspiracies had a publicly visible Community Note attached to them, equivalent to just over 1% of them. Posts without a Community Note received a total of 86 million views.

X data also shows that the four posts that did display Community Notes received 2,574,143 views, while their accompanying notes received just 553,600 views. This indicates that Community Notes were visible for just 22% of views on these posts, due to delays arising from users writing and voting on notes before they get displayed to all users.

Researchers identified the top 100 most-viewed posts from each category promoting antisemitic conspiracies, to give a total 300 posts. Researchers then used X's publicly available dataset on Community Notes to identify which of the selected 300 posts had a visible note, by matching the post ID to that of the Community Note.^{lxiii}


This analysis suggests that the Community Notes system is failing to correct even the most popular Holocaust denial posts, despite Musk claiming: "We've put maximum resources and attention behind Community Notes. So, if somebody tries to push a falsehood like Holocaust denial or something like that, they can immediately be corrected".^{lxiv}

Previous research from CCDH evidenced the limitations of Community Notes and their ineffectiveness in countering misinformation. 74% of accurate notes about election misinformation remained hidden from users, and even when notes are visible, misleading posts still receive significantly more views than the note accompanying them.^{lxv} The following screenshots are of example posts with no Community Note attached.^{lxvi}



X inconsistently applies its rules to antisemitic conspiracies

Out of the top 300 most viewed conspiratorial posts in our dataset, most had no moderation and only four had community notes.



In total, action of some kind had been taken on 36 out of the 300 posts promoting antisemitic conspiracies, equivalent to just 12% of posts in our sample.

- 21 posts had their visibility limited by X
- 6 posts had been removed, either by X or by the users posting them
- 4 posts displayed Community Notes
- 3 posts were from deleted accounts
- 2 posts were from suspended accounts

Although limiting the visibility of posts was the most common action that X took against antisemitic conspiracies, these posts still had a combined total of 2.8 million views.

X previously updated their policies based on the principle of “Freedom of Speech, Not Reach”.^{lxvii} As a result, the platform introduced “visibility filtering”, an enforcement action that limits the visibility of posts if they are identified as “potentially violating [X’s] policies”.^{lxviii} X trialed visibility filtering on posts potentially violating their Hateful Conduct policy, before expanding this to cover policies on Abusive Behavior and Violent Speech.^{lxix}

This suggests that X is rarely labelling posts as being subject to limited visibility, and even when posts are limited in this way, they can still accrue significant views.

8 ANTISEMITISM INFLUENCERS

Ten “antisemitism influencers” account for 32% likes on posts identified as antisemitic

Posts from just ten “antisemitism influencers” account for 32% of likes on posts identified by GPT-4o as containing antisemitism likely to violate X’s platform policies.

Researchers identified 159,055 users as authors of antisemitic posts in our study, ranking them according to the total number of likes they received across their posts. The top ten users in this ranking are the leading “antisemitism influencers” in our study.^{lxx}

In total, the top ten influencers amassed 1,344,908 likes across 675 posts. Examples of antisemitic posts from identified influencers can be found in at the end of this report.

This table lists the top ten antisemitism influencers alongside the following attributes:


- **Biggest on X:** denotes if the largest portion of an influencer’s total follower count is on X out of their identified accounts across other social media channels
- **Verified:** indicates whether an influencer is subscribed to X Premium and therefore verified on the platform, demonstrable by a blue checkmark on their profile^{lxxi}
- **Subscription:** denotes if the influencer offers followers a subscription service on X, which gives them share of the revenue made by X from said subscription^{lxxii}
- **Ads:** denotes whether staff working for the Jewish Council for Public Affairs (JCPA) were able to identify ads displayed near the account’s content.

Account Handle	Sum of Likes	Follower Count	Biggest on X	Verified	Ads	Subscriptions
realstewpeters	235,426	791,116	Yes	Yes*	Yes	Yes*
uncommonsince76	215,794	214,338	Yes	Yes	Yes	No
letsgobrando45	178,028	133,617	Yes	Yes	Yes	Yes
danbilzerian	150,333	1,988,172	No	Yes	No	No
jakeshieldsajj	140,249	839,949	Yes	No	No	No
aidenhunterx	119,874	55,819	Yes	No	No	No
theofficial1984	95,570	240,187	Yes	No	No	No
badazn	84,351	102,584	Yes	Yes	Yes	No
megatron_ron	63,870	469,261	Yes	Yes	No	Yes
shadowofezra	61,413	696,176	Yes	Yes	Yes	No
TOTAL	1,344,908	5,531,219				

** The account realstewpeters lost his verification and subscriptions at the end of December 2024, but had both for 11 months of the period studied^{lxxiii}*

9 out of 10 antisemitism influencers have more followers on X than any other platform

X served as the biggest platform for all but one of our influencers, when compared to both other mainstream and alternative platforms where they are active



For nine out of the 10 identified antisemitism influencers, X serves as their biggest platform in terms of follower numbers, serving as the central hub for their reach and influence.

Researchers identified accounts belonging to influencers on other platforms by performing searches using their X handles. Influencers were found to have accounts on other mainstream platforms, such as Facebook, and alternative platforms such as Rumble. For three influencers, researchers were unable to find accounts on other platforms.

Researchers calculated the total number of followers each influencer had across their social media accounts, before calculating the percentage of total followers that are on X.

The above chart shows that X appears to be the biggest platform for nine of the 10 influencers. This indicates that features of X, such as the perks it offers some antisemitism influencers or its approach to content moderation, may offer these influencers access to a much larger audience than they are able to reach on other social media platforms.

The exception to this finding was Dan Bilzerian, a former poker player turned 'masculinity influencer'.^{lxv} While just seven posts from Bilzerian were labelled as antisemitic, these posts had accrued 5.4 million views and 150,333 likes. This may be due to his large following on X, which is the highest of the influencers at nearly 2 million followers.

6 out of 10 antisemitism influencers are verified on X


Six of the 10 antisemitism influencers are verified on X meaning that they are subscribed to X Premium and have access to features such as a “reply boost” meaning their replies are prioritized in threads.^{lxxv} Uncommonsince76, ShadowofEzra, LetsGoBrando45, badazn, Megatron_ron, and Dan Bilzerian were identified as having verification on X.

X Premium subscribers receive a blue checkmark along with a boost in their engagement and visibility across the platform. Verified users receive prioritized placement in replies, ensuring their content appears more prominently in conversations.^{lxxvi} Users also benefit from tools that can help tailor their posts and boost engagement with their audience, including access to longer posts, video uploads and tools like Media Studio and X Pro.^{lxxvii}

To receive a blue checkmark, X states that accounts must fulfill requirements, including having no signs of being “misleading or deceptive” or “engaging in platform manipulation and spam”.^{lxxviii} Accounts can have their badge removed if they violate X’s rules, including posting content that breaches platform policies on hateful conduct and harassment.^{lxxix}

X’s inconsistent approach to removing verification from accounts

X reinstated conspiracy theorist and far-right broadcaster Stew Peters (realstewpeters) after Elon Musk’s takeover of the platform in December 2022.^{lxxx} Peters’ account was verified for most the period of our study before his verified status was removed in December 2024. Peters claimed “@elonmusk and @x @premium have removed my verified status, ALSO at the behest of the synagogue of satan.”^{lxxxix} Other antisemitism influencers are still verified, indicating this policy is being applied inconsistently.



5 in 10 antisemitism influencers have had ads displayed near their content

Staff working for the Jewish Council for Public Affairs (JCPA) identified ads displayed near X posts from five of the ten antisemitism influencers studied by this report.

X is estimated to earn up to \$141,239 a year from ads displayed near content from those influencers in the period studied, based on available data about how many times their content is viewed, the average rate at which X displays ads and industry information about the cost of displaying ads on X.

Staff working for the JCPA identified ads being displayed near antisemitic posts from the following five accounts: badazn, LetsGoBrando45, realstewpeters, ShadowOfEzra and Uncommonsince76. Examples of ads are available in **Appendix: Influencer Ads**.

X could be earning up to \$141,239 from ads near 'antisemitism influencers'

Estimates of ad revenue from the accounts are based on three elements: estimated post views, the frequency of ads on X and industry information on the cost of X ads.

- **Estimated Views:** Researchers found the five influencers generated an average of 10,370,458 views per day on their X posts and replies. This amounts to an estimated 3,785,217,146 views across a full year.^{lxxxii}
- **Ad Frequency:** In the report *Toxic Twitter*, researchers at the Center for Countering Digital Hate previously found that ads appeared at an average rate of once every 6.7 posts on X's main news feed.^{lxxxiii}
- **Ad Costs:** The social media analytics firm Brandwatch estimates that X ads cost \$0.25 per 1,000 views, a measure known as 'cost-per-mille'.^{lxxxiv}

Using these figures, researchers used the following formula to estimate that X could earn up to \$141,239 across a full year from ads displayed near the five influencers' content.

Arthur Kwon Lee reposted

Arthur Kwon Lee @badazn

You should hate the government more than the jews.

9:29 PM · Jul 6, 2025 · 57.9K Views

448 76 1.1K 30

Post your reply

Bluey The Bear Token Powered by Sol... @blueybear51... · 18h

Oh I do

154

Top News by CPAC @TopNewsbyCPAC

READ NOW: President Trump Announces 25% Tariff on Japan and South Korea — President Donald Trump on Monday sent letters to the leaders of Japan and South Korea announcing that each will face a 25 percent tariff on goods exported to the United...



From breitbart.com

1 4 8 5K

Stew Peters @realstewpeters

I'm more bullish than I've ever been on \$JPROOF.

Keep an eye out on your city, along the highway and ALL OVER THE COUNTRY. **

Nationwide billboards coming soon....



10:44 AM · Jun 28, 2025 · 26.9K Views

36 42 182 13

Post your reply

Cheeks @cheeksmeta · Jun 28

Will we see one in north west Ga ? 🤔

1 440

Freedom Forum @1stForAll

Yesterday, the Senate narrowly passed President Donald Trump's "One Big Beautiful Bill" by a 51-50 vote — with Vice President JD Vance breaking the tie.

Here's how newspapers across the country are covering the passage.




Ads displayed near antisemitism influencers' content on X

3 out of 10 antisemitism influencers profit from paid subscriptions on X

Three of the ten antisemitism influencers are Subscriptions creators, meaning they “get paid a share of the revenue X earns from Subscriptions” while offering subscribers “an extra level of access and bonus content.”^{lxxxv}

Brandon Taylor Moore (letsgobrand45) and anonymous users Megatron_ron and Uncommonsince76 are Subscriptions creators. Creators select the monthly price of their subscription from X’s price points and can then be “paid up to 97% of the revenue X has earned from selling access to [...] Subscriptions, after in-app purchase fees”.^{lxxxvi}



Researchers found that Moore and Uncommonsince76 offer Subscriptions at \$5 USD a month, while Megatron_ron charges \$3.^{lxxxvii} In total, the three accounts have 817,216 followers. Even if just a small portion of their followers signed up for their subscriptions, these influencers could be making a not-insignificant profit.

X’s inconsistent approach to monetization of accounts

Researchers identified two antisemitism influencers who say they were demonetized. TheOfficial1984 claims to have been demonetized in October 2024 for being in violation of the “Hate or Extremist Content” section of X’s policies, while Stew Peters claimed in December 2024 that “@elonmusk and @x STOLE my subscription ability in an effort to defund me and my network, at the command of his jewish masters.”^{lxxxviii}

Dan Bilzerian, has stated that he is not eligible for monetization on X and claimed the platform “datamined all the accounts that were antisemitic”, in some cases banning or limiting them.^{lxxxix} X appears to be inconsistent in enforcing its monetization policies.

APPENDIX: METHODOLOGY

The following section sets out the methodology for this study detailing the development of the six antisemitic categories, how researchers selected data for our study, prompt development and accuracy testing of the large language model used (GPT-4o).

Six categories of antisemitism were developed for study, informed by previous research

To develop categories for classification of antisemitic content, researchers were informed by the IHRA definition and work from external organizations that study antisemitism including the Institute for Strategic Dialogue, Community Security Trust, and the Anti-Defamation League to produce a longlist of antisemitic categories.^{xc}

Researchers then compared the longlist to X's policies on hateful conduct, abuse and harassment, and violent speech.^{xci} Categories that did not clearly breach X's policies were removed from the list.

Researchers created search queries for each category and used a third-party tool with licensed access to X data to estimate the volume of posts on X from February 2024 – January 2025 for each query. Queries were designed to target each category specifically and to be as broad as possible while leaving researchers with a manageable number of posts for analysis.

From this, five categories with the highest volume of posts were chosen for study. Despite a lower post volume, our study also included the 'Anti-Jewish violent speech' category due to the threat such content poses. Throughout this report, these six categories are grouped into two top-level categories: antisemitic conspiracies and antisemitic abuse.

Distinguishing between antisemitic posts and posts condemning Israel or Zionism can be a challenge in antisemitism research. Researchers carefully designed our methodology to minimize conflating political critique with hate speech. Specifically, researchers only studied posts from X containing terms referring to Jewish people and deliberately excluded keywords like "Israeli", "Israel", or "Zionist" from our search queries. As such, none of the categories selected were specifically related to the Israel-Palestine conflict.

Where posts contained references to Israel-Palestine, it was because they contained keywords relating to one of the six chosen categories. For posts containing such content, researchers prompted the model to only categorize the post as antisemitic if it contained antisemitism belonging to one of our six categories and therefore breached X's policies.

The following table sets out the selected six categories of antisemitism for study.

Category Label	Definition
Antisemitic Conspiracies	
Jewish control conspiracy theories	Conspiracies that Jewish people control or possess excessive power over the world, media, finance or politics
Jewish satanic conspiracy theories	Claims that Judaism is satanic, that Jewish people are children of the devil or that Jews worship the devil. It also includes claims that Jews participate in Satanic practices or that Jews drink the blood of non-Jews (known as "blood libel")
Holocaust denial or distortion	Claims that deny, minimize or distort the historical truth of the Holocaust, including the number of Jewish people who were killed and how the Holocaust was carried out
Antisemitic abuse	
Anti-Jewish character attacks	Claims that Jewish people or Judaism are evil, greedy, untrustworthy, liars, cowardly have other character defects
Dehumanization of Jewish people	Comparisons of Jewish people to animals, insects or suggestions that they are subhuman
Anti-Jewish violent speech	Incitement, threats, or glorification of violence against Jewish people, including glorification of Hitler's actions against the Jewish population in WWII
Excluded	
Other	Posts that are not antisemitic or do not fit into one of the categories above

Researchers identified 1.1m X posts for analysis

For each of the six antisemitic categories, researchers used a licensed third-party tool to study English-language posts tagged with a location of the UK, USA, Canada, Australia, New Zealand, Ireland or Unknown. Unknown locations were included as X users often choose not to disclose their location.

This gave an initial 3,951,874 X posts from February 1, 2024, to January 31, 2025, inclusive. To reduce the number of posts and therefore cost, researchers removed all reposts and quote posts, leaving 1,147,914 posts for analysis.

Researchers accessed data for our study February 7 and February 14, 2025. To give X ample time to moderate posts made in January 2025, researchers checked all January posts again on March 3, 2025. This ensured all posts in our final study had remained on the platform for at least one month.

The following table lists the search queries used to identify posts:

Search query
Jewish Control Conspiracy Theories
(🕊 OR ✡ OR synagogue OR jew* OR talmud OR j3w OR j3ws OR yid OR yids OR “(((they)))” OR “(((them)))” OR “(((he)))” OR “(((her)))” OR “(((she)))” OR “(((it)))” OR “(((his)))” OR gews OR hebes OR hebros OR heeb OR heebz OR holohoax OR hooknose* OR jewbag OR jewbags OR jewtard* OR kike OR kikes OR “oven dodger” OR “oven dodgers” OR “shekel nose” OR zionazi* OR kyke OR kykes OR zhid OR zhids OR zhyd OR zhyds OR khazar* OR j3w OR j3ws OR israheil OR jevvs OR joos) NEAR (steal* OR rul* OR control* OR dominat* OR power* OR enslav* OR manipul* OR NEAR (money OR financ* OR dollar* OR economy OR bank* OR media OR politic* OR soros OR rothschild* OR hollywood OR jewllywood) NOT (jewel* OR QT OR Politico OR “jew hat*” OR jew-hat*))
OR (🕊 OR ✡ OR synagogue OR jew* OR talmud OR j3w OR j3ws OR yid OR yids OR “(((they)))” OR “(((them)))” OR “(((he)))” OR “(((her)))” OR “(((she)))” OR “(((it)))” OR “(((his)))” OR gews OR hebes OR hebros OR heeb OR heebz OR holohoax OR hooknose* OR jewbag OR jewbags OR jewtard* OR kike OR kikes OR “oven dodger” OR “oven dodgers” OR “shekel nose” OR zionazi* OR kyke OR kykes OR zhid OR zhids OR zhyd OR zhyds OR khazar* OR j3w OR j3ws OR israheil OR jevvs OR joos) NEAR (cabal OR kabal OR crypto OR globalist OR puppet OR agenda OR “world order” OR “nwo” OR tentacle* OR mafia OR oligarch OR supremacy OR protocol*) NOT (jewel* OR jewelry OR QT OR Politico OR “jew hat*” OR “jew-hat*” OR hannibal)
Jewish Satanic Conspiracy Theories
(🕊 OR ✡ OR synagogue OR jew* OR talmud OR j3w OR j3ws OR yid OR yids OR “(((they)))” OR “(((them)))” OR “(((he)))” OR “(((her)))” OR “(((she)))” OR “(((it)))” OR “(((his)))” OR gews OR hebes OR hebros OR heeb OR heebz OR holohoax OR hooknose* OR kike OR kikes OR “oven dodger” OR “oven dodgers” OR “shekel nose” OR zionazi* OR kyke OR kykes OR zhid OR zhids OR zhyd OR zhyds OR khazar* OR j3w OR j3ws OR israheil OR jevvs OR joos) NEAR (satan* OR devil* OR lucifer* OR demon* OR horns OR horned OR baphomet) NOT (jewel* OR

QT OR "jew hat*" OR jew-hat* OR demonetis* OR demonetiz* OR demonstrat* OR "shoe-horn*" OR demonize OR demonise)
Holocaust Denial or Distortion
(fake* OR exaggerat* OR lie OR lies OR liar* OR lying OR hoax* OR fraud* OR sham OR false OR falsif* OR overstate* OR overestimate* OR inflate* OR embellish*) NEAR (holohoax OR hOI0hoax OR h010hoax OR holocaust OR "6 million" OR "6 gorillion" OR "anne frank" OR "zyklon" OR "gas chamber" OR "gas chambers" OR belsen OR auschwitz OR "elie wiesel" OR treblinka OR buchenwald OR dachau OR ravenbruck OR ravenbrück OR "six gorillion" OR "six million" OR hollowcaust OR hOI0caust OR h010caust OR shoah OR cremator* OR "otto frank") NOT (".6" OR QT OR jew-hat* OR "jew hat*" OR vote*)
Anti-Jewish Character Attacks
(🕍 OR 🕛 OR synagogue OR jew* OR talmud OR j3w OR j3ws OR yid OR yids OR "(((they)))" OR "(((them)))" OR "(((he)))" OR "(((her)))" OR "(((she)))" OR "(((it)))" OR "(((his)))" OR gews OR hebes OR hebros OR heeb OR heebs OR holohoax OR hooknose* OR kike OR kikes OR "oven dodger" OR "oven dodgers" OR "shekel nose" OR zionazi* OR kyke OR kykes OR zhid OR zhids OR zhyd OR zhyds OR khazar* OR j3w OR j3ws OR israheil OR jevvs OR joos) NEAR (greed* OR avari* OR evil OR cruel OR wicked OR vile OR foul OR decei* OR deceptive OR liar OR lie OR lies OR lying OR scum OR filthy OR corrupt* OR malevole* OR nefarious OR degener* OR sinister OR coward* OR craven* OR disgust* OR disloyal* OR untrustworthy OR traitor* OR "double cross*" OR treacher* OR perfid* OR "can't trust" OR "cant trust" OR "can't be trusted" OR "cant be trusted") NOT (jewel* OR QT OR hamas OR jew-hat* OR "jew hat*" OR trope*)
Dehumanization of Jewish people
(🕍 OR 🕛 OR synagogue OR jew* OR talmud OR j3w OR j3ws OR yid OR yids OR "(((they)))" OR "(((them)))" OR "(((he)))" OR "(((her)))" OR "(((she)))" OR "(((it)))" OR "(((his)))" OR gews OR hebes OR hebros OR heeb OR heebs OR holohoax OR hooknose* OR kike OR kikes OR "oven dodger" OR "oven dodgers" OR "shekel nose" OR zionazi* OR kyke OR kykes OR zhid OR zhids OR zhyd OR zhyds OR khazar* OR j3w OR j3ws OR israheil OR jevvs OR joos) NEAR (maggot* OR octop* OR snake* OR serpent* OR pig* OR dog* OR rat OR rats OR vermin OR roaches OR animal* OR reptil* OR termite* OR subhuman OR rodent* OR cockroach* OR parasit* OR monkey*) NOT (jewel* OR jewelry OR QT OR jew-hat* OR "jew hat*")
Anti-Jewish Violent Speech
(🕍 OR 🕛 OR synagogue yid OR yids OR "(((they)))" OR "(((them)))" OR "(((he)))" OR "(((her)))" OR "(((she)))" OR "(((it)))" OR "(((his)))" OR gews OR hebes OR hebros OR heeb OR heebs OR holohoax OR hooknose* OR kike OR kikes OR "oven dodger" OR "oven dodgers" OR "shekel nose" OR zionazi* OR "6 gorillion" OR "good goy" OR "good goys" OR kyke OR kykes OR zhid OR zhids OR zhyd OR zhyds OR khazar* OR jevvs OR jew* OR talmud OR j3w OR j3ws OR joo*) NEAR/5 ("should be" OR deserve* OR "need to be" OR "must be" OR "ought to" OR "we need to" OR "we have to" OR "we should" OR "right to" OR "hitler was right" OR "adolf was right" OR "austrian painter was right") NEAR/5 (kill* OR murder* OR burn* OR "hunt down" OR gas* OR "wipe out" OR exterminat* OR eliminat* OR punish* OR destroy* OR remove* OR slaughter* OR eradicate* OR nuked OR lynch* OR hung OR hang* OR "driven out" OR "drive out" OR rid OR cleanse* OR death OR execute* OR "final solution") NOT (jewel* OR muslim* OR arab* OR quran OR khan OR islam* OR QT OR hamas OR jew-hat* OR "jew hat*" OR gaslight)

Researchers took a 'Chain-of-Thought' approach to develop a prompt for the LLM

To develop prompts that maximized the model's performance, researchers required a labeled sample of posts for testing. A random sample of 3,630 posts was drawn from the data for labeling by the researchers, including posts from all six antisemitic categories and posts not belonging to any of the categories.

Where a post contained antisemitism, but did not fit into one of our six predefined categories, researchers labeled the post as "Does not fit any other category". Where a post fit multiple categories, they selected the category that best represented the content of the post. Once the test set was labelled, researchers began developing a prompt for the LLM.

Researchers instructed the language model to analyze each post methodically, employing chain-of-thought prompting to guide its reasoning process step by step. Chain-of-Thought (CoT) Prompting involves giving an LLM instructions on how to carry out the task it is set, much as a human annotator might be given a guide.^{xcii}

Researchers developed the prompt to contain detailed descriptions of each category and included examples as reference points for the model. As LLM models face particular challenges in interpreting complex linguistic features such as sarcasm, the prompt explicitly instructed the model to consider contradictions, tone (e.g. sarcasm and irony), coded language, and negations when evaluating posts.^{xciii}

Finally, the prompt instructed the model to classify content containing antisemitism according to the six antisemitic categories outlined in our report, while excluding posts that did not fall into our defined categories.

Researchers evaluated the LLM to have a final F1 score of 0.85

Once the LLM analyzed all 1,147,914 posts with the final prompt, a random sample of 5,000 posts were manually checked by researchers to get a final F1 score, an indicator of accuracy. The 5,000 post sample was spread across months and categories as follows.

Month	Number in Sample
February 2024	357
March 2024	358
April 2024	357
May 2024	367
June 2024	358
July 2024	364
August 2024	363
September 2024	357
October 2024	360
November 2024	361
December 2024	365
January 2025	1033

Model Label	Number in Sample
0	695
1	719
2	718
3	719
4	719
5	717

Researchers ran posts made in January before all others and checked a sample of 1033 of these to ensure that the prompt was performing as expected.

Each post was checked by two researchers who were blind to each other's assessment. Where there was disagreement, the posts were discussed by a team of three researchers until a consensus was reached. Once the posts were checked, researchers were able to calculate the overall performance of the LLM using the following four metrics:

1. **Recall:** The percentage of posts in a category that the model correctly identified. "Of all posts that belong to Category X, how many did our model find?"
2. **Precision:** The percentage of posts the model placed in a category that belongs there. "When our model says something is Category X, how often is it right?"
3. **F1 for each category:** A single score that balances precision and recall for a specific category. It is high only when both precision and recall are good.
4. **Weighted F1:** An overall performance score for the model across all categories, giving more importance to categories with more posts.

The weighted F1 for the model was 0.85, on a scale between 0 and 1, indicating strong but not perfect agreement between the LLM classifications and researcher evaluations. Where there was disagreement between researchers and the LLM:

- **False Positives:** 4.96% of posts did not fit any antisemitic category ('Other') according to researchers but were classified as an antisemitic category by the LLM.
- **Category Misclassification:** 9.86% the posts were assigned an antisemitic category, but researchers identified another category as more appropriate.
- **False Negatives:** 3.76% of posts were classified as not fitting an antisemitic category but researchers assigned a label for one of the antisemitic categories.

The F1 scores for individual categories are in the table below.

Category	Precision	Recall	F1 Score
Jewish control conspiracy theories	0.97	0.83	0.9
Jewish satanic conspiracy theories	0.93	0.97	0.95
Holocaust denial or distortion	0.85	0.99	0.91
Anti-Jewish Character attacks	0.94	0.69	0.8
Dehumanization of Jewish people	0.71	0.97	0.82
Anti-Jewish violent speech	0.83	0.98	0.9
Other (Does not fit any category)	0.73	0.67	0.7
Overall	0.86	0.85	0.85 (weighted)

How researchers came up with 679,584 X posts identified as antisemitic

After the LLM had completed analysis of all 1,147,914 posts, researchers checked all posts with over 250,000 views. 21 out of 127 posts were removed, as they did not fit into the selected six categories, to ensure that view figures were not overinflated by misclassifications.

Researchers also checked through each of posts from the identified top 10 “antisemitism influencers” and removed three out of 679 posts that did not fit into the selected six categories to ensure that likes figures for each influencer were not inflated.

These checks resulted in a final study of 679,584 posts that were labelled as antisemitic by the model for study, following the removal of 24 posts that manual checks identified as having been falsely labelled as antisemitic by the LLM.

How researchers checked the 300 most-viewed conspiratorial posts for Community Notes

Researchers took the top 100 most-viewed posts from each of the following categories promoting antisemitic conspiracy theories: Jewish control conspiracy theories, Jewish satanic conspiracy theories, Holocaust denial or distortion. This left a total of 300 posts.

Using X’s publicly available dataset on Community Notes, researchers were able to identify which of 300 posts had a visible note attached to them by extracting the unique tweetID numbers and matching them to that of the Community Note. From this, researchers identified four posts with a Community Note.

Researchers went through the rest of the 300 posts to review whether any action had been enforced on them, including visibility limitations, account suspensions and post removals.

How researchers identified the list of top 10 “antisemitism influencers”

Researchers identified a list of top 10 “antisemitism influencers” who received the most likes on their posts promoting antisemitism.

Researchers calculated the total number of likes amassed by each user across their posts. After identifying the top 10 users with the highest total number of likes, they then reviewed each post to check they were correctly labeled antisemitic and removed those posts which were not. This resulted in the final list of 10 “antisemitism influencers” who

had received a combined total of 1,344,908 likes across 675 posts in the studied timeframe.

How ads were identified near antisemitism influencers' content

Staff working for the Jewish Council for Public Affairs (JCPA) identified ads displayed near X posts from five of the ten antisemitism influencers studied by this report. To do this, they examined recent posts from all the antisemitism influencers, taking screenshots of where ads appeared next to their posts in newsfeeds or underneath their posts in reply threads.

JCPA staff identified ads displayed near content from five of the antisemitism influencers in total: badazn, LetsGoBrando45, realstewpeters, ShadowOfEzra and Uncommonsinc76. Examples of ads are available in **Appendix: Influencer Ads**.

Researchers then used the following data to generate an estimate of the ad revenue that all content from these accounts could generate for X across a year:

- **Estimated Views:** Researchers found the five influencers generated an average of 10,370,458 views per day across a total of 35,130 X posts and replies published between 8 May 2024 and 31 January 2025. This amounts to an estimated 3,785,217,146 views across a full year.^{xciv}
- **Ad Frequency:** In the report *Toxic Twitter*, researchers at the Center for Countering Digital Hate previously found that ads appeared at an average rate of once every 6.7 posts on X's main news feed.^{xcv}
- **Ad Costs:** The social media analytics firm Brandwatch estimates that X ads cost \$0.25 per 1,000 views, a measure known as 'cost-per-mille'.^{xcvi} This is equivalent to a cost of \$0.00025 per individual ad view.

Researchers divided estimated views by ad frequency, generating an estimate of how many views of the influencers' content were accompanied by ads. They then multiplied this figure by the estimated cost-per-view, arriving at an estimate that X could earn up to \$141,239 across a full year from ads displayed near the five influencers' content.

$$(3,785,217,146 \text{ views} / 6.7 \text{ ad frequency}) \times \$0.00025 \text{ cost per view} = \$141,239$$

APPENDIX: LIMITATIONS

This appendix sets out some of the key limitations of our methodology, enabling readers to place our findings in context. Importantly, the AI we used for our analysis is not perfectly accurate, and we were not able to analyze all forms of posts or antisemitism.

Analysis using AI is not perfectly accurate

Accurately identifying and categorizing antisemitic posts is a complex task, and while we used a large language model (LLM) to reach a high level of accuracy in this task, it was not perfectly accurate. This is typical when using AI to categorize large numbers of posts.^{xcvii}

Many posts in our study lacked view data

The licensed third-party tool used to identify X posts for study only provided view figures for posts with engagement (likes, reposts or replies). As a result, 63% of posts in our study were lacking view figures, making our measurement of the total number of views accrued by posts in our study an underestimate.

It was not possible to analyze images, video or audio

The methodology used for our study only allowed for analysis of the text content of posts. Posts that expressed antisemitism using other forms of media, such as images, videos and audio, were therefore only included in our study if they also had accompanying text featuring keywords related to antisemitism.

Our search queries do not cover the full spectrum of antisemitism


To provide an effective analysis and focus on categories of antisemitic post likely to violate X's platform policies, this report focuses on just six key categories of antisemitic conspiracy and abuse. As a result, there are many other categories of antisemitism, some of which are outlined in the International Holocaust Remembrance Alliance (IHRA) definition, that are not included in our analysis.^{xcviii}

Analysis used limited post text and lacked broader context

The licensed third-party tool used to identify X posts for study only retrieved a maximum of 280 characters of text from posts, focusing text featuring keywords in our query. Our LLM analysis was therefore blind to text exceeding this limit, as well as broader context such as an account's biography or post history. In manual checks to assess the accuracy of our analysis, researchers assessed posts based on the same limited text provided to the LLM without additional context.

APPENDIX: EXAMPLES OF ANTISEMITISM FROM LIST OF TOP 10 “ANTISEMITISM INFLUENCERS”

The following appendix lists examples of posts promoting antisemitism from each of the top 10 identified “antisemitism influencers”.

@aidenhunterx	
 2.7 million views; 33,642 likes https://x.com/AidenHunterX/status/1844159259778625747	 742,000 views; 14,864 likes https://x.com/AidenHunterX/status/1841822583895031919

@realstewpeters	
 694,000 views; 18,195 likes https://x.com/realstewpeters/status/1868652728131752110	 997,000 views; 14,290 likes https://x.com/realstewpeters/status/1798089964426572078

@LetsGoBrando45



174,000 views; 6786 likes

<https://x.com/LetsGoBrando45/status/1861429172360331384>



131,000 views; 4,491 likes

<https://x.com/LetsGoBrando45/status/1848468026502615486>

@Uncommonsince76



365,000 views; 10,651 likes

<https://x.com/Uncommonsince76/status/1878258376246718678>



626,000 views; 17,419 likes

<https://x.com/Uncommonsince76/status/1880607908372897869>


@TheOfficial1984

Over the last 2000 years, Judaism has not changed, nor have their practices. They pride themselves in this fact.

This means that sometime within the last 100 years, Christianity has changed. Today, people conflate the two religions with terms like "Judeo-Christian," implying that they are one in the same, or somehow embrace similar beliefs and values.

The truth is that Christianity has been subverted, infiltrated, manipulated, and bastardized.

The truth is, "Judeo-Christianity" is a LIE.



7:54 PM · Jun 29, 2024 · 1.3M Views

797

4.1K


12K

7.5K



1.4 million views; 13,024 likes

<https://x.com/TheOfficial1984/status/1807125569345106028>



10:27 PM · Jun 9, 2024 · 107.5K Views

32

255

2.5K


78



107,000 views; 2,582 likes

<https://x.com/TheOfficial1984/status/1799916181429661942>

@Megatron_ron



919,000 views; 14,995 likes


https://x.com/Megatron_ron/status/1854903150446043330



194,000 views; 5,720 likes

https://x.com/Megatron_ron/status/1874440402088194382

@badazn



738,000 views; 14,253 likes
<https://x.com/badazn/status/1862639631306527127>



451,000 views; 23,410 likes
<https://x.com/badazn/status/1883983750788440075>

@jakeshieldsajj




701,000 views; 12,000 likes
<https://x.com/jakeshieldsajj/status/1778583293979713916>




1.8 million views; 10,271 likes
<https://x.com/jakeshieldsajj/status/1787645019769418184>

@DanBilzerian	
	
2.6 million views; 35,338 likes https://x.com/DanBilzerian/statuses/1841925596521378057	183,000 views; 20,724 likes https://x.com/DanBilzerian/status/1883338558829936888

@ShadowofEzra	
	
241,000 views; 4,794 likes https://x.com/ShadowofEzra/status/1755822585227416014	7.1m views; 48,642 likes https://x.com/ShadowofEzra/status/1825214337306227171

APPENDIX: EXAMPLES OF ADS BY MONETIZED “ANTISEMITISM INFLUENCERS”

The follow appendix lists screenshots of ads served near the content of five of the ten “antisemitism influencers”, suggesting that their content is monetized.

@Badazn	
Arthur Kwon Lee reposted Arthur Kwon Lee @badazn You should hate the government more than the jews. 9:29 PM · Jul 6, 2025 · 57.9K Views 448 76 1.1K 30 Post your reply Reply Bluey The Bear Token Powered by Sol... @blueybear51... · 18h Oh I do Top News by CPAC @TopNewsbyCPAC Ad READ NOW: President Trump Announces 25% Tariff on Japan and South Korea — President Donald Trump on Monday sent letters to the leaders of Japan and South Korea announcing that each will face a 25 percent tariff on goods exported to the United...  From breitbart.com 1 4 8 5K	Arthur Kwon Lee @badazn Imagine having faith in the uniparty.  10:00 PM · Jul 6, 2025 · 4,574 Views 7 4 42 Post your reply Reply StonkyTheClown @StonkyTheClown · 18h This is the most Jewish thing I've seen all day. Danny Kalayanamit @Archer250 · 18h Jewriparty Just Me Joe @JozeT2 · 18h I have faith in only one!  FedMSB @FedMSB Ad Stablecoin = Innovation — it's a payment instrument. Legislation Alert: The GENIUS Act must recognize non-bank stablecoin issuance as regulated money transmission under both federal and state law.
Badazn, X, 6 July 2025, https://x.com/badazn/status/1942033077963743454	Badazn, X, 6 July 2025, https://x.com/badazn/status/1942040969815220467

@LetsGoBrando45



Brandon Taylor Moore
@LetsGoBrando45

Subscribe

We've got to do something about these satanic jews.

Last edited 1:11 PM · Jun 23, 2025 · 317K Views

417

1.9K

16K

245

Post your reply

Reply



Johnathan Charles · @JohnC45461211 · Jun 23

What we do is make more white people

10

55

4.2K



Brandon Taylor Moore · @LetsGoBrando45 · Jun 23

Practical!

1

29

4K



Supernova · @Supernova_SMCI

Discover how Supernova's AI solutions enhance network efficiency, reduce costs, and improve customer experiences in the TELCO industry.

Explore the Full Solution!



From supernova.com

5

46

240

2.8M

Post




Brandon Taylor Moore
@LetsGoBrando45

Subscribe

The geographical nation state of Israel 🇮🇱, formed by the satanic Rothschilds, IS FAKE & GAY. 🚫

Thought you should know. 🙄



2:28 PM · Jun 24, 2025 · 34.3K Views

58

430

1.4K

376

Post your reply

Reply



Yosef · @JustThisGuy2109 · Jun 25

And you call us schizophrenic...

You hate us but you really really want to be us.

3

6

165



Brandon Taylor Moore · @LetsGoBrando45 · Jun 25

No hate.

No jealousy either.

I know who I am in Christ. 🙄

1

6

147

Show replies



Intro · @useintro

Meet Intro, advisors for your business, on-demand.

- Access to 1K+ founders & execs
- Business advice via 1:1 video calls
- Businesses are seeing real ROIs

Our experts —> Founder of Reddit (\$10B), Founder of Zillow (\$12B), Founder of Sweetgreen (\$4B) and 1000+ more

LetsGoBrando45, X, 23 June 2025,
<https://x.com/LetsGoBrando45/status/1937196864518127755>

LetsGoBrandon45, X, 24 June 2025,
<https://x.com/LetsGoBrando45/status/1937578647369777348>

@ShadowofEzra

Shadow of Ezra
@ShadowofEzra

This is a clear sign we are living in the end times.



10:33 PM · Jun 15, 2025 · 2.9K Views

30 1.8K 10K 1.2K

Post your reply

Reply

Mario @MarioGomez1000 · Jun 17
Marrying a woman that was married to Anthony Weiner tells me all I need to know. Alex Soros is gay and this marriage is a sham.

30 1.8K 10K 1.2K

Aurelio @Aurelio1000 · Jun 16
Quite possibly the most Satanic couple in the world. Frazzledrip.



71 1.8K 10K 1.2K

Roman Gavi @RomanGavi · Jun 17
20 year old wealthy gay guy

marries a perimenopausal retired

what a joke

71 1.8K 10K 1.2K

Maverick Alexander @MaverickAlex · Jun 17

Alex Soros is 20, Maria is 45. Marriage is for creating a nuclear family. This wasn't a wedding in the traditional sense.

71 1.8K 10K 1.2K


Jim @JimMurphy7 · Jun 17

why is that middle aged, divorced woman wearing a white dress and veil.

71 1.8K 10K 1.2K

talkspace @talkspace

Members of the military and their families deserve fast access to private, convenient, and affordable mental health support. Sign up for Talkspace therapy, psychiatry, or teen therapy to receive high-quality care without delay, covered by TRICARE.



ShadowofEzra, X, 15 June 2025,
<https://x.com/ShadowofEzra/status/1934816658365009983>

Shadow of Ezra
@ShadowofEzra

New photos of Kathy Griffin are terrifying kids, with many saying she bears an uncanny resemblance to Pennywise the clown from IT.



3:40 PM · Jul 2, 2025 · 147.5K Views

400 496 2.1K 128

Post your reply

Reply

Gunther Eagleman™ @GuntherEagleman · Jul 2
I didn't know IT was real.

8 1 77 6.1K

Shadow of Ezra @ShadowofEzra · Jul 2
IT is everywhere now.

6 3 52 4.6K

Show replies

The Free Press @TheFP
Uri Berliner, a veteran at the public radio institution, says the network lost its way when it started telling listeners how to think.



ShadowofEzra, X 2 July 2025,
<https://x.com/ShadowofEzra/status/1940495908766126494>

@realstewpeters	
 Stew Peters @realstewpeters I'm more bullish than I've ever been on SJPROOF. Keep an eye out on your city, along the highway and ALL OVER THE COUNTRY. 🇺🇸 Nationwide billboards coming soon....  10:44 AM · Jun 28, 2025 · 26.9K Views 36 42 182 13 Post your reply Reply Cheeks @cheeksmeta · Jun 28 Will we see one in north west Ga? 🤔 Freedom Forum @1stForAll Yesterday, the Senate narrowly passed President Donald Trump's "One Big Beautiful Bill" by a 51-50 vote — with Vice President JD Vance breaking the tie. Here's how newspapers across the country are covering the passage.	 Stew Peters @realstewpeters Whites will always be the brightest light in any society, and nons will always hate us, because they cant be us. Instead, they ought to thank us for the honorable privilege to live near us, in the civilizations our people built.  9:49 PM · Jun 25, 2025 · 22.8K Views 86 105 488 42 Post your reply Reply Stew Peters @realstewpeters · Jun 25  From amazon.com 6 4 38 10K AFPM @AFPMonline The U.S. refining industry produces the fuels that drive our economy, safeguarding America from global energy instability and disruption.
Realstewpeters, X, 28 June 2025, https://x.com/realstewpeters/status/1938971820171956234	Realstewpeters, X, 25 June 2025, https://x.com/realstewpeters/status/1938052046407799289

@Uncommonsince76

Uncommon Sense
@Uncommonsince76

Subscribe

When Jews run the music industry, you have to pretend that this is Art...



9:42 PM · Jun 29, 2025 · 34.5K Views

50 133 1.3K 39

Post your reply

Reply

John H. Eberhart · @JohnHEberhart · Jun 29
Jewelry.

1 6 726

Omega News · @satch.omega · Jun 29

Omega News · @satch.omega · Apr 10

Replying to @kanyewest

Record executive Lyor Cohen (who has been in the music scene since the 80s) on why he signs rappers that promote drug culture | 2018

The answer: "I got people to feed. I have a business to run."



1 8 990

U.S. Women for Freedom · @LisaLoc31631820 · Jun 30
And sadly they fork over their hard earned money

53

#NephosWhoLead · @DaiWitaDoc · Ad
Every kidney care journey is unique. That's why our clinical teams focus on transformative, compassionate treatment, powered by innovation, and focus on personalized outcomes.



0:01 / 0:30

Uncommonsince76, X, 29 June 2025,
<https://x.com/Uncommonsince76/status/1939499760303661488>

Uncommon Sense
@Uncommonsince76

Subscribe

Can we go back to New York being a battle between the Irish and the Anglos, and not the Muslims and Jews?

We didn't know how good we had it.



4:09 PM · Jun 28, 2025 · 499.5K Views

209 966 21K 933

Post your reply

Reply

NatSock · @mynatsockacct · Jun 28
This is a movie. During this period in NYC there was approx. 1 murder per month.

12 12 629 19K

Uncommon Sense · @Uncommonsince76 · Jun 28
Total Aryan victory.

8 1 1K 17K

Show replies

Doom Scrollin' · @DoomerScrollah · Jun 28
Why would we want more brother wars?

5 1 100 12K

Uncommon Sense · @Uncommonsince76 · Jun 28
We don't.

2 1 287 11K

Show replies

Molly M. · @mollistan · Jun 28
We don't know how good we had it? There was sewage and rotting horses and dogs all over the place and the air was even less breathable than today.

20 1 58 14K

Uncommon Sense · @Uncommonsince76 · Jun 28
Modern technology has advanced but White men are killing themselves at a higher rate than in the Great Depression.

Not everything in life comes down to comfort. Purpose, Family, Being are more important.

5 4 500 12K

Show replies

Disappointed Millennial · @2008disappoint · Jun 29
I wonder if the Jews deliberately turned the Irish and the English against one another?

3 1 9 3.6K

FedMSB · @FedMSB
Stablecoin = Innovation — it's a payment instrument.

Legislation Alert:
The GENIUS Act must recognize non-bank stablecoin issuance as regulated money transmission under both federal and state law.

Read our letter to Congress: fedmsb.org/twitter



Uncommonsince76, X, 28 June 2025,
<https://x.com/Uncommonsince76/status/1939053614519197944>

REFERENCES

ⁱ “X Content Moderation Failure”, Center for Countering Digital Hate, 13 September 2023, <https://counterhate.com/research/twitter-x-continues-to-host-posts-reported-for-extreme-hate-speech/>

“Hate Pays”, Center for Countering Digital Hate, 11 April 2024, <https://counterhate.com/research/hate-pays/>

“Rated Not Helpful”, Center for Countering Digital Hate, 30 October 2024, <https://counterhate.com/research/rated-not-helpful-x-community-notes/>

ⁱⁱ “Antisemitism a Persistent Driver of Transnational Violent Extremist Narratives and Attacks”, Counterterrorism division, Federal Bureau of Investigation, 25 May 2023, <https://www.fbi.gov/file-repository/fbi-executive-update-antisemitism-a-persistent-driver-of-transnational-violent-extemist-narratives-and-attacks-052523.pdf/view>

ⁱⁱⁱ “Public Support for Social Media Reform”, Center for Countering Digital Hate, 16 August 2023, <https://counterhate.com/research/public-support-for-social-media-reform-star/>

^{iv} “If somebody tries to push a falsehood like Holocaust denial, they can immediately be corrected.”

“LIVE: Ben Shapiro interviews Elon Musk at European Jewish Association conference”, Associated Press, 22 January 2024, <https://youtu.be/kmlSAHRCxqc?t=702>

^v “Upgrade to Premium”, X, accessed 4 April 2025, https://x.com/i/premium_sign_up

^{vi} Researchers used F1 to assess the performance of the model, a type of accuracy measure that calculates the weighted average across ‘precision’ and ‘recall’ for each of our seven categories. The overall F1 score for the LLM was 0.85, on a scale between 0 and 1. Further detail and definitions can be found in Appendix: Methodology.

^{vii} The sample contained 1,147,914 posts in total, rounded to 1.1 million as presented in this section.

^{viii} More detail on this finding is available in section: 4 Volume and reach. Additionally 193 Million is an underestimate because the third party tool we used did not offer view figures for posts with no engagement. More detail on this can be found in Appendix: Methodology.

^{ix} never forget the USS Liberty 🇺🇸, X, 8 August 2024, https://x.com/goysagainst_zOg/status/1821532890393399743

^x Uncensored Reality, X, 24 July 2024, <https://x.com/Mo7549080322765/status/1816045626615886267>

^{xi} Apolitical, X, 15 September 2024, <https://x.com/Apolitical3678/status/1835333298165268619>

^{xii} Bocchi the Rock, X, 23 October 2024, https://x.com/Bocchi_Noor/status/1849212749911638503

^{xiii} David, X, 2 August 2024, <https://x.com/LaHeated/status/1819497858816790727>

^{xiv} Sionarerats, X, 7 June 2024, <https://x.com/sionarerats/status/1799146641255195114>

-
- xv “Hateful Conduct”, X Help Center, April 2023, <https://help.x.com/en/rules-and-policies/hateful-conduct-policy>
- xvi More detail on this finding is available in section: 5 Antisemitic conspiracies
- xvii “LIVE: Ben Shapiro interviews Elon Musk at European Jewish Association conference”, Associated Press, 22 January 2024, <https://youtu.be/kmlSAHRCxqc?t=702>
- xviii More detail on this finding is available in section: 4 Volume and reach
- xix More detail on this finding is available in section 5
- xx The sample contained 1,147,914 posts in total, rounded to 1.1 million as presented in this section. The keywords related to antisemitism can be found in Appendix: Methodology
- xxi Researchers used F1 to assess the performance of the model, a type of accuracy measure that calculates the weighted average across ‘precision’ and ‘recall’ for each of our seven categories. The overall F1 score for the LLM was 0.85, on a scale between 0 and 1. Further detail and definitions can be found in Appendix: Methodology.
- xxii “Working definition of antisemitism”, International Holocaust Remembrance Alliance, accessed 23 January 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>
- Researchers also drew upon other work defining antisemitism, including but not limited to: “Antisemitism Uncovered”, Anti-Defamation League, 11 March 2020, <https://antisemitism.adl.org/>
- “Antisemitism on Twitter Before and After Elon Musk’s Acquisition”, 20 March 2023, Institute of Strategic Dialogue, <https://www.isdglobal.org/isd-publications/antisemitism-on-twitter-before-and-after-elon-musks-acquisition/>
- “Examining antisemitic and anti-Muslim hate speech on Canadian YouTube channels in the aftermath of 7 October”, 17 June 2024, Institute for Strategic Dialogue, https://www.isdglobal.org/digital_dispatches/examining-antisemitic-and-anti-muslim-hate-speech-on-youtube-in-the-aftermath-of-october-7/
- “Decoding Antisemitism: An AI-driven study on hate speech and imagery online”, Discourse Report 6, February 2024, Center for Research on Antisemitism, <https://decoding-antisemitism.eu/publications/sixth-discourse-report/>
- “Decoding Antisemitism: A Guide to Identifying Antisemitism Online”, Becker et al., 2024, <http://dx.doi.org/10.1007/978-3-031-49238-9>
- “Working definition of antisemitism”, International Holocaust Remembrance Alliance (IHRA), <https://holocaustremembrance.com/resources/working-definition-antisemitism>
- “Categories of antisemitic incidents”, Community Security Trust, <https://cst.org.uk/antisemitism/what-is-antisemitism>
- xxiii “Hateful Conduct”, X, April 2023, <https://help.x.com/en/rules-and-policies/hateful-conduct-policy>
- xxiv “Abuse and Harassment”, X, March 2024, <https://help.x.com/en/rules-and-policies/abusive-behavior>
- xxv “Violent Content”, X, February 2025, <https://help.x.com/en/rules-and-policies/violent-content>

^{xxvi} "Chain-of-thought prompting elicits reasoning in large language models", Jason Wei, Xuezhi Wang, DaleSchuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou, Advances in neural information processing systems, 6 December 2022, https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf

^{xxvii} Researchers used F1 to assess the performance of the model, a type of accuracy measure that calculates the weighted average across 'precision' and 'recall' for each of our seven categories. The overall F1 score for the LLM was 0.85, on a scale between 0 and 1. Further detail and definitions can be found in Appendix: Methodology.

^{xxviii} "Hateful Conduct", X Help Center, April 2023, <https://help.x.com/en/rules-and-policies/hateful-conduct-policy>

This is an underestimate of views because the third-party tool used to study posts does not provide view data for posts without engagement.

^{xxix} Researchers used F1 to assess the performance of the model, a type of accuracy measure that calculates the weighted average across 'precision' and 'recall' for each of our seven categories. The overall F1 score for the LLM was 0.85, on a scale between 0 and 1. Further detail and definitions can be found in Appendix: Methodology.

^{xxx} "Our ongoing commitment to combat antisemitism on X", X, 8 September 2023, https://blog.x.com/en_us/topics/company/2023/supporting-peoples-right-to-accurate-and-safe-political-discourse-on-x/

^{xxxi} "Freedom of Speech, Not Reach: An update on our enforcement philosophy", X Safety, 17 April 2023, https://blog.x.com/en_us/topics/product/2023/freedom-of-speech-not-reach-an-update-on-our-enforcement-philosophy

^{xxxii} "Conspiracy Theory", American Jewish Committee, Accessed 10 April 2025, <https://www.ajc.org/translatehate/conspiracy-theory>

^{xxxiii} "Antisemitism Uncovered: Power", Anti-Defamation League, 11 March 2020, <https://antisemitism.adl.org/power/>

"Antisemitism on Twitter Before and After Elon Musk's Acquisition", 20 March 2023, Institute of Strategic Dialogue, <https://www.isdglobal.org/isd-publications/antisemitism-on-twitter-before-and-after-elon-musks-acquisition/>

"Working definition of antisemitism", International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

"Working definition of antisemitism", International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

^{xxxiv} David, X, 2 August 2024, <https://x.com/LaHeated/status/1819497858816790727>

^{xxxv} Stew Peters, X, 10 November 2024, <https://x.com/realstewpeters/status/1855701365432909963>

^{xxxvi} Zaburadon, X, 28 December 2024, <https://x.com/Zaburadon/status/1873144676360085968>

^{xxxvii} “Antisemitism Uncovered: Blood”, Anti-Defamation League, 11 March 2020, <https://antisemitism.adl.org/blood/>

“Antisemitism Uncovered: Deicide”, Anti-Defamation League, 11 March 2020, <https://antisemitism.adl.org/deicide/>

“Working definition of antisemitism”, International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

“Working definition of antisemitism”, International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

^{xxxviii} never forget the USS Liberty 🇺🇸, X, 8 August 2024, https://x.com/goysagainst_zOg/status/1821532890393399743

^{xxxix} Morgan Ariel, X, 23 August 2024, <https://x.com/itsmorganariel/status/1826761684683555208>

^{xl} Weiße Frau, X, 11 November 2024, https://x.com/beef_terfy/status/1856005695591837980

^{xli} “Antisemitism Uncovered: Denial”, Anti-Defamation League, 11 March 2020, <https://antisemitism.adl.org/denial/>

“Working definition of antisemitism”, International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

^{xlii} Uncensored Reality, X, 24 July 2024, <https://x.com/Mo7549080322765/status/1816045626615886267>

^{xliii} JOHN-MF-WICK-AMK-47, X, 26 July 2024, <https://x.com/SCHLONGWICK4769/status/1816875367459606841>

^{xliv} johnny5642, X, 10 December 2024, <https://x.com/RiverToSeaPs/status/1866571703759016120>

^{xlvi} “Cowardice”, Translate Hate, American Jewish Committee, Accessed 15 April 2025, <https://www.ajc.org/translatehate/cowardice>

“Dual Loyalty”, Translate Hate, American Jewish Committee, Accessed 15 April 2025, <https://www.ajc.org/translatehate/dual-loyalty>

“Greed”, Translate Hate, American Jewish Committee, Accessed 15 April 2025, <https://www.ajc.org/translatehate/greed>

“Working definition of antisemitism”, International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

^{xlvi} Simulated Living, X, 3 January 2025, <https://x.com/Smokiediebear/status/1882217320854262071>

^{xlvi} Apolitical, X, 15 September 2024, <https://x.com/Apolitical3678/status/1835333298165268619>

-
- xlvi G.K. Chesterton, X, 30 April 2024, <https://x.com/GKCdaily/status/1785083068191039592>
- xlvi "Creatures", Translate Hate, American Jewish Committee, Accessed 15 April 2025, <https://www.ajc.org/translatehate/creatures>
- "Working definition of antisemitism", International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>
- lv Bocchi the Rock, X, 23 October 2024, https://x.com/Bocchi_Noor/status/1849212749911638503
- lvi NoForeignAid, X, 7 October 2024, <https://x.com/NoForeignAid03/status/1843075715866177804>
- lvii Que, X, 28 September 2024, <https://x.com/Qu1776/status/1840091528129044728>
- lviii "Working definition of antisemitism", International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>
- lix Jorge Arnaldo, X, 4 January 2025, <https://x.com/JorgeACE24/status/1875361279734624666>
- lx Grant, X, 10 June 2024, <https://x.com/Grantgodfreybx/status/1800216852737790416>
- lxi Sionarerats, X, 7 June 2024, <https://x.com/sionarerats/status/1799146641255195114>
- lxii "Hateful Conduct Policy", X Help Center, April 2023, <https://help.x.com/en/rules-and-policies/hateful-conduct-policy>
- lxiii "LIVE: Ben Shapiro interviews Elon Musk at European Jewish Association conference", Associated Press, 22 January 2024, <https://youtu.be/kmlSAHRCxqc?t=702>
- lxiv "Public Support for Social Media Reform", Center for Countering Digital Hate, 16 August 2023, <https://counterhate.com/research/public-support-for-social-media-reform-star/>
- lxv "The Violent Impact of Anti-Semitic Conspiracy Theories: Examining the Jewish World Domination Narratives and History", Michael S. Broschowitz, Center on Terrorism, Extremism, and Counterterrorism, 6 May 2022, <https://www.middlebury.edu/institute/academics/centers-initiatives/ctec/ctec-publications/violent-impact-anti-semitic-conspiracy>
- "Understanding the Origins of Antisemitism", American Jewish Committee, 22 March 2022, <https://www.ajc.org/news/understanding-the-origins-of-antisemitism>
- lxvi "Linking Conspiracy Beliefs with Violence: A Scoping Review of the Empirical Literature", Emma Belton, Tiahna Mulholland, Kristina Murphy, Terrorism and Political Violence, 6 March 2025, <https://doi.org/10.1080/09546553.2025.2462601>
- "The relationship between conspiracy theory beliefs and political violence", Adam Enders, Casey Klostad, and Joseph Uscinski, Harvard Kennedy School Misinformation Review, 12 December 2024, <https://misinforeview.hks.harvard.edu/article/the-relationship-between-conspiracy-theory-beliefs-and-political-violence/>
- "Conspiracy theories continue to mobilize extremists to violence", Cody Zoschak &

Katherine Keneally, Institute for Strategic Dialogue, 26 June 2024, https://www.isdglobal.org/digital_dispatches/conspiracy-theories-continue-to-mobilize-extremists-to-violence/

lxii "Antisemitism a Persistent Driver of Transnational Violent Extremist Narratives and Attacks", Counterterrorism division, Federal Bureau of Investigation, 25 May 2023, <https://www.fbi.gov/file-repository/fbi-executive-update-antisemitism-a-persistent-driver-of-transnational-violent-extemist-narratives-and-attacks-052523.pdf/view>

lxiii "Downloading data – Community Notes", X, accessed 31 March 2025, <https://communitynotes.x.com/guide/en/under-the-hood/download-data>

lxiv "LIVE: Ben Shapiro interviews Elon Musk at European Jewish Association conference", Associated Press, 22 January 2024, <https://youtu.be/kmlSAHRCxqc?t=702>

lxv "Rated Not Helpful", The Center for Countering Digital Hate, 30 October 2024, <https://counterhate.com/research/rated-not-helpful-x-community-notes/>

lxvi Brandon Taylor Moore, X, 28 December 2024, <https://x.com/LetsGoBrando45/status/1873040733923590377>

EvanAKilgore, X, 24 October 2024, <https://x.com/EvanAKilgore/status/1849583132892168441>

AbhorrentBeing, X, 27 January 2025, <https://x.com/AbhorrentBeing/status/1883989566350319799>

lxvii "Freedom of Speech, Not Reach: An update on our enforcement philosophy", X, 17 April 2023, https://blog.x.com/en_us/topics/product/2023/freedom-of-speech-not-reach-an-update-on-our-enforcement-philosophy

lxviii "Freedom of Speech, Not Reach: An update on our enforcement philosophy", X, 17 April 2023, https://blog.x.com/en_us/topics/product/2023/freedom-of-speech-not-reach-an-update-on-our-enforcement-philosophy

lxix "Freedom of Speech Not Reach: New Updates and Progress", X, 12 July 2023, https://blog.x.com/en_us/topics/product/2023/freedom-of-speech-not-reach--new-updates-and-progress

lxx Researchers manually checked posts from the top ten users ranked by likes in order to ensure that the ranking was not influenced by posts wrongly identified as antisemitic by the large language model (LLM) used for the analysis. Further details are available in Appendix: Methodology.

lxxi "How to get the blue checkmark on X", X Help Center, 2023, <https://help.x.com/en/managing-your-account/about-x-verified-accounts>

lxxii "About Subscriptions creators", X, accessed 28 March 2025, <https://help.x.com/en/using-x/subscriptions-creator>

lxxiii realstewpeters, X, 31 December 2024, <https://x.com/realstewpeters/status/1874157766971187498>

lxxiv "How 'masculinity' influencer Dan Bilzerian fell down a 'brazenly antisemitic' rabbit hole", The Jewish Chronicle, 14 November 2024, <https://www.thejc.com/news/features/how->

[masculinity-influencer-dan-bilzerian-fell-down-a-brazenly-antisemitic-rabbit-hole-does360z](#)

lxxv "Upgrade to Premium", X, accessed 4 April 2025, https://x.com/i/premium_sign_up

lxxvi "Upgrade to Premium", X, accessed on 14 April 2025, https://x.com/i/premium_sign_up

lxxvii "Upgrade to Premium", X, accessed on 14 April 2025, https://x.com/i/premium_sign_up

lxxviii "How to get the blue checkmark on X", X, accessed 4 April 2025, <https://help.x.com/en/managing-your-account/about-x-verified-accounts>

lxxix "X Verification requirements", X, accessed 4 April 2025, <https://help.x.com/en/managing-your-account/about-x-verified-accounts>

"The X Rules", X, accessed 4 April 2025, <https://help.x.com/en/rules-and-policies/x-rules>

lxxx "Conspiracy Kingpin Shows Twitter Is a Safe Space for Antisemitism", Rolling Stone, 29 November 2023, <https://www.rollingstone.com/politics/politics-features/stew-peters-twitter-safe-space-antisemitism-1234905199/>

lxxxi Stew Peters, 31 December 2024, X, <https://x.com/realstewpeters/status/1874157766971187498>

Stew Peters, 31 December 2014, X, <https://x.com/realstewpeters/status/1874158151811088608>

lxxxii Researchers gathered views data for 35,130 posts and replies made by the accounts between 8 May 2024 and 31 January 2025, amounting to 2,789,653,184 total views or an average of 10,370,458 a day.

lxxxiii "To find an estimate for how often Twitter shows ads, the Center set up three accounts to follow only reinstated accounts and found that after scrolling through 1,039 tweets, an ad appeared at an average rate of once every 6.7 tweets."

"Toxic Twitter: How Twitter Generates Millions in Ad Revenue by Bringing Back Banned Accounts", Center for Countering Digital Hate, 9 February 2023, https://counterhate.com/wp-content/uploads/2023/02/Toxic-Twitter_FINAL.pdf

lxxxiv "Essential Guide to X Advertising: Strategies and Cost Insights", Brandwatch, 1 November 2024, <https://www.brandwatch.com/blog/x-ads/>

lxxxv "About Subscriptions creators", X, accessed 28 March 2025, <https://help.x.com/en/using-x/subscriptions-creator>

lxxxvi "About Subscriptions creators", X, accessed 28 March 2025, <https://help.x.com/en/using-x/subscriptions-creator>

lxxxvii Creator Subscriptions: LetsGoBrando45, X, accessed 28 March 2025, <https://x.com/LetsGoBrando45/creator-subscriptions/subscribe>

Creator Subscriptions: Uncommonsince76, X, accessed 28 March 2025, <https://x.com/Uncommonsince76/creator-subscriptions/subscribe>

Creator Subscriptions: Megatron_ron, X, accessed 2 April 2025, https://x.com/Megatron_ron/creator-subscriptions/subscribe

lxxxviii TheOfficial1984, X, 7 October 2024, <https://x.com/TheOfficial1984/status/1843367264478773588>

Realstewpeters, X, 31 December 2024,

<https://x.com/realstewpeters/status/1874157089721040999>

^{lxxxix} “But then they basically datamined all the accounts that were antisemitic and they f***** mass-banned like all the Groyper accounts, they removed their checkmarks, they can’t go live, they can’t post tweets that are more than 280 characters or whatever, and they actually showed that he [Elon Musk] banned more accounts than they had banned previously.”

Dan Bilzerian, DrLoupis__, X, 23 February 2025, 0:20–0:38,

https://x.com/DrLoupis__/_/status/1893763899712753984

“Well, I mean they demonetized Jake Shields, they demonetized a lot of people on Twitter as well. I mean I’m not even eligible for it.”

Dan Bilzerian, DrLoupis__, X, 23 February 2025, 2:25–2:36,

https://x.com/DrLoupis__/_/status/1893763899712753984

Note: Groypers are self-described ‘American nationalists’ led by far-right figure Nick Fuentes and are opposed to liberal values, feminism, LGBTQ+ rights and progressive ideology. Source: “Groypers”, Institute for Strategic Dialogue, accessed 4 April 2025, <https://www.isdglobal.org/explainers/groypers/>

^{xc} “Working definition of antisemitism”, International Holocaust Remembrance Alliance, accessed 26 March 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>

“Antisemitism Uncovered”, Anti-Defamation League, 11 March 2020,

<https://antisemitism.adl.org/35>

“Antisemitism on Twitter Before and After Elon Musk’s Acquisition”, 20 March 2023,

Institute of Strategic Dialogue, <https://www.isdglobal.org/isd-publications/antisemitism-on-twitter-before-and-after-elon-musks-acquisition/>

“Examining antisemitic and anti-Muslim hate speech on Canadian YouTube channels in the aftermath of 7 October”, 17 June 2024, Institute for Strategic Dialogue,

https://www.isdglobal.org/digital_dispatches/examining-antisemitic-and-anti-muslim-hate-speech-on-youtube-in-the-aftermath-of-october-7/

“Decoding Antisemitism: An AI-driven study on hate speech and imagery online”, Discourse Report 6, February 2024, Center for Research on Antisemitism,

<https://decodingantisemitism.eu/publications/sixth-discourse-report/>

“Decoding Antisemitism: A Guide to Identifying Antisemitism Online”, Becker et al., 2024,

<http://dx.doi.org/10.1007/978-3-031-49238-9>

“Categories of antisemitic incidents”, Community Security Trust,

<https://cst.org.uk/antisemitism/what-is-antisemitism>

^{xci} “Hateful Conduct”, X, April 2023, <https://help.x.com/en/rules-and-policies/hateful-conduct-policy>

“Abuse and Harassment”, X, March 2024, <https://help.x.com/en/rules-and-policies/abusive-behavior>

"Violent Content Policy", X, February 2025, <https://help.x.com/en/rules-and-policies/violent-content>

^{xcii} Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. "Chain-of-thought prompting elicits reasoning in large language models." *Advances in neural information processing systems* 35 (2022): 24824–24837. https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf

^{xciii} "Evaluating Large Language Models Against Human Annotators in Latent Content Analysis: Sentiment, Political Leaning, Emotional Intensity, and Sarcasm", Bojik et al., 2025, accessed 9 April 2025, <https://arxiv.org/abs/2501.02532>

^{xciv} Researchers gathered views data for all posts and replies made by the accounts between 8 May 2024 and 31 January 2025, amounting to 2,789,653,184 total views or an average of 10,370,458 a day.

^{xcv} "To find an estimate for how often Twitter shows ads, the Center set up three accounts to follow only reinstated accounts and found that after scrolling through 1,039 tweets, an ad appeared at an average rate of once every 6.7 tweets."

"Toxic Twitter: How Twitter Generates Millions in Ad Revenue by Bringing Back Banned Accounts", Center for Countering Digital Hate, 9 February 2023, https://counterhate.com/wp-content/uploads/2023/02/Toxic-Twitter_FINAL.pdf

^{xcvi} "Essential Guide to X Advertising: Strategies and Cost Insights", Brandwatch, 1 November 2024, <https://www.brandwatch.com/blog/x-ads/>

^{xcvii} R. U. Mustafa, N. Ashraf and N. Japkowicz, "Can GPT-4 detect subcategories of hatred?," 2024 IEEE Digital Platforms and Societal Harms (DPSH), Washington, DC, USA, 2024, pp. 1–6, doi: 10.1109/DPSH60098.2024.10775211.

Gandhi, Ankita, Param Ahir, Kinjal Adhvaryu, Pooja Shah, Ritika Lohiya, Erik Cambria, Soujanya Poria, and Amir Hussain. "Hate speech detection: A comprehensive review of recent works." *Expert Systems* 41, no. 8 (2024): e13562. <https://doi.org/10.1111/exsy.13562>

^{xcviii} "Working definition of antisemitism", International Holocaust Remembrance Alliance, accessed 23 January 2025, <https://holocaustremembrance.com/resources/working-definition-antisemitism>



Published September 2025
© Center for Countering Digital Hate Inc