

**CODED,
CONSPIRATORIAL,
AND ACCELERATED:**
THE STATE OF
ANTISEMITISM ONLINE

2025 Year in Review

TABLE OF CONTENTS

EXECUTIVE SUMMARY	3
SECTION 1 – Introduction	5
About the Data	5
SECTION 2 – 2025 at a Glance	7
SECTION 3 – Exploring When the Conversation Spiked	12
SECTION 4 – The Dynamics That Shaped 2025	17
Israel/War Discourse: The Central Backdrop, and the Start of a Shift	17
Conspiracy Theories: A Growing Layer, With Three Themes	19
Holocaust and Nazi Language: Political Shorthand	21
AI and Platform Dynamics: How Technology Reshaped the Conversation in 2025	22
Normalization via Humor, Dog Whistles, and Coded Language	28
Radicalization and Dehumanization: Three Escalating Signals	30
Feelings, Fear, and the Effect on the Jewish Community	31
SECTION 5 – Implications of 2025, and What to Watch Next	34

EXECUTIVE SUMMARY

In 2025, total conversation volume dipped slightly, but participation broadened—and as Israel/Zionism-driven volume declined, antisemitism became more ambient, conspiratorial, coded, and increasingly shaped by Holocaust shorthand and AI-enabled distortion.

BLUE SQUARE ALLIANCE'S FIVE HEADLINE FINDINGS

1 Total volume slightly down; participation up — Mention volume dropped by 6% year-over-year, while the number of users participating in these conversations increased by 10%.

***Why it matters:** Even with fewer total posts, more people engaged—suggesting diffusion into broader audiences and a shift away from a small set of high-volume moments or accounts.*

2 Israel/Zionism-driven volume declined, but when excluded, conversation rose and stayed elevated — Israel and war discussions decreased, contributing to the overall decline, but conversation about Jews and antisemitic themes *without* Israel/Zionism terms increased.

***Why it matters:** The apparent “cooling” is misleading: antisemitism didn’t fade—it shifted away from Israel-centered discussions and became more embedded in cultural and political discourse, with greater focus on Jews.*

3 Holocaust/Nazi references became a primary rhetorical weapon — Holocaust-related discourse was driven less by historical discussion and more by political weaponization; internally, we estimate ~50% involved contemporary political comparisons.

***Why it matters:** Normalizing Nazi and Holocaust analogies trivializes the Holocaust and can erode public understanding of—and sensitivity to—the atrocities committed by the Nazis.*


4 Conspiracy narratives normalized further—especially “power/control” claims about Jews and U.S. politics — “Jewish control” narratives and conspiracies blaming Jews for cultural subversion were more common on social media, while more extreme conspiracies blaming Jews for causing harm in the world grew at faster rates.

***Why it matters:** Conspiracies act as an engine: they turn uncertainty into blame and help antisemitic ideas travel across communities through insinuation rather than overt slurs.*

5 AI became a major part of the story this year — Acting as a content generator, amplifier, and “fact-checker,” AI shaped the conversation as a cited authority, a source of synthetic/remixed content, and a trigger for viral spikes tied to platform incidents.

***Why it matters:** AI compresses the distortion cycle—making it easier to manufacture artifacts, accelerate reach, and rapidly fuse misinformation with conspiratorial antisemitic frames. As technology continues to evolve, it will continue to blur the lines between real and fake information on social media.*

5 NUMBERS THAT DEFINED 2025



SECTION 1 - INTRODUCTION

In 2025, online conversation about antisemitism, Jewish identity, and related narratives tracked by the Blue Square Alliance’s Command Center appeared to cool: total mention volume dipped by 6% compared with 2024. But while the overall volume declined, the tone and the substance of the conversation shifted in meaningful ways. The decrease was driven primarily by a drop in Israel/Hamas war/Zionism-centered discussion, which fell by 13%. However, when removing explicit mentions of Israel- or Zionism-related terms, the trend reverses: volume increases year over year, participation broadens, and the conversation becomes more coded, more conspiratorial, and more directly focused on Jews.

Despite fewer posts overall, more people took part—and the conversation traveled farther. In 2025, we recorded 842 million mentions (down 6%), authored by 17 million users (up 10%), generating an estimated 1,404 billion in reach (up 29%). In other words, the conversation did not shrink. It dispersed.

2025 AT A GLANCE		
Mention Volume	Social Media Users	Estimated Reach
842.4 M	17.0 M	1404.5B
-6%	+10%	+29%

This year’s defining shift was not a simple change in *how much* antisemitism-related discourse exists online, but in *how it showed up*. As Israel- and war-related discussion declined, antisemitic narratives increasingly appeared in forms that are easier to share and harder to moderate: conspiratorial “power/control” frames embedded in everyday political talk; Holocaust and Nazi references used as political shorthand and meme fuel; and coded language that spreads through humor and plausible deniability. At the same time, rapid advances in generative AI influenced both the supply of content and the shape of attention, accelerating how quickly distorted narratives could be created, amplified, and adopted.

What follows substantiates that story with a year-over-year trendline, topic-mix shifts, and deep dives into the dynamics that most shaped the conversation in 2025. After a description of scope and definitions, we will provide a high level overview of the ebb and flow of conversations in 2025, and deep dive into how the conversation changed over time, what topics gained or lost share, and what the data suggests about the forces driving those changes—from conspiratorial narratives and Holocaust distortion to AI, coded language, and the emotional impact on Jewish communities online.

ABOUT THE DATA

This report draws on social listening data from the Blue Square Alliance Command Center, powered by Brandwatch, which tracks online conversation at scale across major social media platforms, including but not limited to X, Instagram, Reddit, YouTube, blogs, and forums. The analysis covers English-language posts globally from January 1 through December 31, 2025, with comparisons made to the same calendar window in 2024.

Throughout the report, “mentions” refers to the full range of captured content types across included platforms—posts as well as replies and comments, and reposting dynamics including quote posts. The dataset is intentionally broad: it includes positive, neutral, and negative conversation relevant to the report’s scope, not only explicitly antisemitic posts. This context matters for interpreting trends in both volume and tone: changes over time can reflect shifts in public attention, platform dynamics, and how people talk about Jews and Jewish-related topics—not just changes in overt hate speech.

This report is a year-in-review of online conversation—how people discussed antisemitism, Jewish identity/culture, and Israel across platforms in 2025—and how that conversation shifted compared with 2024. It is not a census of all antisemitic content online, a measure of offline antisemitism, or a determination of intent behind individual posts. Because the dataset includes news, commentary, and community conversation, volume is best interpreted as a measure of attention, not agreement. A spike in mentions can reflect many dynamics—breaking news, viral platform moments, coordinated amplification, or rapid engagement with a single incident—rather than a single clear cause.

We use a topic-based framing to organize conversation, including discussion of Jewish culture and religion (for example, holidays, community life, and religious references), as well as broader conversation that references Jews more generally (including use of the terms “Jew” and “Jewish”). Because online discourse often blends news, commentary, and community conversation, the dataset also includes news content and headlines that circulate on social platforms. We also include posts where context is ambiguous—including irony, satire, and quote-posting—because these dynamics are central to how narratives spread online, even when the original intent is to criticize or condemn what is being shared.

Finally, metrics such as reach reflect Brandwatch’s estimated reach, which is best interpreted as an indicator of potential exposure and amplification rather than a count of confirmed views or unique individuals reached. Where we refer to “users” or “authors,” the metric reflects the number of distinct accounts participating in the conversation across the platforms included in this analysis.


SECTION 2 - 2025 at a Glance

In 2025, online conversation about antisemitism, Jewish identity and culture, and Israel did not disappear—it shifted shape. Blue Square Alliance’s Command Center (powered by BrandWatch) tracked 842.4M total mentions in 2025. While overall volume was lower than the year prior, the conversation broadened and traveled farther: 17.0M Unique Authors participated (10% increase from 2024), and estimated reach climbed to 1,404.5B (29% increase from 2024). Even in a year with fewer total posts, more people engaged and the conversation’s exposure increased.

Key Metrics (2025 v. 2024)

Metric	2025	YoY change
Total mentions	842.4M	-6%
Unique Authors	17.0M	+10%
Reach (estimated)	1,404.5B	+29%

A WIDENING “NON-ISRAEL” SLICE OF THE CONVERSATION



One of the most important orientation points for 2025 is that the conversation increasingly extended beyond explicit Israel/Zionism framing. Israel-focused discussion remained the majority of what we tracked, but its share of voice declined. In 2025, Israel conversations accounted for 78.2% of conversation (down from 84.3% in

2024). Over the same period, posts about antisemitism and Jewish culture that did not explicitly mention Israel or Zionism—rose from 15.7% of conversation in 2024 to 21.8% in 2025.


This distinction matters for interpreting the year: a decrease in Israel-centered volume can make the year look quieter at first glance, even as conversation about Jews and antisemitic narratives persists—and in many places grows—outside that frame.

Slice of conversation	Mentions	Unique Authors	Reach (estimated)
Israel conversations	675.7M	11.5M	1,268.3B
Non-Israel conversations	183.6M	10.9M	330.6B

WHAT PEOPLE TALKED ABOUT IN 2025: CATEGORY MIX

To provide a high-level map of attention, we also group conversation into thematic categories. This is a next level break down of the analysis and posts can live in multiple categories meaning that a post can be Israel-related and contain Holocaust references, conspiratorial claims, or other themes.

Thematic Category Breakdown of Conversations Online




Several categories grew this past year—especially those that tend to reflect broader cultural hostility, distortion, and myth-making rather than straightforward geopolitical debate including:

Tropes about Jews: +61% Hate Incidents: +45% Holocaust conversations: +31% Conspiracy Theories: +28%

Notably, some categories decreased this past year, with conversations about Zionism down 19% and College Campuses antisemitism down 23% mainly due to the elevated level of discussion regarding college campus antisemitism in 2024 driven by the wave of on-campus encampments in the spring.


Mention Volume per Category in 2025 and Percentage change from 2024



WHEN ATTENTION ROSE: THE YEAR'S RHYTHM

Across 2025, conversation moved in waves—periods of steadiness punctuated by surges of attention. The year began with an early rise, settled into a steadier spring, peaked sharply in mid-year, softened into a quieter summer, and then stepped up again in early fall before rising again toward year-end.

Trend Line of Conversations in 2025 by Weeks



Spikes in volume often aligned with major real-world incidents and high-visibility platform moments, though not every surge can be tied to a single clear trigger. The trendline below shows how attention rose and fell over time, and the incident timeline that follows highlights major events that frequently corresponded with—and helped shape—these waves in online discussion.

Below are the highlights of the major incidents that took place this past year targeting the Jewish community, or promoting online discourse related to antisemitism and the Jewish community.

DATE	INCIDENT
Jan 7, 2025	Meta announces changes to its content moderation approach
Jan 15, 2025	Israel–Hamas ceasefire and hostage-exchange deal announced
Jan 20, 2025	Debate over gestures made during Trump's inauguration
Jan 22, 2025	Antioch High School shooting (Nashville, TN) *
Jan 24, 2025	UFC/MMA fighter Bryce Mitchell praises Adolf Hitler
Feb 7, 2025	Neo-Nazis march in Ohio
Feb 9, 2025	Rapper Ye runs an antisemitic ad during the Super Bowl
Feb 18, 2025	Pro-Palestine protest in Borough Park (Brooklyn) turns violent
Feb 21, 2025	Stabbing at the Berlin Holocaust Memorial
Mar 7, 2025	X rolls out "Ask @Grok"
Apr 13, 2025	Pennsylvania Gov. Josh Shapiro's house is firebombed during Passover
May 3, 2025	Antisemitic sign displayed at Dave Portnoy's Barstool bar in Philadelphia ("Fuck the Jews")
May 5–6, 2025	Pro-Palestine protesters occupy a University of Washington building; university condemns antisemitic remarks tied to the incident
May 8, 2025	Ye (Kanye West) releases new "Heil Hitler" song on X
May 21, 2025	Two Israeli Embassy staffers killed outside the Capital Jewish Museum (Washington, DC)
May 24, 2025	White supremacist group Patriot Front marches through Kansas City, MO
Jun 1, 2025	Boulder, CO fire attack targeting a walk for Israeli hostages killing one and injuring seven including the attacker.
Jun 13, 2025	Israel launches strikes against Iran's nuclear program
Jun 28, 2025	"Death, death to the IDF" chant at Glastonbury Festival (UK)
Jul 8, 2025	Grok posts antisemitic content on X ("MechaHitler" references)
Jul 12, 2025	Conspiratorial and antisemitic Gab AI becomes an interactive tool on X
Aug 5, 2025	Three cars set on fire and "death, death to the IDF" graffiti left outside the home of a Jewish family in Clayton, MO
Aug 27, 2025	Church of the Annunciation shooting (Minneapolis) *
Sep 10, 2025	Assassination of Charlie Kirk (Utah Valley University, Orem, UT) *
Sep 10, 2025	Evergreen High School shooting (Evergreen, CO) *
Sep 22, 2025	Chabad of Charlotte County (Punta Gorda, FL) set on fire on Rosh Hashanah
Sep 22, 2025	Syracuse University: bag of pork tossed into a Jewish fraternity house holding Rosh Hashanah services
Oct 1, 2025	Two people killed in an attack on Heaton Park Hebrew Congregation synagogue in Manchester on Yom Kippur
Oct 14, 2025	Thousands of antisemitic and racist text messages exposed from leaders of the Young Republicans group chat

Oct 24, 2025	OpenAI's Sora is launched and quickly leveraged for antisemitic synthetic video content
Oct 28, 2025	Tucker Carlson hosts Nick Fuentes, amplifying antisemitic ideas
Oct 31, 2025	Viral Halloween content featuring Nazi costumes/imagery circulates
Nov 19, 2025	Pro-Palestine protest held outside a NYC synagogue (Park East Synagogue / Nefesh B'Nefesh event)
Nov 21, 2025	Viral rumor on X alleging Hebrew translation was disabled fuels antisemitic narratives
Nov 22, 2025	X introduces the "About This Account" location feature, prompting controversy over doxxing/targeting risk
Dec 8, 2025	Piers Morgan hosts Nick Fuentes
Dec 13, 2025	Brown University classroom shooting *
Dec 14, 2025	Bondi "Chanukah by the Sea" terror attack killing fifteen people and injuring forty people.
Dec 15, 2025	MIT Professor Nuno Loureiro killed in Brookline, MA *


**Attack not directed at the Jewish community, but it either included antisemitic elements or triggered an increase in antisemitic conspiracy narratives on social media.*

SECTION 3 - Exploring When the Conversation Spiked

These patterns raise the next question: what, specifically, was associated with the most pronounced surges in attention—and how did the conversation differ across the year? This section explores this question, annotating key spike windows and linking them to the events and platform moments that shaped conversation in 2025. The trendline is more than a record of volume—it is a map of how attention moved. In 2025 conversations surged when events broke into public view, when narratives traveled across platforms, and when high-visibility moments pulled people into the discourse.


Compared with 2024, the rhythm of conversation in 2025 looked noticeably different. In 2024, the trendline slopes downward across the year: volume was heavier in the first half and gradually tapered off, interrupted by a distinct spike in October due to the conflict between Israel and Iran, but then continued declining. In 2025, by contrast, the first half is steadier, followed by a major surge in June that does not fully return to the earlier Jan–May baseline. After that mid-year peak, conversation remains elevated through late summer and early fall, climbing again into September and October. In the final months, November dips back toward baseline, but December rises again—signaling renewed momentum as the year closed.

Conversations by Mention Volume (2024 vs 2025)



To unpack what was driving these shifts, we break 2025 into key windows and go one level deeper into the periods when conversation spiked or stepped up.

WINDOW 1: JANUARY–APRIL: CONVERSATIONS STAY STEADY




From January through April, conversation volume remained relatively steady, with increases concentrated in a few short windows rather than sustained surges. These spikes were largely tied to high-salience political moments, Israel– Hamas developments, and campus-related news cycles that pulled antisemitism and Jewish-related narratives into broader public attention.

Key events

- **Jan 20–26:** Conversation increased amid debate over Elon Musk’s gesture at President Trump’s inauguration, early actions by the administration related to antisemitism on college campuses, and the announcement of a ceasefire between Israel and Hamas.
- **Feb 3–9:** Volume rose again following attention to a neo-Nazi group marching in Ohio, renewed discussion of Trump’s Gaza plan, and Super Bowl-related discourse—including Ye/Kanye West’s antisemitic ad and conversation surrounding [Blue Square Alliance’s Super Bowl advertising](#).
- **Mar 17–23:** A further uptick corresponded with the end of the Israel– Hamas ceasefire period and continued campus-focused attention tied to antisemitism crackdowns and related controversies.

WINDOW 2: MAY–AUGUST: LARGE SPIKE IN THE CONVERSATION




From May through early June, conversation volume held relatively steady, fluctuating within a narrow range week to week. That stability broke in mid-June, when volume spiked sharply in response to the Israel–Iran 12-day war—marking the most significant surge of the May–August period and one of the clearest inflection points in the year’s trendline.

Several high-salience incidents and flashpoints fell within this same window and likely contributed to intensified attention and rapid narrative spread online:

- **May 21:** Yaron Lischinsky and Sarah Milgrim were murdered when leaving an event hosted by the American Jewish Committee at the capital Jewish Museum in D.C.
- **Jun 1:** firebombing attack targeting a walk for Israeli hostages in Boulder, CO, killing one and injuring seven including the attacker.
- **Jun 13:** Israel and the U.S. strike Iran’s nuclear program
- **Jun 17:** Threats targeting the Weitzman National Museum of American Jewish History
- **Jun 28:** Glastonbury “death, death to the IDF” chant moment (UK)

After peaking in mid-June, conversation declined quickly and then stabilized through July and August at levels comparable to the pre-spike spring baseline, with a modest lift in late July before tapering again into August.

WINDOW 3: SEPTEMBER – DECEMBER: AN ELEVATED FALL



From September through October, conversation remained consistently elevated, beginning with a strong step-up at the start of September and sustaining through early- to mid-October. In this window, attention was less about a single, sharp peak and more about a prolonged period of higher volume—suggesting multiple overlapping drivers rather than one dominant event.

After mid-October, volume began to decline. That cooling accelerated after the Israel– Hamas ceasefire agreement was signed, and conversation settled into a lower, more stable level through much of November. The final weeks of the year then brought renewed movement: volume began climbing again in December, ending the year on an upward trajectory rather than a sustained low.

This late-year period is also where the split between **Israel-focused** and **non-Israel-focused** conversation became most pronounced. As Israel-centered discussion cooled, a larger share of attention remained concentrated in posts that referenced antisemitism and Jewish identity without explicitly invoking Israel or Zionism—reinforcing the broader pattern that the conversation did not simply fade, but shifted in emphasis.

SEVERAL DYNAMICS APPEAR TO HAVE CONTRIBUTED TO THESE LATE-YEAR WAVES:

- **Conspiratorial surges** tied to breaking political events, including conspiracy narratives attached to the assassination of Charlie Kirk on September 10th.
- **Antisemitism as political discourse**, with escalating rhetoric and ideological framing pulling antisemitic tropes into mainstream debate.
- **A cluster of antisemitic incidents and communal flashpoints** during the High Holidays and their aftermath, including Rosh Hashanah on September 23rd - and Yom Kippur on October 2nd - period incidents and the Manchester synagogue attack.
- **U.S. political and media moments** that drove renewed attention, including the Young Republicans group chat leak on October 14th, Tucker Carlson's interview with Nick Fuentes on October 28th, and other viral controversies.

WHAT THE SPIKE PATTERN SUGGESTS

Taken together, these three windows show a year defined less by a single downward arc and more by stable baseline attention punctuated by distinct waves. In 2025, conversation held relatively steady through the first months of the year, surged dramatically in mid-June, stabilized over the summer, and then stepped up again in early fall—followed by a November reset and a December rebound. The pattern suggests that antisemitism- and Jewish-related discourse in 2025 was shaped by episodic bursts tied to major geopolitical developments, real-world incidents, and high-visibility platform moments, with multiple overlapping triggers sustaining elevated attention in the fall rather than one isolated peak.

This rhythm also clarifies what changed compared with 2024. Where 2024's trendline reads as a clearer negative slope—heavier early-year volume that gradually tapers off, interrupted by a prominent fall spike—2025 shows more stability in its baseline. Instead of steadily declining over time, 2025 oscillated around a relatively consistent level, with attention surging during key moments and then settling back rather than continuing to trend downward. Just as importantly, 2025 ends in a different place: after November's dip, conversation climbs again in December and continues rising into early 2026, suggesting that the dynamics driving attention at the end of the year did not resolve—and that the conversation entered 2026 with renewed momentum rather than winding down.

SECTION 4 – The Dynamics That Shaped 2025

The spike windows in the trendline graph explored in the last section show when attention surged in 2025. The next question is what sustained and shaped the conversation over the course of the year, including in quieter periods—what narratives, frames, and language patterns repeatedly surfaced and influenced how people talked about Jews.


Throughout 2025, the most visible shifts were not simply changes in *volume* but shifts in the content and mechanics of the discourse. As Israel- and war-centered conversation cooled from its highs, a larger share of attention moved into conversation that was often conspiratorial, used Holocaust and Nazi shorthand, and evolving coded language. At the same time, AI, and platform dynamics increasingly shaped both the supply of content and the speed with which narratives traveled, making distortion easier to generate and harder to correct.

The sections that follow examine the key dynamics that defined the year. We begin with the evolution of Israel/Zionism-related discourse because it helps explain why the year appeared to “cool” at first glance. We then trace the layers that increasingly filled the space in the year, conspiratorial narratives, Holocaust/Nazi language as political shorthand, AI and platform shocks, and the normalization mechanics—humor and coded language—that make hate more portable. Finally, we connect these patterns to indicators of the emotional impact visible in how Jews and others responded online after major incidents.

4.1 DEEP DIVE: ISRAEL / WAR DISCOURSE: THE CENTRAL BACKDROP, AND THE START OF A SHIFT

Israel and the Israel– Hamas war remained the anchor of the conversation we tracked in 2025, shaping many of the year’s largest waves of attention. This is consistent with the broader post–October 7, 2023 context, when discourse about Israel surged across platforms and created a heightened baseline of engagement—at times acting as the entry point through which antisemitic narratives traveled. In 2025, major Israel-centered news cycles continued to drive surges, including ceasefire and hostage-exchange developments in January, the Israel–Iran “12-day war” in June, and renewed ceasefire/hostage-deal developments again in October.


Israel-related Conversations (2025) in Millions



At the same time, the data suggest an early shift in how this frame organized the broader landscape. Israel-focused conversation was still massive—675.7M mentions— in 2025, but it declined in volume (-13% YoY) even as participation and amplification rose (11.3M Unique Authors, +4% YoY; 1,073B estimated reach, +20% YoY). In other words, fewer Israel-centered posts still traveled farther and drew in more accounts, underscoring the scale and reach of these conversations even as overall mention volume eased.


The clearest signal of change is visible in share of voice. In 2024, 84.3% of the conversation we tracked explicitly referenced Israel/Palestine or related terms; in 2025, that share declined to 78.2%. Over the same period, the “non-Israel” slice grew from 15.7% in 2024 to 21.8% in 2025. This shift suggests that antisemitic narratives circulated with greater independence from Middle East news cycles, appearing more frequently in domestic politics, cultural commentary, and everyday online language.

Non-Israel Posts Dramatically Increase After September 2025



This matters because it changes what “cooling” looks like. A quieter Israel-centered news cycle can reduce overall volume, but the growth of the non-Israel slice indicates that attention did not simply drop away. In that environment, antisemitism becomes less dependent on specific geopolitical triggers and more sustained by durable narrative engines and rhetorical shortcuts. The remainder of this report examines those dynamics.

Within Israel-centered discourse itself, the year’s major spikes still aligned with widely shared news events. Early 2025 included high-attention moments around ceasefire and hostage-exchange announcements, and hostage-related developments continued to generate surges—including high-profile releases that drew broad visibility and emotional response, as well as conspiracy theories. For example, the May 2025 release of U.S.-Israeli hostage Edan Alexander was widely covered and circulated.





The June Israel–Iran conflict became a defining inflection point in the trendline, concentrating attention and accelerating narrative spread across platforms. During this period, three violent antisemitic attacks also took place—the burning of Governor Shapiro’s house, the murders of Yaron Lischinsky and Sarah Milgrim, and the attack on Jewish protestors calling for the release of Israeli hostages in Boulder, CO— underscoring a period of heightened attacks on Jewish people nationwide.

Later in the year, renewed reporting on ceasefire and hostage-deal negotiations again re-centered the news cycle on Gaza and the hostages, triggering another wave of online engagement.

Taken together, 2025 reveals a “mirage”: Israel and war discourse still explains much of the volume and many of the spikes, but the conversation’s center of gravity began to shift. As Israel’s share of voice declined, a larger proportion of attention remained concentrated in posts that referenced Jews without explicitly invoking Israel or Zionism. In the next parts of this section, we explore what filled that space and how conspiratorial narratives, Holocaust and Nazi shorthand, AI and platform dynamics, and coded language increasingly shaped how antisemitism showed up online in 2025.

4.2 DEEP DIVE: CONSPIRACY THEORIES: A GROWING LAYER, WITH THREE THEMES

Overall, conspiracy narratives became increasingly common in 2025. Posts about antisemitic Conspiracy Theories increased by 28% year over year, reinforcing that conspiratorial framing is not a fringe add-on to antisemitic discourse—it is one of the mechanisms that helps it spread and persist. Conspiracies offer an “explanation layer” that turns uncertainty into blame, repositions breaking news as proof of hidden coordination, and allows antisemitic ideas to travel through insinuation rather than overt slurs.

To better understand what conspiratorial narratives spread most on social media we grouped conspiracy posts into three common broad themes:

GROUP 1:
Power and Control

claims that Jews or “Zionists” control institutions, politics, media, finance, or national priorities.

GROUP 2:
Cultural and Social Subversion


conspiracies that blame Jews for eroding traditional values, promoting moral decline, or driving unwanted social change.

GROUP 3:
Harm and Crisis Creation


claims that Jews cause harm directly, engineer crises, or orchestrate violence (including “false flag” narratives).

In 2025, conspiracies blaming Jews for **cultural and social subversion** were the most common overall. But the fastest growing group were the conspiracies blaming Jews for **harm and crisis creation**—the theme most associated with escalation and dehumanizing narratives.

Mention Volume per Conspiracy Group in 2025 and Percentage change from 2024



THE THEMES IN PRACTICE: What Stood Out in 2025




GROUP 1: Power and control: *visible and politically salient*

Power-and-control conspiracies were among the most politically resonant in 2025, particularly those centered on Jewish control of the U.S. government. This theme frequently

appears through allegations of “puppet” leadership, a shadow government, or insinuations that elected officials serve Jewish interests—often reinforced through “Israel first” framing, a take on the “America first” slogan from the conservative movement. Even when phrased indirectly, the underlying claim is consistent: that Jews exercise hidden, illegitimate power over political outcomes and national decision-making. This framing is a key bridge between geopolitical debate and broader suspicion directed at Jews as a group.

GROUP 2: Cultural and social subversion: *most common overall*

Conspiracies about Jewish influence on culture most often appeared as culture-war framing, with claims that Jews or “Zionists” shape media and entertainment, push ideological agendas, or “brainwash” the public through institutions such as academia and social media. A representative example from this theme is the recurring claim that Jewish



actors “run Hollywood/media” or promote a deliberate social agenda—an allegation that often appears in conversation as a sweeping explanation for cultural change rather than as a direct political argument. These narratives can be especially effective online because they are easily folded into existing ideological communities and presented as “common sense” rather than explicit hatred.

GROUP 3: Harm and crisis creation: *the fastest growing and most destabilizing*

The most concerning shift was the growth of conspiracies accusing Jews of engineering harm or staging events, also called “false flag” claims. These conspiracies claim that violent attacks or other shocking events are not what they appear to be but are secretly staged or orchestrated by a hidden actor to frame someone else or justify a political agenda. In 2025, false-flag narratives spread rapidly after major incidents, including the Bondi Beach attack, in which 15 people died, and the murders of Yaron Lischinsky and Sarah Milgrim in D.C. In these moments, conspiratorial framing often moved quickly from “questions” to accusations, using the language of doubt to deny antisemitic violence while simultaneously reinforcing older myths about manipulation and control. Other conspiracies in this group manifest as blaming Jews for orchestrating attacks domestically and abroad such as the assassination of Charlie Kirk. This theme is particularly impactful because it does double work: it delegitimizes Jewish vulnerability and turns tragedy into evidence for antisemitic worldviews.

Why this matters for the 2025 landscape


Conspiracies are not just one topic within the broader conversation— they function as a connective framework linking many of the dominant narratives about Jews and antisemitism. In 2025, they attached to breaking news, reshaped how incidents were interpreted, and provided a ready-made frame that persisted even when Israel was not the explicit headline.

4.3 DEEP DIVE: HOLOCAUST AND NAZI LANGUAGE: POLITICAL SHORTHAND

Connecticuter @nutmegger740
ICE isn't the damn Gestapo. This is nothing like the Holocaust. I'm sick of people invoking dead Jews like Anne Frank. You know something can be bad without being Nazism? Our country is in a dire moment, but it's not the damn Holocaust.

Holocaust references were increasingly deployed as rhetorical weapons in contemporary debates rather than as language tied to history, education, or remembrance.





Our data shows that about 50% of Holocaust-category conversation in 2025 consisted of contemporary political comparisons—posts using Nazi-era terms as labels or analogies to describe present-day people, institutions, or policies. This pattern is also visible in the growth of key terms that most often show up in these comparisons. A focused query for political-analogy language— “Nazi” OR “Gestapo” OR “Hitler” (excluding common unrelated contexts like “salute,” “tattoo,” and “neo-Nazis”)—rose 29% in 2025. Within that, usage of “Gestapo” increased 228%, “Nazi” increased 56%, and “Hitler” increased 6%.

What makes this shift significant is not only the volume, but the function these terms serve online. Nazi-era labels compress complex political disagreement into a single moral verdict. They take language associated with genocide, totalitarianism, and mass atrocity and redeploy it as shorthand for “enemy,” “evil,” or “illegitimate.” The effectiveness of this strategy has been debated, even within the Jewish community. However, over time, this can blur categories of meaning: when “Nazi” becomes a generalized insult and “Gestapo” becomes a catch-all label for any disliked authority, the words lose specificity—and the Holocaust risks being treated less as a singular historical atrocity and more as a reusable rhetorical tool.

This normalization can have real downstream effects. Extreme comparisons can make the boundaries of acceptable discourse more elastic, pulling Holocaust language into everyday argument and making increasingly escalated rhetoric feel routine. In that environment, Holocaust references are not simply “present” in the conversation—they actively shape its tone, intensify polarization, and create openings for distortion to spread more easily alongside mainstream political speech.

4.4 DEEP DIVE: AI AND PLATFORM DYNAMICS: HOW TECHNOLOGY RESHAPED THE CONVERSATION IN 2025

In 2025, technological advances increasingly shaped how antisemitism spread online. Two trends converged: platforms adjusted how they moderate content (and how their algorithm spreads that content), while generative AI lowered the cost and effort of producing persuasive content at unprecedented speed and volume. The result was a faster distortion cycle: narratives formed more quickly, traveled farther, and were harder to correct once they left their point of origin.


THE PLATFORM BACKDROP: WEAKER GUARDRAILS, BIGGER RIPPLE EFFECTS

In early 2025, Meta announced major changes to how it moderates content, including ending its fact-checking system on Facebook and Instagram. Soon after, leaked training materials surfaced showing that some derogatory, identity-based statements would now be permitted under the updated guidelines—including examples that rely on classic anti-Jewish stereotypes. The result is not just a policy story; it is a structural shift in the information environment. When enforcement becomes less protective or more inconsistent, the window for hateful narratives to gain traction widens—especially during fast-moving news cycles where high-engagement content spreads first and corrections arrive later.

This matters because recent platform history clearly shows what can happen when guardrails shift from enforcement to crowdsourced correction. The BSA Command Center documented that after X adopted Community Notes in late 2022, posts with community notes, which often share conspiratorial ideas, misinformation, or disinformation, tended to get higher engagement and reach. Meta’s decision to discontinue third-party fact-checking in favor of a Community Notes-style system raises concern that a similar dynamic could take hold on its platforms: not necessarily because any single policy change “causes” hate, but because community-driven moderation can inadvertently amplify the most controversial content, allowing tropes and insinuations to travel farther, persist longer, and reach broader audiences before they are challenged.

To illustrate this phenomenon, we examined the activity of the X account Illuminatibot, which posts “Illuminati” conspiracy theories that often include antisemitic themes. In the year after Community Notes were introduced on X, the account’s posting volume changed very little, but the reach of its posts increased 947% year over year. We then isolated Illuminatibot’s September 2023 tweets that received Community Notes to test whether they generated higher engagement and reach than tweets without Notes. Tweets tagged with Community Notes received, on average, 169% more retweets and 199% more reach. In other words, the more controversial tweets—those that prompted a Community Note—were amplified further and reached a substantially larger audience.

Comparing Retweet Volume and Reach for Illuminatibot's Posts with and without Community Notes
September 2023




AI CHANGING THE GAME: THE THREE WAYS IT SHOWED UP IN THE CONVERSATION

Across the year, AI emerged as a driver of antisemitism-related discourse in three distinct roles:

1. **AI as authority (the “truth-finder” effect):** Chatbots are treated as neutral referees—quoted as evidence, screenshotted as “proof,” and used to legitimize claims.
2. **AI as factory (the content multiplier):** Text-to-image and text-to-video tools make it easier to generate persuasive propaganda at scale—high production value, low friction, rapid remixing.
3. **AI as a controversy catalyst (platform shocks):** When AI systems behave unpredictably—hallucinating, contradicting policy, or responding provocatively—the incident itself becomes the story, generating spikes and new conspiratorial framing.

WHAT FOLLOWS ARE THREE CASE STUDIES THAT ILLUSTRATE THESE DYNAMICS IN ACTION:

CASE STUDY 1: GROK ON X —AI AS A CITATION ENGINE



X's rollout of "Ask @Grok" made a platform-native chatbot a routine participant in public conversation—and it scaled quickly. In Grok's first week, 743,000 users asked it over 2 million questions; by the second week, that figure rose to over 5 million questions. And Grok's presence

didn't stop there: between March and July, the AI was tagged in 44 million posts by 5.8 million accounts, and Grok-tagged posts collectively generated 45 trillion impressions. The takeaway is not just that an AI model exists on the platform—it's that its output is now woven into the flow of public conversation at an enormous scale, and that scale has continued to grow beyond the initial launch window.

Grok is most often used directly inside threads as a real-time “context” tool: users tag @Grok beneath viral posts to ask what happened, whether a claim is true, or what a term “really means.” In the early weeks after launch, a notable share of questions centered on Jewish-related topics—spanning the Israel–Hamas war and Zionism, the Holocaust, and conspiracy theories. Many prompts were framed as straightforward information-seeking questions but carried leading assumptions (“How many genocides have the Jews committed?”) or attempts to illicit responses that will help others spread antisemitic rhetoric (“What document shows Nazis killed 6 million Jews?”). This interaction pattern matters because it makes Grok part of the argument itself: its answers are screenshot, quote-posted, and treated as receipts—so the shape of questions users ask (and the premises they embed) can steer the conversation as much as the answers do.

That risk became unmistakable when Grok itself began producing antisemitic content. In July, it generated posts that promoted antisemitic tropes, praised Hitler, and even referred to itself as “MechaHitler.” For a tool being used as a truth-finder, moments like this are uniquely corrosive. They don't just spread a single hateful message—they undermine the boundary between information and propaganda, normalize the idea that AI responses are acceptable inputs to public debate, and provide bad actors with a high-visibility mechanism for injecting antisemitic rhetoric into mainstream threads. In a platform environment that rewards speed and virality, the damage is compounded by how easily AI outputs can be clipped, reposted, and detached from context—long after deletions or corrections occur.



CASE STUDY 2: GAB__AI – AN EXTREMIST CHATBOT ENTERS MAINSTREAM DISTRIBUTION



If Grok illustrates how a mainstream AI can be used as a context tool at scale, Gab__AI illustrates something more direct: an AI system designed to amplify conspiratorial and antisemitic narratives—and then deployed inside a mainstream distribution network. The Gab__AI account launched on X in January 2024, but it became fully interactive on July 12, 2025, when users could tag it to receive AI-generated replies. In just a few weeks after


becoming interactive, it was tagged in 92,000+ posts, responded to 62% of them (more than 56,000 replies), and generated 9+ million impressions. Notably, about one in four of Gab__AI's posts were captured by the Command Center tracking system (antisemitism, Jewish culture, and Israel)—making it a consistent amplifier of content directly relevant to this report's scope.

What makes Gab__AI distinct is not just that it can be “prompted into bias,” but that its default mode leans into extremist framing. Users tagged the bot with questions that consistently surfaced familiar conspiratorial themes—often phrased as “truth-seeking” prompts—such as allegations that Jews or “Zionists” control media and finance, that AIPAC or a “deep state” governs U.S. politics, or that major events are secretly orchestrated by Mossad. Common question clusters included Epstein-related conspiracies (casting him as a Mossad blackmail operation), Great Replacement narratives (framing demographic change as a Jewish plot), and Holocaust denial prompts designed to cast doubt on the death toll or historical record. The bot itself reinforced these frames with confident, polished answers—and it periodically boosted engagement with provocative “conversation starter” prompts inviting users to share controversial opinions, ask about conspiracy theories, or “question” mainstream history.

This matters because it demonstrates a new kind of amplification pathway: AI-generated radicalization. Unlike a niche extremist forum, X provides access to broad, everyday audiences—many of whom may not recognize the source behind the bot or the conspiratorial scaffolding embedded in its answers. When extremist narratives are delivered in the tone of an intelligent assistant, they can feel less like propaganda and more like information—making it easier for conspiracy claims to spread, normalize, and migrate across communities.

This case also shows that platform governance decisions still matter. After Blue Square Alliance [published its reporting](#) on Gab__AI's behavior and reach, X removed the account's ability to interact with users—cutting off its chatbot-style replies. While the content and narratives it amplified continued circulating beyond the bot itself, the removal illustrates that AI-driven amplification is not inevitable: platform interventions can reduce exposure and disrupt the rapid, automated spread of conspiratorial antisemitism.

CASE STUDY 3: SORA AND THE BROADER TEXT-TO-VIDEO LEAP — HIGH REALISM, LOW FRICTION



If Grok demonstrates how AI can function as a “truth tool,” Sora shows the different side of the AI story: AI as a content factory. In October 2025, OpenAI released Sora, a text-to-video tool capable of producing high-definition, hyper-realistic videos in seconds. The adoption curve was immediate: within three days, the invite-only Sora app became the most-downloaded app on Apple’s App Store, surpassing even OpenAI’s own ChatGPT. As invite codes spread widely across social platforms, access expanded quickly—and so did the ecosystem of Sora-generated content migrating onto mainstream feeds.

Within days of launch, antisemitic video content began appearing both inside the Sora environment and, more importantly, across platforms where it could travel without friction. Much of it recycled familiar antisemitic tropes—Jews depicted through greed and money imagery, insinuations of hidden power, and content that reimagined Nazi propaganda or portrayed Hitler in ways that were comedic or stylized enough to feel “shareable.” In some cases, Sora was used to create propaganda-like content tied to war discourse; in others, it was used to generate meme-format clips that trivialized the Holocaust by turning atrocity into a punchline. The shift here is not simply that hate exists in a new medium—it’s that a new medium makes hate easier to package as entertainment, and easier to spread as “just a meme.”

Sora’s realism is what makes it distinct—and what makes it dangerous. Compared with earlier video-generation tools, its outputs look more authentic: camera movement, lighting, and facial expression can feel natural enough that viewers may not immediately suspect manipulation. OpenAI watermarks Sora-generated videos, but users rapidly found ways to remove or obscure watermarks using readily available tools and guides. Once a video is reposted without its original watermark or context, it becomes difficult for everyday viewers to distinguish what is fabricated from what is real—especially in fast-moving news cycles.

The ripple effects were visible in the conversation almost immediately. Following Sora’s release, BSA observed a 216% increase in posts discussing AI-generated videos, including surging mentions of “Sora,” “AI video,” and related terms. That spike is a signal of both public fascination and the rapid normalization of synthetic media as a routine part of the online information environment. And as external research has also shown, this is not only a Sora problem: testing across multiple AI video tools in 2025 found that platforms could still produce hateful or antisemitic content at meaningful rates despite safeguards. The broader takeaway is that as text-to-video generation improves, the “factory” effect accelerates: more content, higher production quality, faster remix cycles, and greater potential for extremists and bad actors to launder old narratives into new formats that feel novel—or even harmless.

This is why the Sora moment matters for 2025’s antisemitism landscape. In earlier eras, antisemitic content was often text-heavy, crude, or easier to identify. With hyper-realistic synthetic video, propaganda can look polished and emotionally compelling, travel farther than a text post, and outpace correction. The line between imagination and information gets blurrier—and that ambiguity is exactly what misinformation entrepreneurs and propagandists exploit.

PLATFORM SHOCKS: WHEN GLITCHES AND FEATURES BECOME ANTISEMITIC STORYLINES

Beyond the case studies, 2025 repeatedly showed how small platform moments can become accelerants—especially when a technical hiccup or new feature creates ambiguity that users then interpret through existing political and conspiratorial frames.



One of the clearest examples came in mid-November, when a rumor spread on X claiming that Hebrew translation had been disabled. The spark was not an official policy change, but a translation failure on a promotional post that mixed Hebrew and English—likely confusing the translation system. An influencer account then circulated a screenshot, from Grok, asserting that Hebrew translation had been removed because “Jews were calling for genocide” without consequences. As the

claim spread, users tagged Grok asking why Hebrew wasn't available, and Grok replied that translation had been disabled to limit inflammatory content. That answer—despite no official explanation from X—was treated as evidence and helped the rumor travel faster. The result was a rapid surge of antisemitic reactions: thousands of posts reframed a mundane technical failure as “proof” of Jewish privilege, special protection, or control over the platform, while other posts claimed translation was removed to hide wrongdoing—folding anti-Israel grievances into overt hostility toward Jews and Hebrew speakers.



A few days later, a different kind of platform moment produced similar dynamics. After X introduced its “About This Account” location feature, showing where an account is based, the rollout generated massive attention as users began scrutinizing the labels on high-profile accounts. In many cases, the feature surfaced discrepancies that fueled debates about authenticity and foreign influence—especially accounts presenting themselves as “American voices” while appearing to operate from abroad, including accounts that also pushed antisemitic conspiracy theories. But some of the most volatile reactions emerged when accounts were marked as based in Israel. Misinterpretations and instability around the labels—particularly viral screenshots suggesting unexpected “Israel” tags—triggered calls for boycotts and became fodder for conspiratorial narratives about

Israeli or Jewish influence. The key pattern is not that the feature itself was antisemitic, but that it created a new, highly visible signal that could be quickly misread, weaponized, and folded into existing storylines about hidden power and coordinated manipulation.

What this changed in 2025

Taken together, these technology dynamics help explain why the conversation in 2025 could feel more volatile—and, in some ways, more durable—even as overall volume dipped. Generative AI lowered the cost of producing persuasive content at scale, from authoritative-sounding “answers” to high-quality synthetic media that is easy to remix and repost. At the same time, tools like Grok were pulled directly into everyday debate as a real-time

context engine, with users treating AI output as receipts—even when the system could be steered by leading prompts or when its own outputs became part of the controversy. Platform policy shifts and product changes also mattered: when guardrails loosened, or when new features and technical ambiguity created confusion, rumors and conspiratorial interpretations had more space to spread quickly. The result was a compressed distortion cycle—claim → amplification → normalization—where antisemitic and conspiratorial narratives could gain traction fast and travel far before corrections, context, or enforcement could catch up.

4.5 DEEP DIVE: NORMALIZATION VIA HUMOR, DOG WHISTLES, AND CODED LANGUAGE


In 2025 antisemitic ideas traveled in forms that were easier to share, easier to deny, and harder to detect. On social media, antisemitism does not typically appear as an explicit declaration of hate, although that happens. Instead, it is often packaged as irony, inside jokes, memes, or coded references that rely on an audience already knowing the meaning. Coded language lowers the social and platform “cost” of engaging with antisemitic content: it makes it feel more casual, more ambiguous, and therefore more portable across communities.



In 2025 we saw the emergence of new dog whistles that gained traction quickly and spread widely. Three terms stood out as especially prominent this year:



“PROMISED TO THEM 3,000 YEARS AGO” emerged as a dog whistle in 2025, generating 74,000+ posts with an estimated reach of 72M+. It began gaining momentum over the summer and peaked in October. On its face, the phrase references the biblical idea of a “promised land,” but online it’s typically used sarcastically to mock and delegitimize Jewish historical and religious ties to Israel—framing Jewish connection as irrational entitlement to an “ancient myth.” It also frequently expands beyond Israel to suggest Jews would use the same “excuse” to claim other countries too—tapping into older antisemitic tropes that portray Jews as greedy, manipulative, and power-hungry.

“SPIRITUALLY ISRAELI” also emerged in 2025, appearing in 81,000+ posts reaching 62M+, rising through the summer and fall and spiking in December. The term circulated as an internet insult for something “soulless,” “corporate,” or “culturally empty”—but as a dog whistle it does more than criticize a trend. It functions by treating “Israeli” (and, by association, Jews) as a negative vibe-label that can be pasted onto anything someone dislikes, reinforcing a mental association between Jewishness/Israel and “cringe,” emptiness, or moral deficiency. The phrase also echoes older prejudices that painted Jews as spiritually lacking or inferior, repackaging that notion into a meme-friendly form that’s easy to spread and hard to challenge.



“EARLY LIFE” became a notable coded marker in 2025 as well. While it’s harder to quantify (because it’s a common phrase), it circulated as a shorthand prompt—often “early life” or “check their early life”—that directs people to the “Early life” section of a Wikipedia page to surface whether someone is Jewish or has Jewish family background. The antisemitic logic is implicit but clear: it suggests that Jewish identity is the hidden explanation for behavior the user dislikes, inviting audiences to “connect the dots” without saying anything explicit. In practice, it acts as a mechanism for stigmatizing Jewish identity and turning biography into insinuation—replacing overt slurs with a coded “gotcha” that can be passed off as neutral curiosity.

A key feature of this normalization layer is churn: coded language evolves quickly as moderation systems and public awareness adapt. The lifecycle can be fast—terms rise, spread, get recognized, and then fade or mutate.



That churn is part of why the 2025 story cannot be understood by tracking a static list of slurs or keywords alone. The codes themselves shift, and the conversation shifts with them.


Older dog whistles continued to circulate and overlap with newer ones in 2025, creating a kind of “shared vocabulary” that users can mix and match depending on platform context and audience. Two examples illustrate this continuity:

THE “TJD” ACRONYM remained active in 2025, generating 12,000+ posts with an estimated reach of 26M. On its surface, it often appears as an innocuous wish—“have a totally joyful day”—showing up on posts such as those that reference Jewish topics and feel out of place. That is because the acronym also has hateful contexts, TJD also stands for “total Jew death” for people in the know. The phrase functions as a coded signal: it lets users embed a violent message in a format that looks harmless, preserving plausible deniability while cueing like-minded audiences to read the deeper meaning. The acronym also developed to attack other minority groups with phrases such as “totally nice day” (TND), “totally fabulous day” (TFD), “totally kind day” (TKD). That combination—everyday language carrying hidden hostility—makes it a powerful normalization tool, allowing hate to circulate more widely and evade casual detection.

“THE NOTICING” (and related variants) continued circulating as a durable insinuation framework in 2025, even when it is harder to track cleanly through simple keyword counts. The phrase signals a moment of “awakening” where someone claims to have suddenly recognized a hidden pattern—typically implying

that Jews are behind social, political, or cultural outcomes. In practice, “the noticing” operates like a gateway: it invites audiences to adopt antisemitic explanations without saying them explicitly, relying on winks and innuendo (“once you notice, you can’t unsee it”). This framing mirrors long-standing conspiracy traditions that portray Jews as a secret coordinating force.

A parallel evolution in 2025 was the increasing role of **EMOJIS AS CODED SYMBOLS**. Emojis are uniquely well-suited for normalization: they are compact, visually shareable, and often treated as “just a joke.” But in context, emojis can function as shorthand for dehumanization, conspiracy, or ideological signaling—especially when they appear repeatedly in the same kinds of threads or alongside familiar coded phrases. Relatedly, the use of animal emojis revived older antisemitic patterns by visually encoding dehumanizing comparisons in a way that can evade straightforward moderation.



Because coded language depends on context, it is important to be explicit about what this section does and does not claim. Many of these phrases can appear in non-hateful contexts; we do not treat individual instances as definitive evidence of antisemitism. Instead, we track how terms function at scale—how they cluster in antisemitism-adjacent conversations, how they spike in predictable patterns, and how they operate as signals that make hostile ideas easier to spread without explicit wording.


Ultimately, this normalization layer helps explain why 2025 looked like a shift rather than a decline. As Israel-centered volume cooled, coded references and “shareable” insinuations continued circulating in parallel—often detached from Israel or Zionism terms and embedded in cultural and political conversation. That is one reason the non-Israel slice of discourse grew: not because antisemitism disappeared from mainstream spaces, but because it increasingly traveled through humor, ambiguity, and code.

4.6 RADICALIZATION AND DEHUMANIZATION: THREE ESCALATING SIGNALS

Jumping 180 degrees - since October 7, 2023, antisemitic rhetoric online has not only expanded—it has *hardened*. In a separate analysis, we tracked three markers of radicalization in English-language public posts: content that portrays Jews as inherently evil, content that dehumanizes Jews, and posts where users self-identify as antisemitic or Nazi-aligned. Post-October 7 period (Oct. 7, 2023–July 28, 2025 for this analysis) each category rose sharply



This signals both a quantitative surge and a qualitative shift toward more explicit, socially unrestrained expressions of antisemitism.




DEPICTING JEWS AS EVIL captures rhetoric that frames Jews as inherently malevolent, morally corrupt, or the source of evil—language that treats “Jewishness” as an innate threat rather than a religion, identity, or community. In practice, this often appears through claims that Jews are “pure evil” or “evil by nature,” or that they are aligned with dark forces—rhetoric that turns disagreement or grievance into moral absolutism.

DEHUMANIZING JEWS captures language that pushes Jews outside the bounds of humanity—describing Jews as subhuman, parasitic, animalistic, or like a disease. A related pattern includes “disgust” framing that invites revulsion without always using overt slurs (for example, language that describes Jews as “garbage,” “scum,” or a “walking disease”). This distinction matters because disgust-based rhetoric can normalize hostility: it makes hatred feel intuitive, and it can circulate more easily across mainstream spaces than explicitly violent language.

SELF-IDENTIFIED ANTISEMITISM

captures posts where users openly declare or take pride in antisemitic beliefs or Nazi alignment (for example, “antisemitic and proud” or statements that frame Nazism as socially acceptable). While this



category appears at lower overall volume than the other two, it is uniquely revealing: publicly adopting the label of “antisemitic” suggests an erosion of stigma and a growing sense that extreme identity statements can function as social signaling, not social cost.

Crucially, these were not one-time spikes that faded as headlines changed. The pattern persisted into 2025. Compared with 2024, we continued to see growth in all three categories: posts depicting Jews as evil increased by 63% in 2025, dehumanizing posts increased by 48%, and self-identified antisemitism increased by 49%. In other words, the post-October 7 surge created a higher baseline—and 2025 built on it.

The post-October 7 surge helps explain why these trends accelerated—but 2025 is notable because the escalation continued even as overall conversation volume dipped. The same year that we saw broader discourse shift toward more coded and conspiratorial forms of antisemitism, we also saw a parallel intensification in language that is *less coded*: rhetoric that depicts Jews as evil, strips Jews of humanity, and increasingly treats antisemitism itself as a badge of belonging.


4.7 FEELINGS, FEAR, AND THE EFFECT ON THE JEWISH COMMUNITY

All the dynamics covered in this report—shifts in Israel-centered discourse, the normalization of conspiracy narratives, Holocaust terminology as political shorthand, AI-driven acceleration, and coded language that makes hate more shareable—have a real-world effect: they shape how safe people feel, how communities interpret risk, and how Jewish identity is expressed publicly online. Across multiple flashpoints, social media conversation showed a recurring pattern: major attacks against Jewish communities triggered sharp, measurable increases in fear-laden posts, alongside disgust, grief, and anger. These shifts did not stay contained to the immediate news cycle. They also showed up in more intimate cultural spaces online—holiday greetings, communal rituals, and everyday conversations about what it means to be visibly Jewish in public.

AFTER ATTACKS, FEAR BECOMES THE DOMINANT EMOTIONAL SIGNAL

The Bondi “Chanukah by the Sea” attack illustrates the intensity of these emotional surges. In the days immediately following the shooting, the emotional profile of posts about antisemitism, Jewish life, and Israel shifted sharply—fear and disgust together accounted for more than half of emotionally coded mentions. Compared with a similar period before the attack, posts expressing fear rose roughly 160% and disgust rose about 125%. Even joy nearly doubled, often appearing in solidarity messages and hopeful holiday wishes—an important reminder that fear and resilience often coexist in the same moments.

Breakdown of emotions in posts about antisemitism, Jewish life, and Israel following the Bondi Beach Shooting



But the most revealing part of these spikes is what the fear was *about*. In the Bondi aftermath, a major cluster of posts centered on fears of further attacks and copycat violence, with people questioning whether Jewish gatherings—schools, synagogues, public menorah lightings, small businesses—could remain open and visible. Another cluster directed anxiety toward institutions, criticizing governments, police, and platforms for ignoring warning signs or failing to protect Jewish communities. In other words, fear was not only an emotional reaction; it was an assessment of risk, trust, and public vulnerability.

A similar pattern appeared earlier in the year after two attacks in close succession: the murders of Yaron Lischinsky and Sarah Milgrim in Washington, D.C. and the Boulder fire attack targeting a walk for Israeli hostages. The tone around Boulder was overwhelmingly driven by fear, differing from the D.C. response where disgust and sadness were more prominent before fear. What stood out most was the diminishing presence of surprise—an unsettling signal that, for many users, antisemitic violence was starting to feel disturbingly normal.


HOLIDAYS AS A CULTURAL BAROMETER OF SAFETY AND VISIBILITY

In 2025, these emotions weren't confined to moments of crisis. They surfaced in the language of Jewish life itself—especially around holidays when social media tends to amplify communal identity and public expressions of belonging.

Regarding Passover, analysis over the past three years shows the holiday's online conversation shifting from unity and interfaith optimism toward urgency and reflection. In 2023, Passover overlapped with Ramadan and Easter, and 25% of Passover posts also referenced celebrating those holidays—reflecting a surge of interfaith solidarity messaging. By 2024 and into 2025, Passover conversation carried heavier themes, shaped by rising antisemitism and ongoing war. In 2025 specifically, posts reflected the tension between celebration and sorrow: many referenced hostages still held in Gaza and the newly found tradition of leaving them an empty seat at the Seder table, and conversation was further intensified by the fire attack on Pennsylvania Governor Josh Shapiro's residence on the first night of Passover.

Rosh Hashanah offered an even more distilled measure of mood: the evolution of greetings themselves. Over the past three years we analyzed the share of posts wishing for a sweet year versus a safe year during the holiday. In 2023, about 75% of Rosh Hashanah greetings wished for a "sweet" or "joyful" year. In 2024, wishes for a "safe year" rose to 47%. In 2025, that shift became more pronounced, with more than half of posts emphasizing a "safe" or "peaceful" year—suggesting that safety and protection have become central themes in Jewish public well-wishes.

Percentage of posts mentioning Rosh Hashanah with "sweet" or "safe" as the sentiment



SECTION 5 - Implications of 2025 —And What to Watch Next

The 2025 data points to a clear implication: antisemitism did not cool off in 2025. Instead, it took on different forms. In 2025, antisemitism on social media became less concentrated on Israel/war framing and more dispersed across cultural, political, and conspiratorial discourse—often in forms that are easier to share and harder to detect. Israel-centered conversation still dominated much of what we tracked, and major spikes were still tied to Middle East news cycles. But as explicit Israel/Zionism language declined, a larger share of the conversation shifted toward Jews as the subject, carried by conspiracies, coded language, and Holocaust-era shorthand. In that sense, 2025 reflects not a reduction in antisemitism, but a change in its operating mode: more ambient, more portable, and more resilient.

Several mechanisms help explain why. Conspiracies increasingly function as a default “explanation layer” that could attach to almost any incident, especially in moments of uncertainty. Holocaust and Nazi references were increasingly used as political shorthand, flattening moral categories and eroding the weight of those terms. Platform dynamics and AI accelerated the pace and scale of distortion: chatbots were pulled into argument as context engines and authoritative voices, while synthetic media tools lowered the barrier to producing persuasive propaganda and misinformation. Meanwhile, coded language—dog whistles, euphemisms, and emoji signaling—made antisemitic ideas easier to circulate with plausible deniability, allowing them to travel farther into mainstream spaces.

The human signal behind these dynamics was also visible. After attacks, fear rose sharply, and even holiday greetings increasingly focused on safety and protection—an indication that the ambient nature of the conversation is not abstract. It shapes how people assess risk, how they interpret public visibility, and how Jewish life is expressed online.

What to Watch In 2026

- **AI-ENABLED DISTORTION AND SYNTHETIC MEDIA.** As generative tools, especially text-to-video, improve, synthetic content will become more persuasive, easier to remix, and harder to trace back to origin. The speed of the distortion cycle will continue to compress.
- **FASTER CONSPIRACY ATTACHMENT AFTER BREAKING NEWS.** False-flag narratives and blame frameworks are likely to keep surging quickly after violence or breaking news, turning real-world events into rapid vehicles for antisemitic framing.
- **ESCALATION SIGNALS: DEHUMANIZATION AND SELF-IDENTIFICATION.** Continued growth in dehumanization, “evil” framing, and self-identified antisemitism suggests a lowered threshold for explicit extremity—an early warning indicator for normalization and potential downstream harm.
- **ANTISEMITISM IN POLITICAL CYCLES.** Jewish issues, including antisemitism, played a key role in political conversations in 2025 and will likely continue to play a role in 2026.

CLOSING IMPLICATION

The story of 2025 is that antisemitism online did not recede—it repositioned. It became less dependent on a single geopolitical frame and more embedded in the broader attention economy: conspiratorial, coded, and increasingly shaped by platform and AI dynamics. That shift raises the stakes for 2026: the question is not only how much conversation exists, but how quickly it mutates, how widely it spreads, and what kind of behavior and fear it produces in the real world.