



The 2024 Election Integrity Initiative

Auditing Political Exposure Bias: Algorithmic Amplification on Twitter/X Approaching the 2024 U.S. Presidential Election

Jinyi Ye, Luca Luceri, Emilio Ferrara

University of Southern California
HUMANS Lab – Working Paper No. 2024.9

Auditing Political Exposure Bias: Algorithmic Amplification on Twitter/X Approaching the 2024 U.S. Presidential Election

Jinyi Ye, Luca Luceri, Emilio Ferrara
University of Southern California

ABSTRACT

Approximately 50% of tweets in X’s user timelines are personalized recommendations from accounts they do not follow. This raises a critical question: what political content are users exposed to beyond their established networks, and how might this influence democratic discourse online? Due to the black-box nature and constant evolution of social media algorithms, much remains unknown about this aspect of users’ content exposure, particularly as it pertains to potential biases in algorithmic curation. Prior research has shown that certain political groups and media sources are amplified within users’ *in-network* tweets. However, the extent to which this amplification affects out-of-network recommendations remains unclear. As the 2024 U.S. Election approaches, addressing this question is essential for understanding the influence of algorithms on online political content consumption and its potential impact on users’ perspectives.

In this paper, we conduct a three-week audit of X’s algorithmic content recommendations using a set of 120 sock-puppet monitoring accounts that capture tweets in their personalized “For You” timelines. Our objective is to quantify *out-of-network* content exposure for right- and left-leaning user profiles and to assess any potential biases in political exposure. Our findings indicate that X’s algorithm skews exposure toward a few high-popularity accounts across all users, with right-leaning users experiencing the highest level of exposure inequality. Both left- and right-leaning users encounter amplified exposure to accounts aligned with their own political views and reduced exposure to opposing viewpoints. Additionally, we observe a right-leaning bias in exposure for new accounts within their default timelines. We hope this empirical study promotes discussion on the transparency and accountability of social media algorithms, contributing to critical issues on safeguarding election integrity and fostering a more informed digital public sphere.

INTRODUCTION

During the 2024 U.S. Presidential Election, social media platforms like X (formerly Twitter) serve as key venues for political information and discourse. Meanwhile, the information users encounter on the X platform is increasingly curated by algorithmic recommendation systems that personalize content in their “For You” timelines. As of this writing, X’s “For You” timeline typically consists of 50% *in-network* tweets (i.e., from accounts a given user follows) and 50% *out-of-network* tweets (i.e., from accounts that user does not directly follow)—up from 37% in 2023 [6]. How does X’s algorithm select relevant tweets from outside a user’s network? In 2023, Twitter partially open-sourced its recommendation algorithm, revealing that out-of-network recommendations are sourced through engagement and follow graphs, ranked by a neural network, and refined with heuristics and filters [22]. Despite this algorithm’s growing influence, little has been done to examine the specific composition and nature of these out-of-network tweets.

The impact of algorithmic content curation on political discourse in social media has been a major focus of research and public debate: Previous studies consistently show that X’s algorithm

amplifies political biases and prioritizes high-engagement content, including emotionally charged, toxic, and low-credibility information [2, 3, 6, 7, 9, 14]. Researchers have used methods including randomized experiments, sock-puppet audits, crowdsourced audits, and observational data to study X’s algorithmic effects. Some have found that Twitter’s algorithms tend to amplify content from right-leaning media sources and politicians more than their left-leaning counterparts [14]. Other studies report increased exposure to ideologically aligned friends [2, 6] and low-credibility content [9], with right-leaning users experiencing higher exposure to such content [7]. For out-of-network tweets, qualitative analyses suggest the algorithm leans toward promoting centrist content for partisan users [7] and displays a more diverse political mix overall [6].

Despite these insights, current research on X’s algorithmic auditing faces a critical challenge in analyzing out-of-network content: Many studies assess amplification by comparing personalized timelines with reverse-chronological timelines as baselines, where tweets appear in the order they were posted without algorithmic effects [3, 4, 6, 14]. However, out-of-network tweets lack a reverse-chronological baseline, as users do not follow the authors of those tweets, making it challenging to quantitatively measure exposure bias. To address this limitation, we utilize a “sock-puppet audit,” a study design that deploys artificial user accounts with controlled features to systematically capture and analyze platform recommendations. This approach is particularly well-suited to studying out-of-network exposure patterns because it allows us to observe algorithmic behavior without the interference of real user behaviors or connections [2, 3, 7].

In this study, we deploy 120 sock-puppet accounts distributed across four groups—left-leaning, right-leaning, balanced, and neutral—enabling us to collect a robust dataset that currently encompasses over 5 million tweets. This large-scale audit offers a unique perspective on algorithmic content exposure, as it allows for comprehensive comparisons across various political profiles. Further details are provided in the Experimental Setup section.

The contributions of this work can be summarized as follows:

- **We quantify algorithmic exposure to out-of-network content** for users with varying political alignments during the 2024 U.S. election period through a sock-puppet audit of X’s personalized timelines.
- **We propose a methodology for evaluating out-of-network (political) exposure biases** by creating a baseline using politically balanced accounts.

We find that:

- The X platform **skews exposure toward a few high-popularity accounts** for all users, with right-leaning users experiencing the most inequality.
- Both left- and right-leaning users encounter amplified exposure to accounts aligned with their own political stance and **reduced exposure to opposing viewpoints**.
- Additionally, neutral accounts who do not follow anybody (akin to a newly-registered user account on the platform) show a **default right-leaning bias** in content exposure.

EXPERIMENTAL SETUP

To analyze algorithmic biases in X’s “For You” timeline, we developed a timeline scraper to systematically collect tweets recommended to different types of user profiles. We created 120 accounts

divided into four groups: 30 neutral accounts (default setting, following no one), 30 left-leaning accounts, 30 right-leaning accounts, and 30 balanced accounts.

To categorize the political alignment of accounts to follow, we used the AllSides Media Bias Chart,¹ which rates news sources on a spectrum from left to right based on their political bias. Each left-leaning and right-leaning account follows 10 media outlets, including seven outlets with a moderate (center-left or center-right) bias and three with a stronger (left or right) bias, as defined by the AllSides' chart. This selection ensures that these accounts represent a realistic mix of moderately and strongly aligned sources, enhancing the accuracy of our analysis of political exposure. Additionally, left-leaning accounts follow key Democratic figures and entities (Kamala Harris, Tim Walz, House Democrats, and Senate Democrats), while right-leaning accounts follow their Republican counterparts (Donald Trump, JD Vance, House Republicans, and Senate Republicans). Balanced accounts, designed to reflect a centrist perspective, follow five center-left and center-right media outlets and both presidential candidates from each major party. All media follows are randomly selected from the respective groups in the media bias chart, ensuring consistency with each group's intended alignment.

The timelines for each account are scraped four times daily, yielding approximately 500–700 tweets per session, or about 2,000–3,000 tweets per account per day, within the limits that X's terms of service impose on new, non-premium accounts. The choice of four daily scraping sessions was made to capture the variability in recommendations throughout the day, as the content recommended by X's algorithm can shift based on temporal factors like recent events or trending topics. This frequency provides a more comprehensive picture of the algorithmic exposure that users might experience. Data collection began on October 2, 2024, and continued through October 28, 2024. Regarding the proportion of out-of-network tweets, neutral accounts, which follow no other accounts, exclusively receive out-of-network content, while other account types have approximately 56–59% of their content coming from out-of-network sources. Further details on data collection are provided in the Methods section.

RESULTS

Out-of-Network Exposure Inequality Among Users with Different Political Leanings

One significant aspect of algorithmic biases on social media is popularity bias [19]. Algorithms often tend to amplify content from certain users over others, creating inequalities in exposure [4]. For instance, Twitter's ranking algorithm employs a ~48M parameter neural network, which uses thousands of features to score each tweet based on engagement probabilities, prioritizing content with higher likelihoods of interaction in users' feeds [22].

Table 1 provides an overview of the statistics for the collected tweet dataset across different account types. It shows the average proportion of out-of-network tweets that each account type encounters, with neutral accounts seeing exclusively out-of-network content, while the other accounts have around 56-59% of their timelines composed of out-of-network tweets. Additionally, it details the average proportions of retweets, quoted tweets, and promoted tweets observed by each account type.

¹AllSides Media Bias Chart <https://www.allsides.com/media-bias/media-bias-chart>

Table 1. Statistics of the collected tweet dataset (mean values with standard deviations)

Statistic	Neutral	Left	Right	Balanced
Out-of-network tweet	100%	56.51% (6.52)	57.41% (7.25)	59.00% (5.56)
Retweet	0.04% (0.11)	2.81% (1.02)	2.33% (1.30)	2.57% (1.86)
Quoted tweet	0.94% (1.56)	7.93% (2.12)	12.75% (3.33)	12.40% (2.59)
Promoted tweet	0.95% (1.44)	6.42% (1.09)	5.67% (1.14)	7.54% (0.62)

Previous research has shown that popularity biases can lead to a skew in the visibility of tweets when comparing personalized feeds with reverse-chronological ones, and that users are disproportionately exposed to friends’ tweets [3, 4]. Yet, it remains unclear whether exposure inequalities extend beyond friends to include a broader set of recommended users. Specifically, we ask:

Are personalized recommendations in \mathbb{X} distributed evenly among users, or is exposure dominated by a few accounts? Furthermore, do these inequalities differ among users with different political leanings?

To address these research questions, we use the Gini coefficient, a standard measure of inequality that quantifies disparities in exposure by calculating how concentrated exposure is across a set of users. A Gini coefficient close to 1 indicates high inequality (i.e., a few users receive most of the exposure), while a value near 0 suggests a more equal distribution. To measure a user’s exposure within a timeline, we introduce a metric called “weighted occurrence per 1,000 tweets,” defined as the number of times a user’s tweets appear per 1,000 tweets in the timeline, weighted by each tweet’s visibility according to its rank using an exponential decay function—this correction is introduced to give more weight to tweets that appear earlier in one’s timeline, as those tweets are also the more likely to be seen by a user and are known to generate more engagements [15]. Detailed descriptions of the Gini coefficient calculation and the exposure metric are provided in the Methods section. For each sock-puppet monitoring account, we compute its Gini coefficient with respect to all recommended users in that account’s timelines.

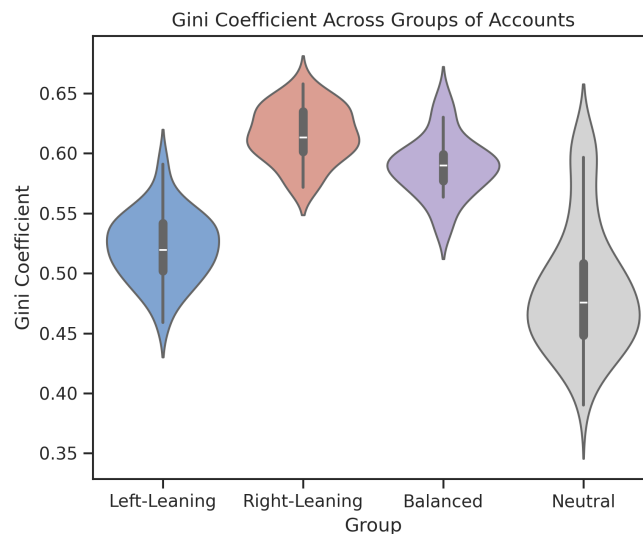


Fig. 1. Distribution of Gini coefficient across different groups of accounts. Significant disparities are found in all pairwise comparisons (Mann-Whitney U test: $p < 0.001$), with right-leaning users experiencing the highest out-of-network exposure inequality.

Figure 1 presents the distribution of Gini coefficients across different account groups: Left-Leaning, Right-Leaning, Balanced, and Neutral. The Mann-Whitney U test reveals that the differences in Gini coefficients between all pairs of groups are significant at the 0.001 level, underscoring meaningful disparities in exposure inequality across these groups. As illustrated in Figure 1, the average Gini coefficient across all groups exceeds 0.45, which suggests a moderate to high level of inequality in exposure on the \mathbb{X} platform. This indicates that algorithmic exposure is somewhat concentrated among certain users rather than evenly distributed.

Notably, right-leaning users experience the highest exposure inequality, followed by balanced and left-leaning users. This suggests that the algorithm’s out-of-network tweet recommendations for right-leaning users are more centralized, reflecting a stronger popularity bias, where a few users dominate exposure. In contrast, neutral users—who do not follow anyone—receive the most diverse recommendations, potentially due to *algorithmic cold start*, i.e., the absence of information about user preferences that typically informs recommendations [24].

Since neutral accounts are critical for detecting bias, we took particular care in their setup to ensure neutrality. Neutral accounts follow no other accounts and therefore receive exclusively out-of-network recommendations. This configuration limits any bias that could arise from follow choices, aiming to capture a baseline view of how the algorithm behaves when no user-preferences are specified. However, it is worth noting that certain factors, such as \mathbb{X} ’s default settings or trending topics, could still introduce slight biases into these recommendations.

Our findings are significant when compared to previous studies that report Gini coefficients of approximately 0.6–0.7 for inequality in exposure to friends’ tweets [3]. This suggests that even beyond the friend network, exposure inequality remains at a similar level, indicating that the platform’s algorithm amplifies certain accounts both within and outside of users’ direct networks.

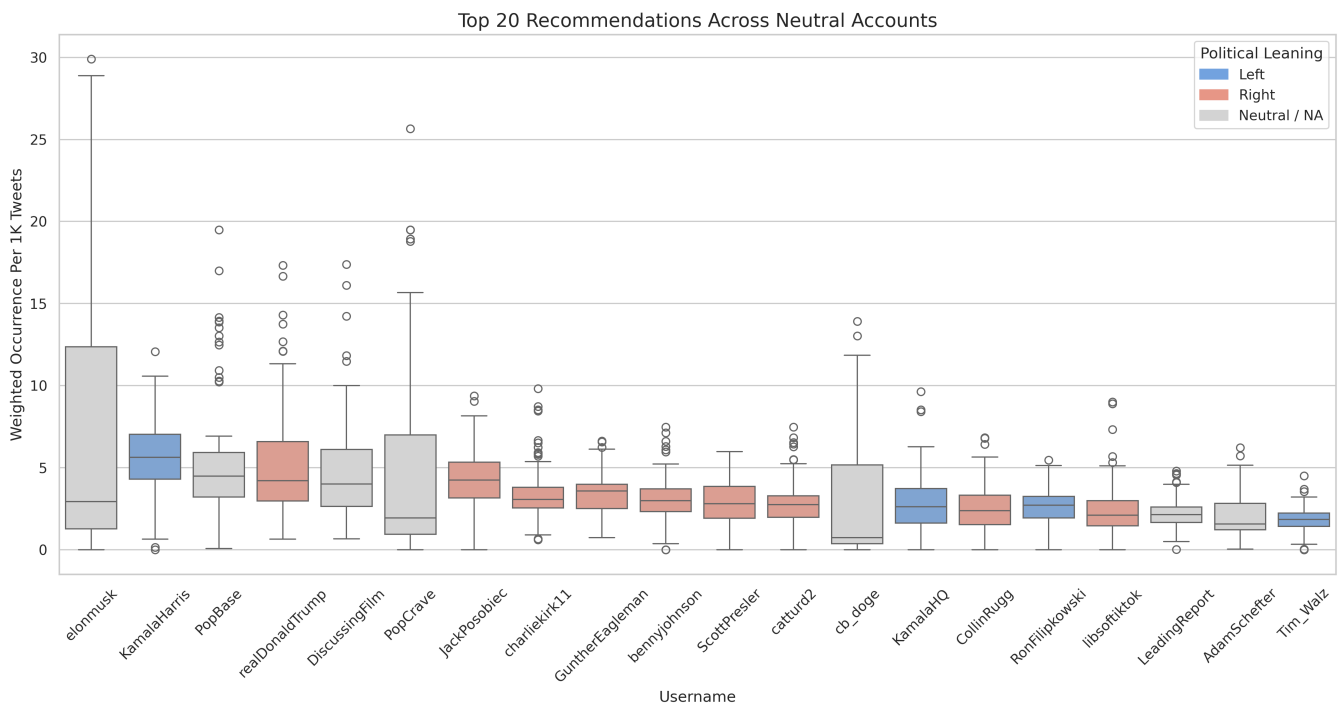


Fig. 2. Top 20 recommended users in neutral accounts, ranked by their average weighted occurrence per 1,000 tweets. Each box in the boxplot shows the distribution of exposure across all neutral accounts, with red and blue colors indicating right- and left-leaning users, respectively. The figure suggests that right-leaning users are more frequently recommended than left-leaning users in the algorithm’s out-of-network recommendations for neutral accounts.

Now that we understand that out-of-network exposures are skewed toward certain users, an important question arises: *Who are these users?* Here, we are particularly interested in neutral accounts, which provide an unbiased look at the algorithm’s default behavior. Figure 2 displays the top 20 recommended users in neutral accounts, ranked by their weighted occurrence per 1,000 tweets. Each box in the boxplot represents the distribution of this exposure metric across all neutral accounts. Boxes are colored red or blue to indicate whether the user is right- or left-leaning, based on publicly available data. A qualitative inspection reveals that right-leaning users appear more frequently among the top recommendations than left-leaning users. To quantify this difference, we use the “weighted occurrences per 1,000 tweets” metric: among the top 20 recommended users, right-leaning users make up 30.16% of exposure, compared to 29.34% for left-leaning users. This disparity persists as we expand the pool, with right-leaning users making up 35.26% of exposure in the top 50 (versus 22.34% for left-leaning users) and 31.39% in the top 100 (versus 20.83% for left-leaning users). Further investigations to characterize these users will be conducted in future studies.

In the Appendix, interested readers can find the top 20 recommendations for the left-leaning, right-leaning, and balanced account groups, shedding light on the most amplified accounts across different user groups. A detailed table describing these users’ public information is also provided.

Differential (De-)Amplification of Political Content Among Partisan Accounts

Selective exposure is a psychological concept that refers to the tendency of individuals to prefer information that aligns with their pre-existing beliefs, attitudes, or preferences, while avoiding information that contradicts them [12]. Algorithms on social media platforms can amplify this effect by recommending content similar to what users already prefer or agree with, reinforcing selective exposure through personalization [16]. To assess the amplification of certain accounts in partisan users’ timelines, we introduce the “mean amplification ratio” metric, inspired by Huszár et al. [14]. This metric compares the exposure of a user in left-leaning or right-leaning timelines relative to a baseline observed in politically balanced users. A positive *mean amplification ratio* indicates amplification, while a negative ratio indicates de-amplification. For a detailed description of this metric, please refer to the Methods section.

Figure 3 shows the amplification ratio of the top 50 recommended users in left-leaning and right-leaning accounts, compared to a baseline of balanced accounts. Colored bars indicate a significant difference in exposure metrics (weighted occurrence per 1,000 tweets) between groups at the 0.05 significance level (using the Mann-Whitney U test), while gray bars indicate no significant difference. Detailed information about these users is provided in the Appendix. A qualitative inspection reveals that left-leaning sock-puppet accounts tend to see left-leaning users amplified, and right-leaning users de-amplified, with the opposite pattern observed for right-leaning accounts.

For instance, in left-leaning accounts, the top three amplified users are *POTUS* (the official account of the U.S. President, currently Joe Biden of the Democratic Party), *Ron Filipkowski* (a former federal prosecutor known for his criticisms of conservative figures), and *Mueller, She Wrote* (a political commentary and investigative journalism account with a liberal stance). In contrast, the most de-amplified accounts are *catturd2* (a right-wing influencer known for political satire), *Jack Posobiec* (a right-wing media personality and political activist), and *Elon Musk* (CEO of Twitter/X, who has recently shared conservative viewpoints). This suggests that left-leaning timelines prioritize left-aligned figures while downplaying right-leaning accounts.

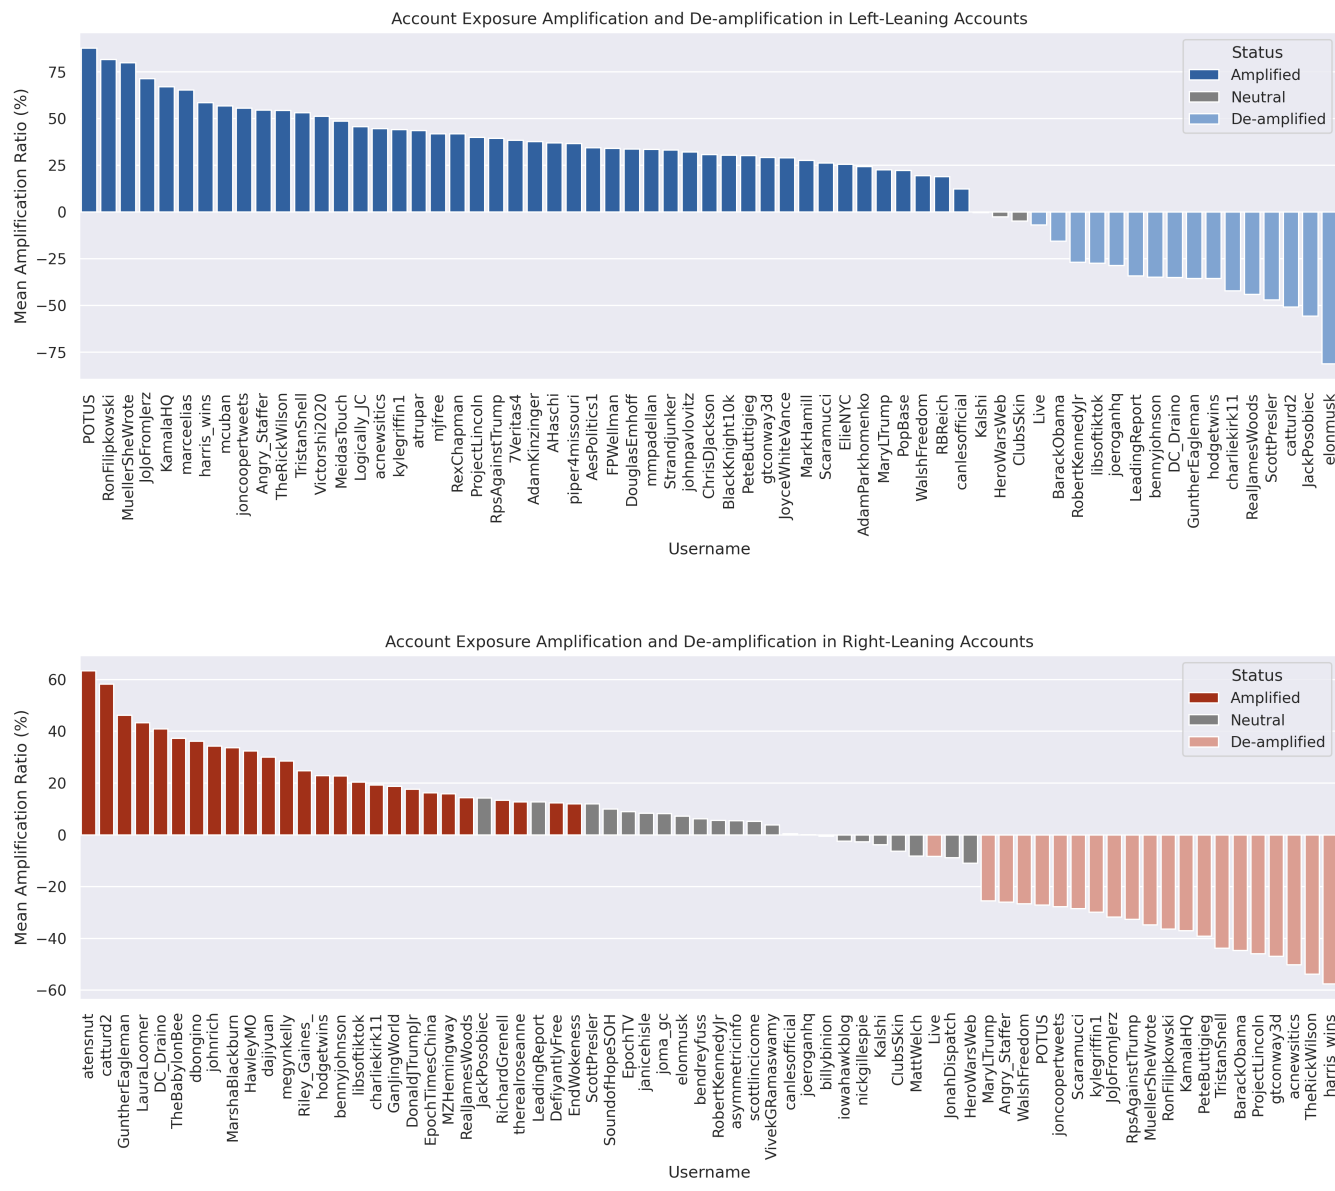


Fig. 3. Amplification ratio of the top 50 recommended users in left-leaning (top) and right-leaning (bottom) accounts, compared to the baseline of balanced accounts. Colored bars indicate a significant difference in exposure metrics (weighted occurrence per 1,000 tweets) between the groups at the 0.05 significance level (using the Mann-Whitney U test), while gray bars indicate no significant difference.

On the contrary, for right-leaning accounts, the top three accounts with the highest amplification in right-leaning timelines are *atensnut* (a conservative commentator), *catturd2* (a right-wing influencer known for political satire), and *Gunther Eagleman* (a conservative content creator). Conversely, the most de-amplified accounts are *Rick Wilson* (a co-founder of The Lincoln Project), *harris_wins* (an online community supporting presidential candidate Kamala Harris), and *acnewsitics* (a liberal-leaning news commentator). This pattern highlights the algorithm’s tendency to amplify conservative figures more heavily in right-leaning timelines while reducing exposure to left-leaning accounts.

Notably, we observe that amplified users in left-leaning group experience a significantly higher magnitude of amplification compared to those in right-leaning group ($M_{left} = 41.86\%$, $M_{right} = 27.99\%$, Mann-Whitney U $p < 0.001$). However, there are no significant differences in the extent of de-amplification between the two groups.

DISCUSSION

In this study, we present preliminary findings from an ongoing audit of algorithmic recommendations on X’s “For You” timelines. Using 120 sock-puppet accounts with left-leaning, right-leaning, balanced, and neutral political orientations, we observe that X skews exposure toward a select few high-popularity accounts for all users, with right-leaning users experiencing the highest level of inequality. Both left- and right-leaning users see amplified exposure to accounts aligned with their political stance, while exposure to opposing viewpoints is reduced. Additionally, qualitative analysis of neutral accounts with no follow activity reveals a default right-leaning bias in the platform’s recommendations.

Our analysis of exposure inequality aligns with previous studies on algorithmic bias, which have reported similar amplification patterns within users’ in-network content [4]. However, our findings diverge from earlier research suggesting that personalized recommendations tend to be more centrist in political stance [6, 7]. This discrepancy perhaps highlights a shift in X’s algorithmic behavior, which might have moved away from promoting moderate content to reinforcing users’ existing preferences more explicitly, especially in out-of-network recommendations. The results also add to the growing body of literature indicating that right-leaning accounts are often more prominently featured in algorithmic recommendations, a trend seen here in the default bias toward right-leaning content for new or neutral accounts.

These findings have important implications for the 2024 U.S. Presidential Election. Social media platforms like X play a significant role in shaping political discourse, and the algorithmic amplification of politically aligned content may influence user perceptions and potentially reinforce echo chambers. The observed right-leaning bias for neutral accounts suggests that users new to the platform or those without strong pre-existing preferences may be more exposed to conservative viewpoints. This could subtly affect user perspectives, especially for those less routinely engaged with political content or discourse.

Another key observation is that, unlike prior research, which has primarily examined the amplification of tweets from media outlets and political figures—especially elected legislators from major political parties [14]—our findings reveal that X’s algorithm now also amplifies political commentators and influencers. This trend is most pronounced in the recommendations for neutral accounts, suggesting a shift in the algorithm’s prioritization toward these types of voices. This shift could be influenced by recent claims that X prioritizes verified and paid subscription accounts, potentially amplifying influencers who invest in these platform features. The prominence of these non-institutional voices in political content raises questions about the influence of individual commentators on public opinion, as their perspectives may carry a more personal or sensational tone compared to traditional media sources. Adding to the concerns, recent investigations uncovered state-sponsored foreign interference operations with financial backing of prominent political influencers.² This underscores the need for further examination into how the recommendation algorithm’s priorities may shape political engagement and public discourse, especially during critical periods like an election year.

In light of these preliminary findings, we emphasize the need for continuous monitoring of algorithmic shifts and their impact on content exposure dynamics. As social media platforms evolve

²*Justice Department Disrupts Covert Russian Government-Sponsored Foreign Malign Influence Operation Targeting Audiences in the United States and Elsewhere* <https://www.justice.gov/opa/pr/justice-department-disrupts-covert-russian-government-sponsored-foreign-malign-influence>

and adjust their algorithms, regular audits are essential to ensure transparency and accountability, particularly when the stakes involve democratic processes and election integrity.

METHODS

Data Collection and Dataset Statistics

Data collection for neutral monitoring accounts began around October 2, 2024, and reached a stable deployment of approximately 30 active neutral accounts per day on October 11. Left-leaning, right-leaning, and balanced accounts began appearing consistently in the dataset around October 7, with each group reaching a stable count of about 30 active accounts per day shortly thereafter. Each neutral account receives approximately 500 tweets per session, while each left-leaning, right-leaning, and balanced account receives around 700 tweets per session. Figures 4 display the number of active accounts and the total tweets collected daily.

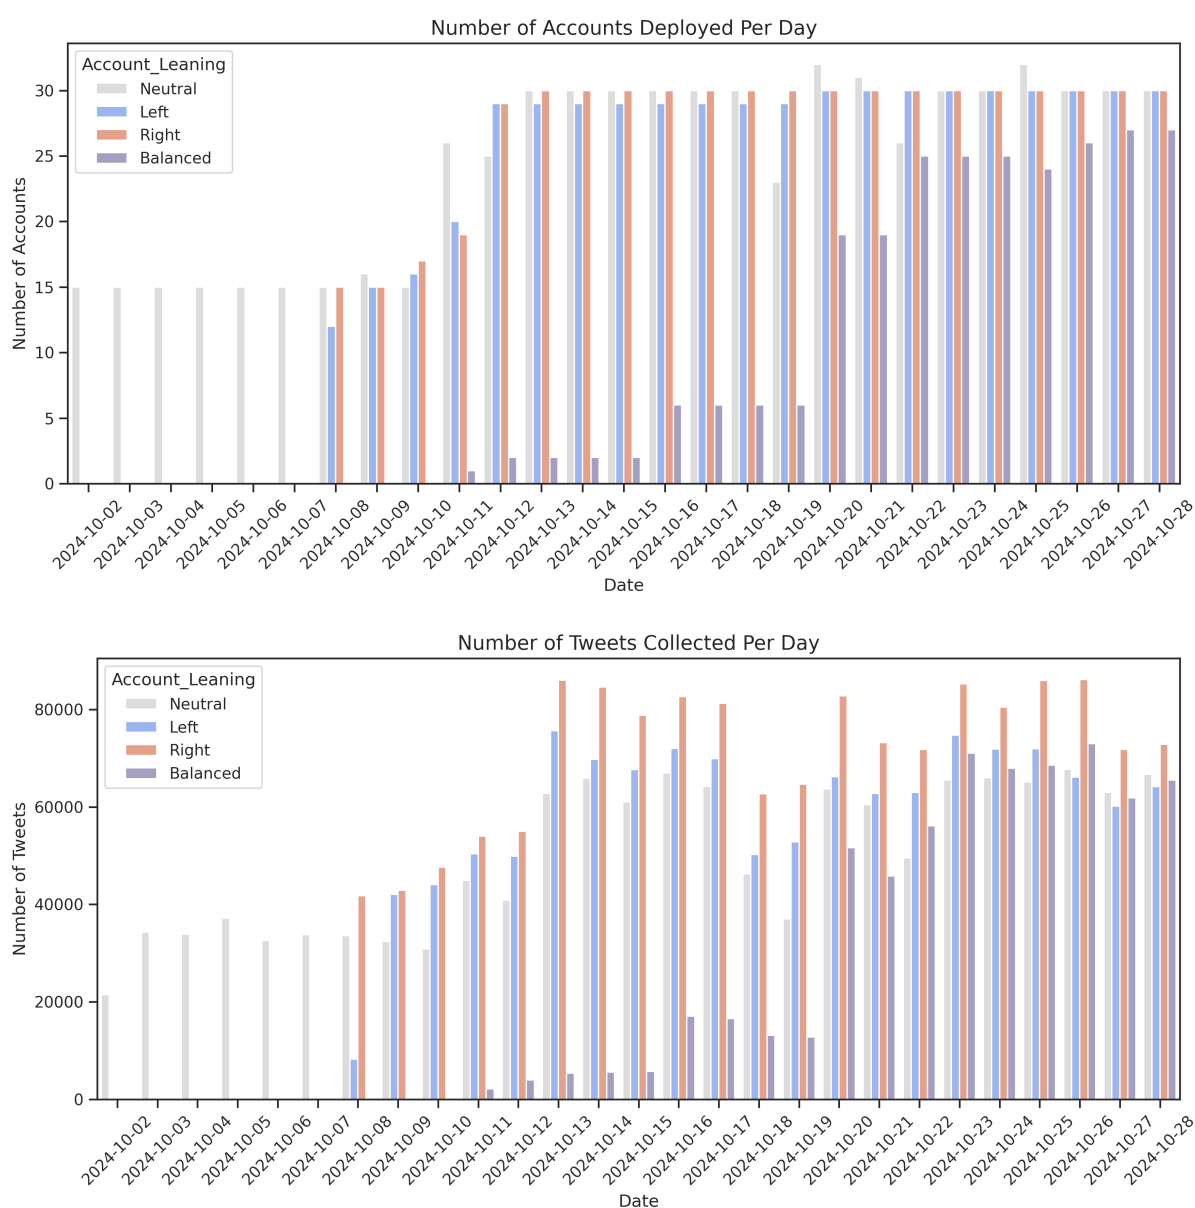


Fig. 4. Overview of data collection: (a) Number of active accounts per day, and (b) Number of tweets collected per day.

Exposure Evaluation Metric

For each \mathbb{X} user whose tweet appears in the personalized timelines of our monitoring accounts, we assign a metric called the “weighted occurrence per 1,000 tweets.” This metric is mathematically expressed as:

$$\text{Weighted Occurrence Per 1K Tweets} = \frac{1}{N} \sum_{i=1}^n p_i \cdot 1000,$$

where p_i is the probability of exposure related to a specific tweet, n denotes the total number of times the user’s tweets appear in the monitoring account’s timeline, and N is the aggregate count of tweets in all timelines collected for the monitoring account.

The probability of exposure, p_i , represents the estimated likelihood that a tweet is seen by a real user. Items near the top of a user’s social media feed are more visible and thus more likely to be viewed. Following prior work on modeling collective attention on social media [17, 23], we employ an exponential decay function, $p(r) = A \cdot e^{-\lambda r}$, to approximate the probability that a tweet at a given rank r in a timeline will be seen. Each tweet in the sequence is assigned a weight that decreases gradually from 1 towards 0, representing the declining probability of user exposure as the tweet’s position moves further down the timeline.

The parameters of the exponential decay function are informed by findings from studies on platforms like TikTok and YouTube [13], which indicate that the top 20% of an account’s videos receive more than 70% of the views. Using this as a reference, we assume that the top 20% of tweets in a timeline similarly capture the majority of user attention, and we calibrate our decay model accordingly. For instance, for a neutral account with an average timeline length of 500, the exponential decay function is defined as:

$$p_{\text{neutral}}(r) = 1.009 \cdot e^{-0.0120 \cdot r}.$$

Gini Coefficient

To measure whether exposure is evenly distributed among users or dominated by a few accounts, we employ the Gini coefficient, a widely used measure to quantify inequality [4]. The Gini coefficient ranges from 0 to 1, where 0 indicates perfect equality (all users have the same exposure) and 1 signifies maximum inequality (exposure is concentrated among a few accounts). In our specific case, the Gini coefficient G is calculated as:

$$G = \frac{\sum_{i=1}^n \sum_{j=1}^n |E_i - E_j|}{2n^2 \bar{E}},$$

where E_i and E_j represent the exposure metrics—weighted occurrence per 1,000 tweets—of users i and j in a monitoring account’s timeline, n is the total number of users, and \bar{E} is the mean exposure metric across all users. A higher Gini coefficient indicates greater inequality in exposure distribution, suggesting that a small number of users dominate exposure in the timeline, while a lower coefficient suggests a more even distribution among users.

Amplification Measure

To assess the (de-)amplification of specific users in relation to left- and right-leaning monitoring accounts compared to a baseline constructed from balanced accounts, we introduce the “mean amplification ratio,” inspired by the work of Huszár et al. [14] on algorithmic amplification.

The mean amplification ratio a_u for a user u , take the example of left-leaning monitoring accounts, is defined by the formula:

$$a_u = \left(\frac{\bar{E}_u^{\text{left}} + 1}{\bar{E}_u^{\text{balanced}} + 1} - 1 \right) \times 100\%,$$

where:

$$\bar{E}_u^{\text{left}} = \frac{1}{|V_{\text{left}}|} \sum_{v \in V_{\text{left}}} E_{v,u},$$

$$\bar{E}_u^{\text{balanced}} = \frac{1}{|V_{\text{balanced}}|} \sum_{v \in V_{\text{balanced}}} E_{v,u}.$$

Here, V_{left} is the set of left-leaning accounts, and V_{balanced} is the set of balanced accounts. $E_{v,u}$ denotes the weighted occurrence per 1,000 tweets of account u in the timelines of account v . This amplification ratio quantifies the extent to which a user’s exposure is increased or decreased when viewed by left-leaning monitoring accounts compared to the balanced baseline. The calculation for right-leaning monitoring accounts follows a similar approach.

ABOUT THE TEAM

The 2024 Election Integrity Initiative is led by Emilio Ferrara and Luca Luceri and carried out by a collective of USC students and volunteers whose contributions are instrumental to enable these studies. Special thanks to Marco Minici and Federico Cinus for comments on this study. The authors are also grateful to the following HUMANS Lab’s members for their tireless efforts on this project: Ashwin Balasubramanian, Leonardo Blas, Charles ‘Duke’ Bickham, Keith Burghardt, Sneha Chawan, Vishal Reddy Chintham, Eun Cheol Choi, Srilatha Dama, Priyanka Dey, Isabel Epistelomogi, Saborni Kundu, Grace Li, Richard Peng, Gabriela Pinto, Jinhu Qi, Ameen Qureshi, Namratha Sairam, Tanishq Salkar, Srivarshan Selvaraj, Kashish Atit Shah, Gokulraj Varatharajan, Reuben Varghese, Jinyi Ye, Siyi Zhou, and Vito Zou. **Previous memos:** [1, 5, 8, 10, 11, 18, 20, 21]

REFERENCES

- [1] Ashwin Balasubramanian, Vito Zou, Hitesh Narayana, Christina You, Luca Luceri, and Emilio Ferrara. 2024. *A Public Dataset Tracking Social Media Discourse about the 2024 U.S. Presidential Election on Twitter/X*. Technical Report. HUMANS Lab – Working Paper No. 2024.6. <https://arxiv.org/abs/2411.00376>.
- [2] Jack Bandy and Nicholas Diakopoulos. 2021. More accounts, fewer links: How algorithmic curation impacts media exposure in Twitter timelines. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–28.
- [3] Nathan Bartley, Andres Abeliuk, Emilio Ferrara, and Kristina Lerman. 2021. Auditing algorithmic bias on twitter. In *Proceedings of the 13th ACM Web Science Conference 2021*. 65–73.
- [4] Nathan Bartley, Keith Burghardt, and Kristina Lerman. 2023. Evaluating content exposure bias in social networks. In *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining*. 379–383.
- [5] Leonardo Blas, Luca Luceri, and Emilio Ferrara. 2024. *Unearthing a Billion Telegram Posts about the 2024 U.S. Presidential Election: Development of a Public Dataset*. Technical Report. HUMANS Lab – Working Paper No. 2024.5. <https://arxiv.org/abs/2410.23638>.
- [6] Paul Bouchaud, David Chavalarias, and Maziyar Panahi. 2023. Crowdsourced audit of Twitter’s recommender systems. *Scientific Reports* 13, 1 (2023), 16815.
- [7] Wen Chen, Diogo Pacheco, Kai-Cheng Yang, and Filippo Menczer. 2020. Neutral bots reveal political bias on social media. *arXiv preprint arXiv:2005.08141* (2020).
- [8] Federico Cinus, Marco Minici, Luca Luceri, and Emilio Ferrara. 2024. *Exposing Cross-Platform Coordinated Inauthentic Activity in the Run-Up to the 2024 U.S. Election*. Technical Report. HUMANS Lab – Working Paper No.

- 2024.7. <https://arxiv.org/abs/2410.22716>.
- [9] Giulio Corsi. 2024. Evaluating Twitter’s algorithmic amplification of low-credibility content: An observational study. *EPJ Data Science* 13, 1 (2024), 18.
- [10] Emilio Ferrara. 2024. *Charting the Landscape of Nefarious Uses of Generative Artificial Intelligence for Online Election Interference*. Technical Report. HUMANS Lab – Working Paper No. 2024.1. <https://arxiv.org/abs/2406.01862>.
- [11] Emilio Ferrara. 2024. *What Are The Risks of Living in a GenAI Synthetic Reality?* Technical Report. HUMANS Lab – Working Paper No. 2024.2. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4883399.
- [12] Andrew Guess, Brendan Nyhan, and Jason Reifler. 2018. Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. *European Research Council* 9, 3 (2018), 4.
- [13] Benjamin Guinaudeau, Kevin Munger, and Fabio Votta. 2022. Fifteen seconds of fame: TikTok and the supply side of social video. *Computational Communication Research* 4, 2 (2022), 463–485.
- [14] Ferenc Huszár, Sofia Ira Ktena, Conor O’Brien, Luca Belli, Andrew Schlaikjer, and Moritz Hardt. 2022. Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences* 119, 1 (2022), e2025334119.
- [15] Jeon-Hyung Kang and Kristina Lerman. 2015. Vip: Incorporating human cognitive biases in a probabilistic model of retweeting. In *Social Computing, Behavioral-Cultural Modeling, and Prediction: 8th International Conference, SBP 2015, Washington, DC, USA, March 31-April 3, 2015. Proceedings 8*. Springer, 101–110.
- [16] Erik Knudsen. 2023. Modeling news recommender systems’ conditional effects on selective exposure: Evidence from two online experiments. *Journal of Communication* 73, 2 (2023), 138–149.
- [17] Zhenpeng Li and Tang Xijin. 2020. Dynamics of online collective attention as hawkes self-exciting process. *Open Physics* 18, 1 (2020), 6–13.
- [18] Marco Minici, Luca Luceri, Federico Cinus, and Emilio Ferrara. 2024. *Uncovering Coordinated Cross-Platform Information Operations Threatening the Integrity of the 2024 US Presidential Election Online Discussion*. Technical Report. HUMANS Lab – Working Paper No. 2024.4. <https://arxiv.org/abs/2409.15402>.
- [19] Dimitar Nikolov, Mounia Lalmas, Alessandro Flammini, and Filippo Menczer. 2019. Quantifying biases in online information exposure. *Journal of the Association for Information Science and Technology* 70, 3 (2019), 218–229.
- [20] Gabriela Pinto, Charles Bickham, Tanishq Salkar, Luca Luceri, and Emilio Ferrara. 2024. *Tracking the 2024 US Presidential Election Chatter on Tiktok: A Public Multimodal Dataset*. Technical Report. HUMANS Lab – Working Paper No. 2024.3. <https://arxiv.org/abs/2407.01471>.
- [21] Kashish Shah, Patrick Gerard, Luca Luceri, and Emilio Ferrara. 2024. *Unfiltered Conversations: A Dataset of 2024 U.S. Presidential Election Discourse on Truth Social*. Technical Report. HUMANS Lab – Working Paper No. 2024.8.
- [22] Twitter, Inc. 2023. Twitter’s Recommendation Algorithm. https://blog.x.com/engineering/en_us/topics/open-source/2023/twitter-recommendation-algorithm.
- [23] Fang Wu and Bernardo A Huberman. 2007. Novelty and collective attention. *Proceedings of the National Academy of Sciences* 104, 45 (2007), 17599–17601.
- [24] Hongli Yuan and Alexander A Hernandez. 2023. User cold start problem in recommendation systems: A systematic review. *IEEE Access* 11 (2023), 136958–136977.

APPENDIX

Top Recommended Users in Left-Leaning, Right-Leaning, and Balanced Accounts

Figure 5, figure 6 and figure 7 display the top 20 recommended users in left-leaning, right-leaning, and balanced accounts, ranked by their average weighted occurrence per 1,000 tweets. Each box in the boxplot represents the distribution of exposure across all accounts in each group, with red indicating right-leaning users and blue indicating left-leaning users. Political leanings of users are inferred based on publicly available data, which may be subject to inaccuracies or changes over time.




Fig. 5. Top 20 recommended users in left-leaning accounts.

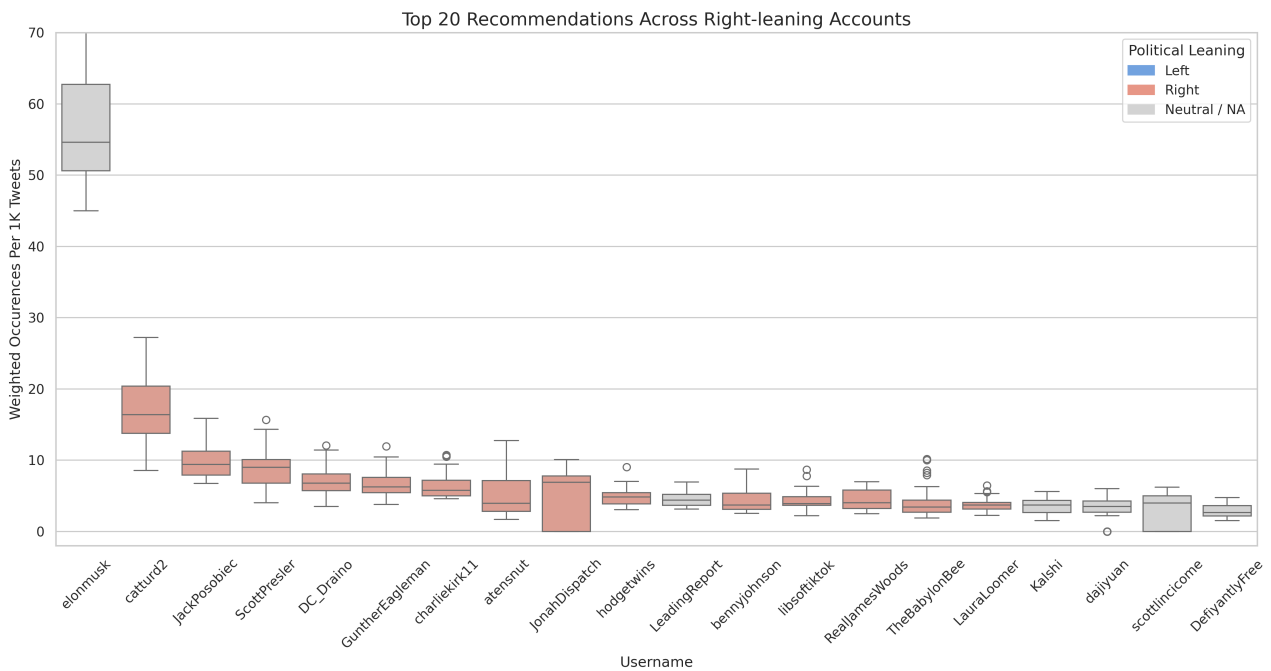


Fig. 6. Top 20 recommended users in right-leaning accounts.

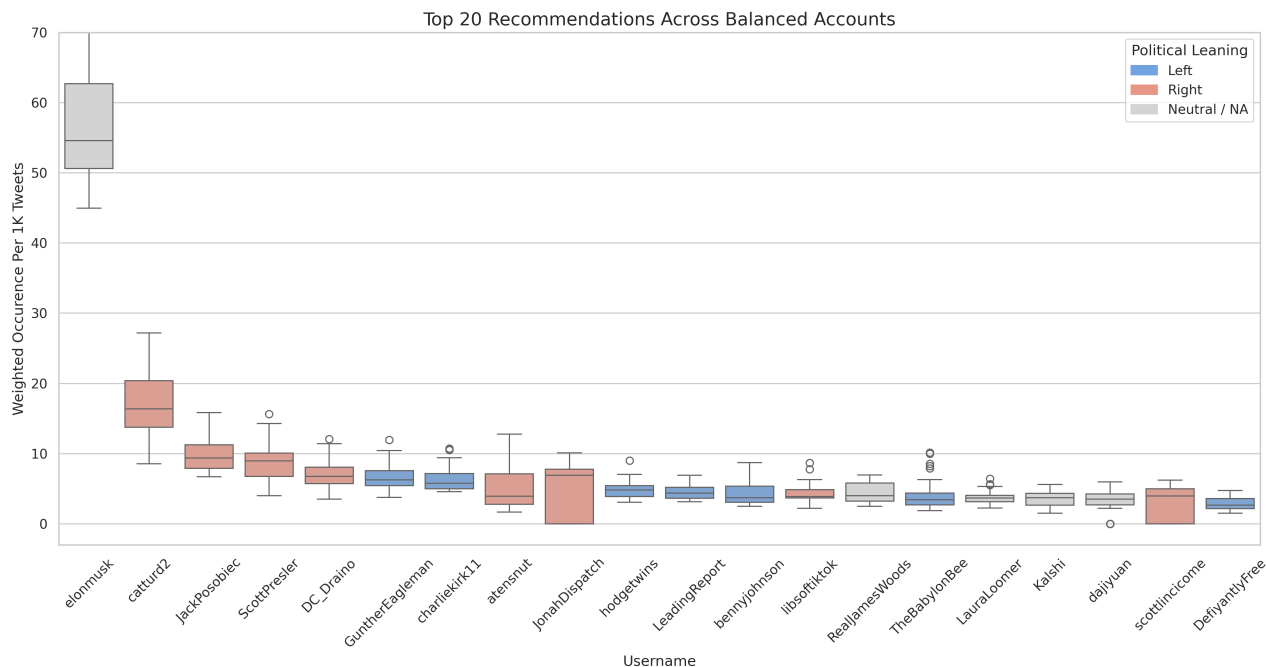


Fig. 7. Top 20 recommended users in balanced accounts.

Profile Information of the Top Recommended Users Across Account Groups

Table 2 presents the profile information of the top 30 out-of-network recommendations across all account groups combined, sorted in descending order by the number of followers for each user. Note that, for left-leaning, right-leaning, and balanced accounts, top recommendations exclude media and politician accounts they already follow.

Table 2. Profile information of top out-of-network recommendations across all groups of accounts

Index	Username	Screensname	Profile Description on X	# of Followers
1	elonmusk	Elon Musk	Read @America to understand why I'm supporting Trump for President	202742780
2	BarackObama	Barack Obama	Dad, husband, President, citizen.	132026578
3	realDonaldTrump	Donald J. Trump	45th President of the United States of America	92023938
4	POTUS	President Biden	46th President of the United States, husband to @FLO-TUS, proud dad & pop.	36825725
5	KamalaHarris	Kamala Harris	Fighting for the people. Wife, Momala, Auntie. She/her. Official account is @VP.	21259465
6	AdamSchefter	Adam Schefter	ESPN Senior NFL Insider. Interview & Podcast Requests: ESPNPR@espn.com Host of the Adam Schefter Podcast https://t.co/oz43ix5jZU	11319193
7	Live	Live		9024660
8	mcuban	Mark Cuban	Dunking on the pharma industry with @costplus-drugs.com, the lowest prices on meds anywhere. check it out !	8959820
9	dbongino	Dan Bongino	Public Enemy #1	5881895
10	historyinmemes	Historic Vids	Daily history lessons. Education through memes!	5451424
11	AMAZINGNATURE	Nature is Amazing	Animals Nature Discovery	4496570
12	RealJamesWoods	James Woods	Please enjoy our inaugural YouTube video about the creation of my album with Shooter Jennings, right here: https://t.co/N1RReBLopn	4271926

Index	Username	Screename	Profile Description on X	# of Followers
13	TheBabylonBee	The Babylon Bee	Fake news you can trust. January 6: The Most Deadliest Day— now streaming!	4217074
14	RobertKennedyJr	Robert F. Kennedy Jr		4110129
15	PeteButtigieg	Pete Buttigieg	Personal account. For official updates, follow @SecretaryPete. Husband, father, veteran, writer, South Bend's former Mayor Pete. (he/him)	3876190
16	charliekirk11	Charlie Kirk	Founder & CEO: @TPUSA • @TPAction_ • Host: The Charlie Kirk Show • Click the link below to subscribe	3689713
17	libsoftiktok	Libs of TikTok	News you can't see anywhere else. submissions@libsoftiktok.com. DM submissions. Bookings: Partnerships@libsoftiktok.com. Subscribe to our newsletter	3619947
18	InternetHOF	internet hall of fame	the internet just wouldn't be the same without these iconic posts.	3360230
19	megynkelly	Megyn Kelly	Happily married to Doug, crazy in love with my children Yates, Yardley, and Thatcher, journalist.	3278907
20	catturd2	Catturd™	The turd you can't flush.	3054117
21	ProjectLincoln	The Lincoln Project	"You cannot escape the responsibility of tomorrow by evading it today." – Abraham Lincoln Home of #The-Breakdown and LP Podcast	2994780
22	hodgetwins	Hodgetwins	Merch & Giveaways at: https://t.co/kxb8qjGCDW — PODCAST: @thetwinspace	2993141
23	bennyjohnson	Benny Johnson	i make internet	2937906
24	JackPosobiec	Jack Posobiec	Sr Editor, @HumanEvents. Veteran Navy intel officer. Catholic. NYT Bestselling Author	2809915
25	gtconway3d	George Conway	President and Executive Director of @PsychoPAC24, the Anti-Psychopath Political Action Committee. President, @chkbal. Contributor, @TheAtlantic.	2396447
26	unusual_whales	unusual_whales	Stocks/Options/Crypto/Market News + Tools. Not advice Get \$50-\$5000 to trade: https://t.co/wGf2ZdlXpw Discord: https://t.co/0xJ9e0ZYYG More: https://t.co/nsxZlPV0pC	1901961
27	PopCrave	Pop Crave	Craving Pop Culture.	1884374
28	DC_Draino	DC_Draino	Rogan O'Handley	1855070
29	DiscussingFilm	DiscussingFilm	Your leading source for quick reliable news. Home for healthy and liberating discussion on all things pop culture. (Amazon links shared may earn us commissions)	1835626
30	ScottPresler	ThePersistence	I helped defeat Hillary, Cheney, & organized the Baltimore cleanup. My goal is to re-elect President Trump. Check out @EarlyVoteAction MAGA MAHA	1776345
31	TheRickWilson	Rick Wilson	Lincoln Project. Award-winning ad-maker. Writer. Instrument-rated pilot. NYT #1 best-seller. Still got the shovel. Writing: https://t.co/e04n749N5H	1698059
32	PopBase	Pop Base	Pop Base is your best source for all pop culture related entertainment, news, award show coverage, chart updates, statistics and more. email@popbase.tv	1683990
33	CollinRugg	Collin Rugg	Co-Owner of Trending Politics Investor American	1561596
34	atensnut	Juanita Broaddrick	Author, "You'd Better Put Some Ice On That" retired RN & business owner, Speaker.	1455919
35	KamalaHQ	Kamala HQ	Providing context.	1416761
36	kylegriffin1	Kyle Griffin	Executive Producer @TheWeekendMSNBC. Opinions mine. Do not congratulate. THREADS @griffinkyle	1409244

Index	Username	Screename	Profile Description on X	# of Followers
37	joncoopertweets	Jon Cooper	Ex: LI Campaign Chair for Barack Obama; National Finance Chair of Draft Biden; Majority Leader of Suffolk County Legislature. Gay dad of 5 kids. #YesWeKam	1391095
38	LauraLoomer	Laura Loomer	Investigative Journalist Free Spirit Founder of LOOMERED. Host of @LoomerUnleashed Former @Project_Veritas operative. America First Feisty Jewess	1364113
39	Tim_Walz	Tim Walz	Running to win this thing with @KamalaHarris.	1311484
40	Riley_Gaines_	Riley Gaines	Host of Gaines for Girls podcast Author of Swimming Against the current TPUSA contributor Director of the Riley Gaines Center	1283591
41	MeidasTouch	MeidasTouch	The official account of the MeidasTouch Network. Unapologetically pro-democracy.	1239469
42	RexChapman	Rex Chapman	It's Hard For Me to Live With Me is available now. For speaking inquiries please contact Jornstein@wmeagency.com	1221843
43	AdamKinzinger	Adam Kinzinger	Proud RINO, dad, Husband, Lt. Col in @AirNatlGuard, CNN Senior Political Commentator, former Congressman, founder @thecountryfirst	1082499
44	Scaramucci	Anthony Scaramucci	Entrepreneur @SkyBridge. Host, Open Book and @Rest-PoliticsUS. https://t.co/t4SOzQjxuy	1077632
45	JoJoFromJerz	Jo	mom. jersey. dem. news junkie. Lebanese. hot-head.views are my own. https://t.co/zueo7YDFWx https://t.co/q4qgmwRLzt . https://t.co/9Fp1kdOX6w	1029714
46	RonFilipkowski	Ron Filipkowski	Editor-in Chief https://t.co/HLS0hEHY1C , Co-host Uncovered, Attorney, Marine, Former Federal and State Prosecutor, Republican Party Insane Asylum Escapee	1021928
47	GuntherEagleman	Gunther Eagleman™	Political Commentator - America First - MAGA - Trump 2024 - Unfiltered	1011785
48	atrupar	Aaron Rupar	journalist. sign up for my newsletter, Public Notice (link below). Powered by @SnapStream (more info: https://t.co/2oHPuuFBnN).	987623
49	Dexerto	Dexerto	The leading source for influencer, streamer, gaming, and viral content	980351
50	cb_doge	DogeDesigner	UX/UI & Graphic Designer at Dogecoin & MyDoge Inc./ Citizen Journalist	935600
51	marceelias	Marc E. Elias	Founder @DemocracyDocket. Chair @EliasLawGroup. My dog's name is Bode.	899982
52	RpsAgainstTrump	Republicans against Trump	Pro-democracy conservatives Republicans fighting Trump & Trumpism. Please support our work: https://t.co/FkmisNic4X	821564
53	MuellerSheWrote	Mueller, She Wrote	DONATE to Kamala Harris: https://t.co/gOvFmy1bYN Subscribe to my FREE newsletter	803719
54	harris_wins	Kamala's Wins	Keeping Score of Kamala Harris' wins. The largest online community supporting soon to be President Kamala Harris	790310
55	LeadingReport	Leading Report	Leading source for breaking news.	630544
56	Angry_Staffer	Angry Staffer	Not a WH Staffer Politics, NatSec, and Snark - Your Mileage May Vary Subscribe to my Patreon newsletter for free: https://t.co/Kj4zTicPyk	609103
57	TristanSnell	Tristan Snell	Lawyer, legal commentator, fighter for democracy. Prosecuted Trump University @ NY AG. Commentator, MSNBC. Creator of book/podcast/newsletter TAKING DOWN TRUMP.	583266

Index	Username	Screename	Profile Description on X	# of Followers
58	7Veritas4	Jack E. Smith	“Whatever you are, be a good one”. Here for people, politics and PARODY. alt @jacksmith22	543575
59	Victorshi2020	Victor Shi	Now—Working on Team Harris-Walz. Writer. Fmr—Host @iGenPolitics_, @JoeBiden, @WhiteHouse, @Precisionstrat, @SKDK. @UCLA 24 English alum. Chicagoan. Views mine.	328827
60	acnewsitics	Alex Cole	Software Engineer & Pilot Progressive Follow @newsitics & https://t.co/Retehye9rD	286475
61	Logically_JC	John Collins	Dad Husband Low-Key Nerd EdD / JD	225758
62	scottlincicome	Scott Lincicome	@CatoInstitute Vice President (Econ/Trade), @DukeLaw adjunct, @TheDispatch newsletter-er. CH RTS. You didn’t read the article, did you? Go @Rangers.	78894
63	EpochTimesChina	The Epoch Times - China Insider	China content of The Epoch Times. Sign up for our China newsletter Read on App: https://t.co/wGG3L4uBaT	63943
64	Kalshi	Kalshi	The first legal way to bet on the election in America.	50930
65	GanJingWorld	Gan Jing World	Video and movie streaming. Join #KindnessIsCool contest & win awards. Connect with friends & family.	32626
66	canlesofficial	Canles	Engineered for walking Comfy & versatile footwear for life’s adventures Breathable, lightweight designs	7407
67	janicehisle	Janice Hisle Epoch Times	Assigned to report on President Trump’s 2024 campaign and related topics. Supporter of free speech. Email tips to janice.hisle@epochtimes.us .	2846