

X'S ENFORCEMENT GAPS ON DEHUMANIZATION AND THREATS OF VIOLENCE AGAINST MUSLIMS



1629 K St. NW, Suite 300
Washington, D.C. 20006

The Center for the Study of Organized Hate (CSOH) is a nonprofit, nonpartisan think tank based in Washington, D.C. CSOH is strongly driven by its mission to advance research and inform policies that combat hate, violence, extremism, radicalism, and online harms.

Our research, strategic partnerships, and community engagement programs are guided by the vision of a more inclusive and resilient society against all forms of hate and extremism.



The Bridge Initiative is a multi-year research project on Islamophobia housed at Georgetown University.

Author

Meriem Mahdhi, Researcher, CSOH

©2026 Center for the Study of Organized Hate (CSOH).
All rights reserved

No portion of this report may be reproduced in any form without permission from CSOH, except as permitted under applicable U.S. copyright laws.

For inquiries, please email: press@csohate.org

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY	4
2. FINDINGS	6
2.1. IDENTICAL THREATS INCONSISTENT ENFORCEMENT IN IPHONE	8
2.2. BLUE BADGES AND REPEAT VIOLATORS	9
3. X'S POLICIES AND THE ENFORCEMENT GAP	11
4. CONCLUSION	13
5. RECOMMENDATIONS	15
6. END NOTES	17

1. EXECUTIVE SUMMARY

This study was conducted against a backdrop of a sharp surge in bigotry and incitement targeting Muslim Americans across social media platforms, alongside a documented rise in hate crimes and bias incidents targeting the community. In 2025, the Council on American-Islamic Relations (CAIR) recorded 8,683 civil rights complaints nationwide,¹ including 555 direct hate incidents and 33 instances of vandalism, damage, harassment, intimidation, or zoning disputes at Islamic institutions such as Mosques, community centers, and prayer spaces, the highest number of single-year complaints CAIR has recorded since its first civil rights report covering the year 1996.

The threat posed by this broader climate of anti-Muslim hostility became starkly evident on May 18, 2026, when two gunmen aged 17 and 18 opened fire at the Islamic Center of San Diego and killed three people. The two shooters left behind a manifesto² that named the 2019 Christchurch mosque shooting as the inspiration for the attack and mentioned conspiracy theories claiming that Muslims are trying to impose “Sharia” in Texas and the UK, as well as claiming that the Somali community in Minnesota is responsible for scamming the government out of millions of dollars.

CAIR documented 17 incidents targeting Mosques³ and other Islamic institutions in the three months following the Islamic Center of San Diego attack, a 183% increase over the same period the previous year, when it recorded six. On July 5, less than two months after San Diego, a man threw an incendiary device at the Northeast Philadelphia Islamic Center⁴, causing a fire inside the structure. Following the arrest of the man, investigators tied him to a Facebook account⁵ which he had used to discuss “using fire to damage Mosques”.

The spike in threats of violence against Muslims and their places of worship was also evident in the digital space, particularly on X (formerly Twitter), where CSOH has previously documented a higher baseline of anti-Muslim content, including bigotry, dehumanization, incitement, and conspiracy theories, amid shifts in the platform’s content moderation infrastructure and enforcement approach.

For this report, we identified and reported 464 posts published between January 10 and September 9, 2026, by US-based accounts. Those contained explicit threats of violence against Muslims and their places of worship, as well as explicit dehumanizing language that has historically preceded that level of violence.

This study therefore examines two distinct layers of content moderation on X: proactive moderation, in which the platform detects and acts on policy-violating content without first receiving a user report, and reactive moderation, in which the platform reviews and acts on content after it has been formally reported through its user-reporting system.

All 464 posts were located using keyword queries through X's Advanced Search. They fell into two broad categories: extreme dehumanization and calls for violence. Of the 464, 194 posts (42%) used extreme dehumanizing language, describing Muslims as vermin, cockroaches, rats, parasites, a cancer, or an infestation. The remaining 270 posts (58%) called for violence against Muslims or Islamic institutions: 160 posts (34% of the dataset) called for Mosques to be burned and 110 (24%) called for Muslims to be killed.

The manner in which this content was identified is itself significant. None of the posts in the dataset employed coded language, deliberate obfuscation, or other techniques intended to evade detection. They relied instead on the most direct formulations available: the phrases "burn the Mosques" and "kill Muslims," alongside explicit dehumanization. These are precisely the terms any enforcement system designed to detect violent speech and dehumanizing content would be expected to prioritize.

The fact that researchers could readily locate these posts using the platform's own search functionality raises significant questions about the effectiveness of proactive moderation. The content remained live, publicly searchable, and accessible long enough for researchers to systematically identify and compile it into a dataset.


The second phase tested X's reactive moderation. All posts were formally reported through the platform's user-reporting system for either threats of violence (calls to kill Muslims and burn Mosques) or hate, abuse, and harassment (for posts containing dehumanizing language). While the platform does not have an official response time for user-submitted reports, we checked the outcome of all posts flagged to X 7 days after the report was made. The platform's reporting guidelines confirm that all reports are acknowledged within 24 hours of receipt and state that "resolution times vary, but it may take thirty days".⁶

2. FINDINGS

This brief draws on a dataset of 464 posts published on X between January 10 and September 9, 2026, which contained dehumanization or threats of violence against Muslims and their places of worship. Researchers reported each post to X through the platform's standard reporting tool, selecting the Hateful Conduct or Violent content policy as appropriate.


The dataset falls into two broad categories: posts dehumanizing Muslims as “vermin,” “parasites,” or referencing an “infestation,” and posts containing explicit threats of violence, including calls to “kill Muslims” or “burn Mosques.” Across all posts, 200 (43%) remain fully accessible, 89 (19%) are visibility-limited but still accessible, 171 (37%) were removed, and 4 (1%) were deleted by their authors. Overall, 289 posts, 62% of the dataset, remain accessible on X in some form.

Figure 1 (overall 464 posts): Enforcement Outcomes for Reported Anti-Muslim Content on X




X took some form of action on 88% of posts containing explicit threats of violence. Of the 270 posts in this category, X removed 165 (61%) and limited the visibility of 72 (27%), while 32 (12%) remained fully accessible and one was deleted by its author.

**Figure 2 (comparison by content type):
X's Enforcement of Reported Violent and Dehumanizing Anti-Muslim Content**



"Violent threats" combines calls to burn Mosques (160) and calls to kill Muslims (110).
 "Deleted by author" denotes posts removed by the author, not by platform enforcement.

Broken down further, 93 of 160 posts calling for Mosques to be burned (58%) and 72 of 110 calling for Muslims to be killed (65.5%) were removed. The response to dehumanizing content differed markedly: only six of 194 posts (3%) were removed, while 168 (87%) remain fully accessible. This category alone accounts for 84% of all posts in the dataset that remain online without any visibility restrictions.




Examples of Violent and Dehumanizing Posts That Remain Online

2.1 Identical Threats, Inconsistent Enforcement

Posts containing the same narratives received inconsistent removal rates once reported to the platform, with materially different responses even when the wording was nearly identical.

For example, not all violent threats, such as calls to “burn the Mosques,” were equally actioned: this exact phrase, or close variants of it, appears in 160 posts in the dataset, and those posts split several ways: 93 were removed, 49 were visibility-limited, 17 remain fully live and untouched, and the author deleted one.

One post reading simply “BURN THE MOSQUES” has been live since January 10 with no restriction at all, while another reading “Burn the Mosques down” was visibility-limited, and a third, “Now burn the Mosques,” was removed outright. All three say functionally the same thing, and each received a different response.



Similarly, a post reading “Kill all Muslims; impale them” has remained fully accessible since April 14, while another stating “Kill Muslims before they kill you” was visibility-limited, and a third stating “Hell no. Kill the Muslims and make dogs happy” was removed outright.

This Post violated the X Rules. [Learn more](#)



Jason Horn
@JasonHorn636



Visibility limited: this Post may violate X's rules against Violent Speech. [Learn more](#)

Kill Muslims before they kill you.

6:04 PM · Apr 22, 2026 from Imperial, MO · 3 Views



SayWhatNow
@saywhatnowet



Kill all muslims; impale them.


9:59 AM · Apr 14, 2026 · 6 Views

2.2 Blue Badges and Repeat Violators

The analysis identified 45 accounts that posted more than one piece of violative content. Although these accounts represented just 12% of the 386 accounts in the dataset, they accounted for 123 posts, or 26.5% of all violative content analyzed.

Posts from Premium accounts were removed at a rate of 26%, compared to 41% for unverified accounts. On the contrary, while non-subscribers are more likely to have their posts removed by X, blue-badge accounts are more likely to have the visibility of their posts limited by the platform, with data showing users with a subscription having their posts' visibility limited at a 25% rate compared to the 17% rate for users with a non-premium account.

Figure 3. Enforcement by Account Type: Premium vs. Standard



Repeat posting was also closely tied to having a paid X Premium subscription. Among accounts with a blue badge, 20.5% posted more than once, compared to just 9% of standard accounts, meaning Premium subscribers were more than twice as likely to come back and post again. Of everything posted by Premium accounts in the dataset, 43% came from these repeat posters, meaning more than two in five Premium-account posts came from a small group of accounts posting repeatedly, rather than from one-off posts.

3. X'S POLICIES AND THE ENFORCEMENT GAP

The findings show both a proactive and reactive enforcement gap in moderating dehumanization and threats of violence against Muslims, when measured against X's own policies⁷ that prohibit the type of content covered in this study. The platform's Hateful Conduct policy strictly prohibits targeting people or groups based on protected characteristics such as religion; this includes dehumanizing posts describing Muslims as "vermin", "cockroaches", and any other language related to "infestation". Its Violent Content policy⁸ defines violent speech as "content that threatens, incites, glorifies, or expresses desire for violence or harm", but also includes violent threats expressing a desire "to damage civilian homes and shelters, or infrastructure that is essential to daily, civic, religious, or business activities", such as violent threats to Muslims and their places of worship.

Yet, X's capacity and stated philosophy for policing this kind of content have both shifted since Elon Musk acquired ownership of the platform (then Twitter) in October 2022. Musk disbanded Twitter's Trust and Safety Council,⁹ an independent advisory board on hate speech, child exploitation, and self-harm, that had operated since 2016, without replacing it with another oversight body. According to 2024 figures, the company's own global trust and safety team shrank by 30%, from 4,062 to 2,849¹⁰ employees and contractors. Its in-house content moderation team¹¹ was cut in half from 107 to 51, while engineers¹² focused on trust and safety globally fell from 279 to 55 (an 80% reduction). X admitted to Australian regulators¹³ that responses to reported hateful posts slowed by 20%, and to hateful direct messages by 75%.

Despite identified gaps in content policy enforcement, X's transparency¹⁴ report describes a moderation system that combines machine learning and human review. Its strongest claims about proactive detection are specifically geared for child exploitation and terrorist content, categories it says are caught by automated systems before most users ever see or report them; it makes no equivalent claim for hate speech or violent speech targeting a religious group.

X's own data further confirms this: user reports of accounts and posts surged roughly 1,830%¹⁵ between late 2021 and the first half of 2024, while account suspensions grew about 300%, and suspensions specifically for hateful conduct fell from 104,565 to 2,361 over the same period, a decline of roughly 98%. Its EU-specific disclosures sharpen that same picture: over the six months covered by its transparency report¹⁶ for the period 1 October 2024 to 31 March 2025, automated systems proactively applied more than 284,000 reach-restriction labels for Hateful Conduct across EU countries, while automated removal under that same policy totaled a single digit across all 27 listed countries.

X's transparency practices have also drawn direct regulatory scrutiny. In December 2025, the European Commission issued its first DSA non-compliance decision against X,¹⁷ a

120 million Euro fine over separate failures in its blue checkmark account verification system, advertising repository, and researcher data access, while a broader Commission investigation¹⁸ into X's handling of illegal content and disinformation remains ongoing.

The evidence in this brief describes a systematic enforcement gap and failure in platform accountability. X prohibits every category of content identified and reported, yet the content remained visible for several months despite the platform's automated checks, and 62% remains accessible after manual reporting. Our data suggests that this gap is not evenly distributed. It is widest for posts containing dehumanizing language, 87% of which remain fully live while just 3% are removed. It also varies by account type: posts from blue checkmark accounts paying X for enhanced visibility survive at nearly twice the rate they are removed (47.5% stay up against 26% removed), while posts from standard accounts are closer to an even split (41.5% stay up against 41% removed).

X states that it enforces its Hateful Conduct and Violent Content policies based on the severity and harm of content, which could require swift removal. Nonetheless, in April 2023, the company announced its new "Freedom of Speech, not Reach" philosophy,¹⁹ meaning that rather than removing content, it would restrict its visibility through reduced recommendation, exclusion from search, or a warning label, but often leaving the post itself up. Consequently, X's Hateful Conduct policy reserves visibility limiting for what it calls "moderate, isolated usage" of harassment and intimidation, rather than broader incitement.

X's Violent Content policy ties the response to how severe it judges the content to be, and describes this directly as taking "proportionate action based on the severity and likelihood of harm." Violent Speech that doesn't meet that bar is handled more flexibly; X states it "may remove or reduce the visibility" of this content instead, reflecting an aim to keep conversations open while still limiting exposure to violent themes for users who don't wish to see them.

4. CONCLUSION

CSOH's April 2026 report *Manufacturing the Muslim Threat*²⁰ documented a coordinated wave of anti-Muslim rhetoric, tracking 1,111 posts by 46 elected officials since February 2025. Nearly half of the posts advanced claims that Muslims were seeking to impose Islamic law in the United States, often alongside Great Replacement narratives portraying cities with large Muslim populations as being under civilizational conquest.

Our research has also linked measurable spikes in anti-Muslim content on X to specific trigger events. Zohran Mamdani's rise to become the first Muslim mayor of New York City²¹ drew 35,522 Islamophobic posts in the run-up to the general election alone, reaching an estimated 1.5 billion impressions, with 72% of that volume framing him outright as a terrorist or extremist threat because of his faith, while Abdul El-Sayed's campaign²² in Michigan's 2026 Senate race saw a 20-fold Islamophobic surge after his August 4 primary win.

The intensification of the US-Iran conflict produced a comparable surge with explicit dehumanizing language and direct calls for deportation, internment, and violence against Mosques.

The dehumanizing language documented in this report is a serious concern. Researchers at the Dangerous Speech Project²³ classify this specific register as Dangerous Speech, a category built on the observation that violence against a group is rarely committed or condoned spontaneously, as a population must first be taught to see the targeted group as "vermin", "pests", or an "existential threat". A substantial share of the content identified went further still, into explicit threats of violence, including direct calls to burn Mosques and to kill Muslims; precisely the category of content that the platforms' own violent speech policies are designed to address.

The relationship between this kind of speech and subsequent offline violence is well documented, not speculative. FBI officials have confirmed that the two teenage perpetrators of the San Diego mosque shooting met online and were radicalized there²⁴, absorbing white-supremacist ideology before carrying out the attack. Experts have said the manifesto the perpetrators left behind, which invoked the 2019 Christchurch shooting, first circulated on a Discord server before being shared to another online forum.²⁵ That case is not an isolated one: as recently as September 19, 2026, a Spotsylvania County man was arrested and charged²⁶ after posting a comment from his Instagram profile, threatening to bomb the ADAMS Center in Sterling, Virginia, according to the Loudoun County Sheriff's Office.

With the 2026 midterm elections approaching amid elevated levels of anti-Muslim bigotry and incitement online and offline, the enforcement gaps documented in this report warrant particular concern. CSOH's previous research has found measurable increases in anti-Muslim content on social media platforms. A nationwide election cycle

in which immigration, religion, and questions of national belonging feature prominently could create conditions for further increases in anti-Muslim incitement online and resulting in offline violence.

X, Meta, Alphabet, and other social media platforms must enforce their policies on hateful and violent content, disclose how their AI moderation systems handle threats and dehumanizing content against Muslims, and act proactively against violent threats while the community faces heightened risk of hate-fueled violence.

5. RECOMMENDATIONS

- **Enforce existing policies consistently:** X's Hateful Conduct and Violent Content policies already prohibit every category of content documented in this brief, including dehumanizing language and threats against people and places of worship based on religious affiliation. X should enforce these policies as written. That means acting on dehumanizing content, 87% of which remained fully accessible after being reported, and applying consistent outcomes to content that is functionally identical, so that a call to "burn the Mosques" is not removed in one instance and left live in another.
- **Disclose how proactive enforcement works:** X should disclose what proactive detection systems it uses for hate speech, dehumanization, and incitement targeting religious groups; which policy categories those systems cover; and the extent to which enforcement in the United States relies on proactive detection versus user reports. Transparency reports and reporting tools should disaggregate Hateful Conduct by violation type, including dehumanization, rather than treating it as a single category that can obscure enforcement gaps.
- **Strengthen accountability for verified users:** Verified accounts have played a significant role in amplifying direct threats of violence and dehumanizing narratives against Muslims, and because verification brings algorithmic prioritization and access to monetization, these accounts carry greater reach and greater responsibility. X should enforce its policies consistently against verified repeat violators, up to suspension; remove verification privileges, including amplification and monetization, from accounts spreading hate-based content.
- **Transparency in user reporting:** Users reporting violative content must be able to track the real-time status of the reported posts. They must be notified of the platform's decision with a reasonable explanation. The platform must also disclose the kind of action taken, including visibility restrictions. They must disclose whether the decision was made solely by an automated system. The absence of this information undermines the complainant's right to appeal the decision in internal and external independent grievance redressal fora.
- **Adopt surge protocols for elections and other trigger events:** Anti-Muslim content on X rises in sharp, identifiable bursts tied to specific events. X has stated that it conducts risk assessments for national elections and applies its Civic Integrity policy²⁷ during election periods, but that framework primarily addresses voter participation, not hate speech or threats against religious communities. Ahead of the 2026 midterms, X should extend these assessments to cover targeted hate and incitement, and adopt surge protocols around elections and other foreseeable trigger events, including expanded proactive detection, faster review

of user reports, and public disclosure of when protocols are activated and what enforcement resulted.

- **Support counterspeech alongside enforcement:** Counterspeech is defined as speech that challenges harmful narratives. Platforms have long sought to promote counterspeech as a better solution to hateful content as opposed to bans and restrictions on violative content. But counterspeech cannot substitute for enforcement, particularly against direct threats of violence, and its effectiveness depends on whether platforms actively support it. X should develop a coherent counterspeech strategy that works alongside, not in place of, consistent enforcement of its hate speech policies. This may include labeling hateful and racist content, increasing the visibility of counterspeech that challenges it, and offering premium features free or at a discount to vetted counterspeakers and organizations working in this space.

6. END NOTES

1. Council on American-Islamic Relations, "CAIR's 2026 Civil Rights Report Shows the 'Right to Be Different' Narrowed in the Past Year," press release, March 10, 2026, https://www.cair.com/press_releases/cairs-2026-civil-rights-report-shows-the-right-to-be-different-narrowed-in-the-past-year/.
2. Southern Poverty Law Center, "San Diego mosque shooters left document revealing racist, neoNazi beliefs," document submitted for the record, U.S. House Committee on the Judiciary, 119th Cong., June 9, 2026, <https://www.congress.gov/119/meeting/house/119365/documents/HHRG-119-JU00-20260609-SD017-U17.pdf>.
3. Council on American-Islamic Relations, "Mosque Incidents May-June-July 2026," Google Sheets, accessed October 6, 2026, <https://docs.google.com/spreadsheets/d/1G1iqFTRDVJM7OvjB6-JzUtc0L-hUiZIRjhG0yvYohL0/edit?gid=0#gid=0>.
4. U.S. Department of Justice, Office of Public Affairs, "Philadelphia Man Indicted on Civil Rights Charge in Connection with July Arson at Northeast Philadelphia Islamic Center," press release, August 6, 2026, <https://www.justice.gov/opa/pr/philadelphia-man-indicted-civil-rights-charge-connection-july-arson-northeast-philadelphia>.
5. David Charns, "Suspect in Philly Mosque Fire Showed Anti-Islamic Views Before Attack, Feds Say," NBC10 Philadelphia, July 16, 2026, <https://www.nbcphiladelphia.com/news/local/suspect-in-philadelphia-mosque-fire-expressed-anti-islamic-views-before-attack-investigators-say/4434260/>.
6. "Report Violations," X Help Center, accessed October 6, 2026, <https://help.x.com/en/rules-and-policies/x-report-violation>.
7. "Hateful Conduct," X Help Center, last modified April 2023, <https://help.x.com/en/rules-and-policies/hateful-conduct-policy>.
8. "Violent Content Policy," X Help Center, last modified February 2025, <https://help.x.com/en/rules-and-policies/violent-content>.
9. Associated Press, "Musk's Twitter Has Dissolved Its Trust and Safety Council," NPR, December 12, 2022, <https://www.npr.org/2022/12/12/1142399312/twitter-trust-and-safety-council-elon-musk>.
10. Rod McGuirk, "X Corp. has slashed 30% of trust and safety staff, an Australian online safety watchdog says," Associated Press, January 2024, <https://apnews.com/article/x-corp-musk-australia-staff-safety-bc4772369cab1fe8dd975132fd8d61ed>.
11. Thomas Brewster, "Musk's X Fired 80% of Engineers Working on Trust and Safety, Australian Government Says," Forbes, January 10, 2024, <https://www.forbes.com/sites/thomasbrewster/2024/01/10/elon-musk-fired-80-per-cent-of-twitter-x-engineers-working-on-trust-and-safety/>.
12. Rod McGuirk, "X Corp. has slashed 30% of trust and safety staff, an Australian online safety watchdog says," (Associated Press; see note 10).
13. Evelyn Manfield, "Online Safety Regulator Lashes X, Formerly Twitter, over Failure to Police Hate," ABC News (Australia), January 11, 2024, <https://www.abc.net.au/news/2024-01-11/online-safety-x-twitter-failure-online-hate/103307246>.
14. X Corp., Global Transparency Report H1 2024 (X Corp., 2024), <https://transparency.x.com/content/dam/transparency-twitter/2024/x-global-transparency-report-h1.pdf>.
15. Anisha Sircar, "X's Latest Content Findings Reveal Troubling Trends in AI Moderation," Forbes, October 18, 2024, <https://www.forbes.com/sites/anishasircar/2024/10/18/xs-latest-content-findings-reveal-troubling-trends-in-ai-moderation/>.

16. X Internet Unlimited Company, "DSA Transparency Report – April 2025," X Transparency Center, April 2025, <https://transparency.x.com/dsa-transparency-report-2025-april.html>.
17. European Commission, "Commission Fines X €120 Million Under the Digital Services Act," press release IP/25/2934, December 5, 2025, https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_25_2934/IP_25_2934_EN.pdf.
18. "European Commission Investigation into X's Compliance with DSA Requirements to Address Illegal Content and Disinformation," Digital Policy Alert, accessed October 6, 2026, <https://digitalpolicyalert.org/change/7366-european-commission-investigation-into-x-formerly-twitter-compliance-with-dsa-requirements-to-address-illegal-content-and-disinformation>.
19. "Freedom of Speech, Not Reach: An Update on Our Enforcement Philosophy," X Blog, April 17, 2023, https://blog.x.com/en_us/topics/product/2023/freedom-of-speech-not-reach-an-update-on-our-enforcement-philosophy.
20. Center for the Study of Organized Hate, Manufacturing the Muslim Threat (CSOH, 2026), <https://www.csohate.org/wp-content/uploads/2026/04/Manufacturing-the-Muslim-Threat.pdf>.
21. Center for the Study of Organized Hate, "Islamophobia and the New York City Mayoral Election," November 3, 2025, <https://www.csohate.org/2025/11/03/islamophobia-and-new-york-city-mayoral-election/>.
22. Center for the Study of Organized Hate, "Anti-Muslim Hate Targeting Abdul El-Sayed Surged 20-Fold After Primary Win," August 12, 2026, <https://www.csohate.org/2026/08/12/abdul-el-sayed-islamophobia/>.
23. Dangerous Speech Project, homepage, accessed October 6, 2026, <https://www.dangerousspeech.org/>.
24. Associated Press, "San Diego Mosque Shooters Met Online and Left Writings Expressing Hate, FBI Says," ABC News, May 19, 2026, <https://abcnews.com/US/wireStory/san-diego-mosque-shooters-met-online-left-writings-133097787>.
25. Andrew Dyer, "Experts: Mosque Shooters Followed Familiar Path of Far-Right Radicalization," KPBS, May 20, 2026, <https://www.kpbs.org/news/public-safety/2026/05/20/experts-mosque-shooters-followed-familiar-path-of-far-right-radicalization>.
26. Loudoun County Sheriff's Office, "Spotsylvania Man Charged in Connection with Threat Against ADAMS Center," news release, September 21, 2026, <https://sheriff.loudoun.gov/m/newsflash/home/detail/10881>.
27. "X's Civic Integrity Policy," X Help Center, last modified November 2025, <https://help.x.com/en/rules-and-policies/election-integrity-policy>.

