

# Toxic Synergy Between Hate Speech and Fake News Exposure

Munjung Kim\*, Tuğrulcan Elmas, Filippo Menczer

Observatory on Social Media, Indiana University, Bloomington

## Abstract

Hate speech on social media is a pressing concern. Understanding the factors associated with hate speech may help mitigate it. Here we explore the association between hate speech and exposure to fake news by studying the correlation between exposure to news from low-credibility sources through following connections and the use of hate speech on Twitter. Using news source credibility labels and a dataset of posts with hate speech targeting various populations, we find that hate speakers are exposed to lower percentages of posts linking to credible news sources. When taking the target population into account, we find that this association is mainly driven by anti-semitic and anti-Muslim content. We also observe that hate speakers are more likely to be exposed to low-credibility news with low popularity. Finally, while hate speech is associated with low-credibility news from partisan sources, we find that those sources tend to skew to the political left for antisemitic content and to the political right for hate speech targeting Muslim and Latino populations. Our results suggest that mitigating fake news and hate speech may have synergistic effects.

## Introduction

Hate speech is a barrier to inclusivity and harmonious coexistence in a diverse society. The proliferation of misinformation, on the other hand, challenges healthy public dialogue. Although both issues are harmful and should be mitigated on their own, we conjecture that they may also synergistically amplify each other. In this paper, we explore this possible link by asking: “*Is the use of hate speech associated with exposure to fake news?*” We attempt to answer this question on Twitter (later renamed X) by studying two groups of users, *hate speakers* who use hate speech on social media toward certain target populations and *non-hate speakers* who also speak about the same populations but in non-hateful ways. We start from a preexisting dataset in which posts are annotated to contain hate speech or not. We expand this data by (i) identifying the authors of those posts, (ii) collecting the accounts followed by them and their tweets, (iii) identifying the news articles shared in those tweets, and finally (iv) using the credibility of the news sources to measure fake news exposure. We focus on the following research questions:

**RQ1** Are hate speakers exposed to more fake news through their following connections?

Is the relationship between hate speech and fake news exposure affected by...

**RQ2** ... the groups targeted by the hate speakers?

**RQ3** ... the popularity of the fake news?

**RQ4** ... the political orientation of the fake news?

In answering these questions, we found that hate speakers are exposed to a higher proportion of news from low-credibility sources through the accounts they follow when compared to non-hate speakers. This finding holds for hate speech targeting Jews and Muslims; differences for other target groups are not significant. We also found that the difference in exposure to news from low-credibility sources was mainly due to unpopular tweets linking to far-right sources for users targeting Muslims and far-left sources for those targeting Jews.

## Related Work

**Exposure to Fake News.** Fake news pose research challenges to both social science and computer science (Lazer et al. 2018). Previous work tackled their detection by analyzing their spread or exposure. For instance, Grinberg et al. (2019) found that the consumption of fake news was highly concentrated among users and that user age and the number of political links in the news feed were strong predictors of exposure to fake news. Bovet and Makse (2019) reported that during the 2016 U.S. elections, the activity of Clinton supporters was influenced by top influencers spreading traditional news, while Trump supporters influenced the dynamics of the top fake news spreaders. Both studies underline that certain individuals (e.g., older and more conservative voters) are more vulnerable to fake news. We expand upon these studies by showing that individuals who use hate speech are also more susceptible to fake news.

**Use of Hate Speech.** Perera et al. (2023) reported that hate speakers often maintain anonymity, display infrequent account verification and geotagging, and tend to have more followers/following, favorites, group memberships, and statuses than non-hate speakers. Despite these characteristics, engagement with hate content from general Twitter users is low. Cinelli et al. (2021) found that “serial haters” —active

\*Corresponding author. Email: munjkim@iu.edu  
Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

users posting exclusively hateful comments— are uncommon. Mathew et al. (2019) reported that content by hate speakers spreads faster and reaches a broader audience. Hate speakers also exhibit denser connections among themselves, emphasizing the cohesive nature of hate-driven online communities. Yousefi et al. (2023) reported that toxic tweets in different contexts have different engagement as measured by the total number of retweets, likes, quotes, and replies. Vaidya, Nagar, and Nanavati (2024) observed that toxicity increases over time on Twitter and developed a model to capture this pattern.

**Association Between Hate Speech and Fake News.** Noor et al. (2023) analyzed COVID-19 vaccination conversations on Twitter and observed a significant presence of conspiracy theories among individuals with highly toxic attitudes. Mosleh, Cole, and Rand (2024) delved into the use of generally harmful language among Twitter users who share misinformation. Posts containing links to lower-quality news outlets tend to contain more harmful language, and false headlines are more likely to include harmful language compared to true headlines. Moreover, Twitter users who share links to low-quality news sources also tend to use more harmful language. Rather than analyzing harmful language by people who share misinformation, our present investigation focuses on exposure to misinformation by users who produce hate speech. Following someone means subscribing to their updates or posts on social media platforms, indicating a level of interest or affinity with the content they produce. Therefore, studying the posts generated by the users followed by the hate speakers can elucidate the type of information that the hate speakers choose to engage with.

## Data and Methodology

We aim to study the correlation between the use of hate speech and exposure to low-credibility news. To this end, we collect three datasets from Twitter: (i) users who do or do not produce hate speech, (ii) posts by the accounts followed by those users, and (iii) the credibility of those posts.

**User Dataset.** While hate speech encompasses many forms of discrimination including gender, sexual orientation, and disability among others, our analysis primarily focuses on racism-related hate speech due to the opportunity presented by the availability of two existing annotated hate speech datasets. The first dataset comprises 5,880 tweets including the keywords “Asians,” “Blacks,” “Jews,” “Latinos,” and “Muslims.” These tweets were annotated by 120 human annotators. Each tweet was annotated by 3–7 people, depending on the keyword. As a result, 357 tweets (6.1%) were labeled as hate speech against minority groups (Jikeli et al. 2023). In cases of disagreement, the labels were selected based on the majority of the annotators. The second dataset comprises 6,941 tweets related to Jews including derogatory terms,<sup>1</sup> spanning the period from January 2019 to December 2021. Each tweet was annotated by two experts, out of a pool of 10 different experts. Among these tweets, 1,250 (18%) were classified as hate speech, adhering to the IHRA

<sup>1</sup>We do not list the terms here as they are offensive.

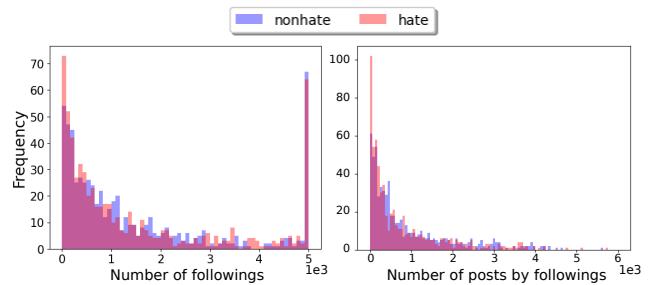


Figure 1: The distributions of the number of followings (left) and posts by followings (right) of hateful and non-hate speakers. The peak in the distribution of the number of followings at 5,000 is attributed to our imposed limitation, restricting the number of followings to a maximum of 5,000.

definition of antisemitic content (Jikeli, Karali, and Soemer 2023). In cases of disagreement, the annotators reached a consensus through discussion. After combining these two datasets, we extracted two sets of users: one that mentions hate speech and one that does not, despite discussing the same topic. Specifically, we randomly sampled 630 users who authored at least one tweet annotated as hate speech. We label these users as *hate speakers*. 444 of them targeted Jews, 68 targeted Blacks, 74 targeted Muslims, 23 targeted Latinos, and 21 targeted Asians. We also sampled from this dataset 630 users who did not post hate speech as the control group, labeling them as *non-hate speakers*. 502 of them mentioned keywords about Jews, 35 about Blacks, 34 about Muslims, 34 about Latinos, and 25 about Asians.

**Exposure Dataset.** We assume that the posts by the users followed by hate and non-hate speakers are their main sources of exposure to low-credibility news. As such, we use those posts as a proxy for exposure. Thus, we first collected the accounts that are followed by the users in our dataset. Due to Twitter API limits, we were able to collect up to the most recent 5,000 followings for both hate and non-hate speakers (52 hate speakers and 53 non-hate speakers had more than 5,000 followings). We then collected the tweets that these followed accounts had authored in the previous 100 days. For each hate/non-hate speaker, the 100-day period preceded their post (or, if they posted more than one hate/non-hate tweet, it preceded one of these posts selected at random). As a result, we collected 48,637,629 posts from 943,671 users followed by 630 hate speakers and 55,435,787 posts from 960,752 users followed by 630 non-hate speakers. The distributions of the number of users followed by each hate speaker and the number of posts produced by those users are shown in Figure. 1.

**Post Credibility.** As it is challenging to evaluate the credibility of each post with high accuracy, we focus on posts that link to news sources. We use source credibility as a proxy for post credibility. We use source credibility scores provided by NewsGuard, which evaluates news and websites using nine journalistic criteria. These criteria are converted into a *trust score* ranging from 0 to 100. Sources with score be-

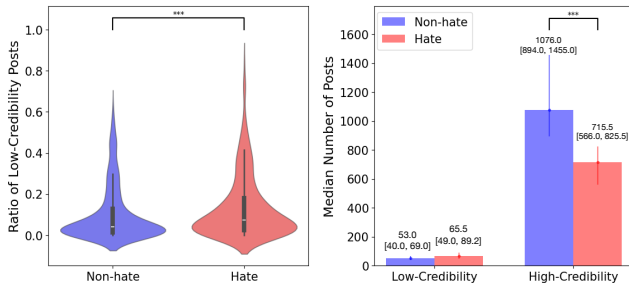


Figure 2: Left: Distributions of the ratio of low-credibility posts to which hate and non-hate speakers are exposed. Right: Median numbers of posts linking to low-credibility and credible news sources to which hate and non-hate speakers are exposed. Error bars indicate 95% Bias-Corrected and Accelerated (BCa) confidence interval (Efron 1987). Significant differences in this and the following figures are indicated (\*\*\*:  $p < 0.001$ , \*\*:  $p < 0.01$ , \*:  $p < 0.05$ ).

low 60 are considered unreliable because they fail to adhere to several basic journalistic standards (NewsGuard 2022). In this paper, we classify these sources as low-credibility. We also use the political alignment (far-left, slight-left, slightly-right, far-right) of the sources, also provided by NewsGuard. Among our 104,073,416 posts, 3,512,022 (3.4%) linked to one of 5,224 distinct sources annotated by NewsGuard. We focus on users exposed to at least 10 tweets linking to these sources, excluding 40 hate and 35 non-hate speakers.

## Results

### RQ1: Exposure of Hate vs. Non-Hate Speakers

We first investigate whether hate speakers are exposed to more fake news compared to non-hate (control) speakers. For each user, we collect the tweets to which they are exposed and compute the proportion linking to low-credibility sources. The median proportion of low-credibility posts is 0.073 for hate speakers and 0.040 for the control group. The distributions are shown by the violin plots in Fig. 2. To statistically compare the proportions of low-credibility tweets to which hate and non-hate speakers are exposed, we use two-tailed Mann-Whitney U tests. Since the distributions of the ratios of low-credibility posts are highly skewed, the Mann-Whitney U test is appropriate as it does not assume normal distributions. The proportion for hate speakers is significantly higher than for non-hate speakers ( $p < 0.001$ ).

The proportion of low-credibility news masks the differences in the actual exposure values. Therefore, let us compare the numbers of posts linking to news articles from sources with low and high credibility to which users in the two groups are exposed. The median numbers of posts are also shown in Fig. 2. We find that the number of posts linking to low-credibility sources for hate speakers is not significantly higher than for non-hate speakers ( $p = 0.36$ ). However, a significant difference exists when we compare the numbers of posts that link to credible news sources ( $p < 0.001$ ). These results suggest that hate and non-hate speakers are exposed to similar (small) volumes of low-

credibility news but hate speakers are exposed to fewer posts linking to credible news sources.

### RQ2: Hate Speakers with Different Targets

We reported that hate speakers are exposed to a higher proportion of news from low-credibility sources. In this section, let us test if this result holds for hate speakers targeting different populations. Fig. 3 presents the distributions for each target population. We compare hate and non-hate speakers who posted tweets including the same population keywords (e.g., “Asians”) to make the control group contextually relevant in each case. Mann-Whitney U tests show that only hate speakers targeted Jews ( $p < 0.001$ ) and Muslims ( $p < 0.01$ ) are exposed to a significantly higher proportion of low-credibility tweets than the corresponding non-hate speakers. This suggests that the pattern observed across all hate speech is mainly driven by data about anti-semitic and anti-Muslim content; this may be partly explained by the predominance of anti-semitic content in our dataset.

Next, we compare the raw numbers of posts with low- and high-credibility news to which users are exposed for each target population. The results are also shown in Fig. 3. We observe that hate speakers targeting Jews are exposed to significantly fewer credible posts than non-hate speakers ( $p = 0.036$ ). Hate speakers targeting Latinos have significantly lower exposure to credible tweets ( $p < 0.01$ ) than non-hate speakers, but also significantly lower exposure to low-credibility tweets ( $p = 0.02$ ). We do not observe significant differences for the remaining target populations; this is possibly due to the sparsity of annotated hate speech targeting those populations in our dataset.

### RQ3: Exposure to Popular Fake News

We next investigate the popularity of the low-credibility news to which hate speakers are exposed. Let us define *popular posts* as low-credibility tweets with more than 500 engagements (likes + retweets) and *unpopular posts* as those with less than 10 engagements. Hate speakers were exposed to 190,911 popular posts and 685,212 unpopular posts, while non-hate speakers were exposed to 249,194 popular posts and 981,478 unpopular posts. While the thresholds are somewhat arbitrary, we considered different thresholds: 1,000 for popular and 20 or 50 for unpopular posts; the results presented below are robust. We compute the proportion of low-credibility posts in the two popularity categories, for each user. Fig. 4 shows that when we focus on popular low-credibility posts, the difference between hate and non-hate speakers is not significant ( $p = 0.22$ ). However, we observe a significant difference for unpopular low-credibility posts ( $p < 0.001$ ). This suggests that hate speakers are more likely to be exposed to low-credibility content that is *not* spreading virally. While this content may reach fewer users, it is also less likely to be fact-checked. This highlights a potential challenge in fact-checking, emphasizing the need to address misinformation across various levels of online visibility.

### RQ4: Exposure to Political Fake News

We finally study the connection between the use of hate speech and exposure to low-credibility news with differ-

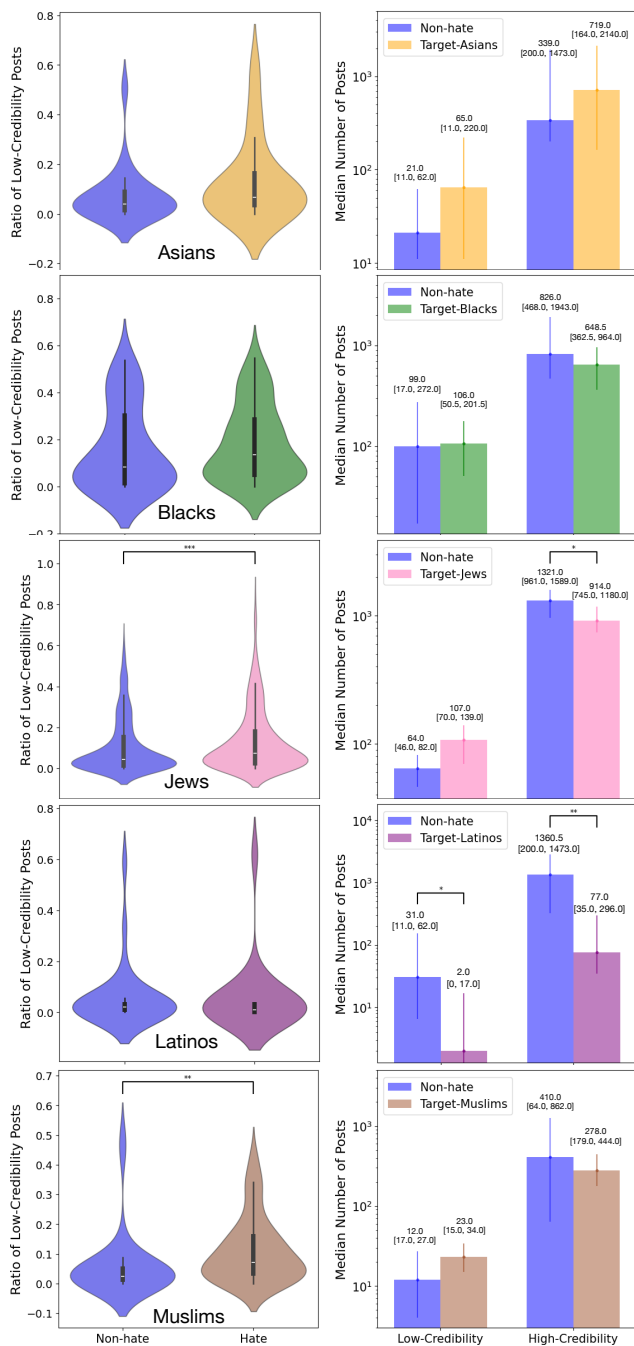


Figure 3: Left: Proportions of posts linking to low-credibility news sources to which hate and non-hate speakers are exposed, for each population targeted by hate speech. Right: Median numbers of posts linking to low-credibility and credible news sources to which hate and non-hate speakers are exposed, for each target population.




Figure 4: Median proportions of tweets linking to low-credibility news sources to which hate and non-hate speakers are exposed, broken down by popularity: less than 10 vs. more than 500 likes and retweets.

ent political alignment. Since such a connection with political alignment can depend on the target population of hate speech, we compare hate and non-hate speakers overall as well as for each targeted population.

The median proportions of low-credibility tweets from news sources in four political alignment categories are reported in Fig 5. Overall, we observe that hate users tend to be exposed to a significantly higher proportion of low-credibility news from far-left ( $p < 0.001$ ), slightly right ( $p = 0.011$ ), and far-right ( $p = 0.012$ ) partisan sources. The proportions are higher from conservative than liberal sources for both hate and non-hate speakers. Looking at specific target populations, individuals producing anti-semitic content are exposed to significantly higher proportions of low-credibility news from far-left ( $p < 0.001$ ) and slightly left ( $p = 0.018$ ) sources. Conversely, those engaging in hate speech against Muslims show a significantly higher exposure to slightly right ( $p = 0.030$ ) and far-right ( $p < 0.001$ ) low-credibility sources. We also note higher exposure to slightly right ( $p = 0.035$ ) and far-right ( $p = 0.056$ ) low-credibility sources for hate speakers targeting Latinos

## Discussion

In this paper, we study the relationship between the use of hate speech and exposure to fake news on social media. We compare the credibility of the posts exposed to 630 hate speakers and 630 non-hate speakers by collecting all the tweets posted by the accounts they follow. We find that hate speakers are exposed to significantly higher proportions of tweets linking to low-credibility news sources, which holds for hate speech targeting Muslims and Jews. We also find

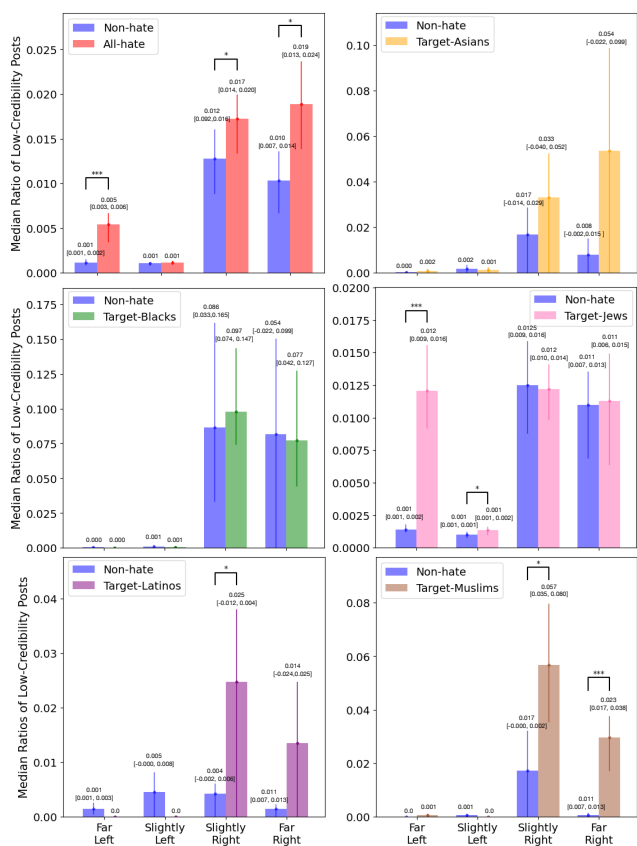


Figure 5: Median proportions of low-credibility news to which hate and non-hate speakers are exposed, broken down by the political alignment of the source for different targets.

that hate speakers are exposed to higher proportions of unpopular posts linking to low-credibility news sources, while both groups are exposed to roughly the same proportion of popular fake news. Finally, anti-semitic speakers are exposed to significantly more posts featuring low-credibility news sources with far-left political alignment, while hate speakers targeting Muslims are exposed to significantly more low-credibility posts from conservative sources.

Although our study design ensures that exposure to fake news occurs *before* the use of hate speech, we cannot prove a *causal* link from fake news to hate speech. There are multiple possible explanations for the observed correlation between hate speech and low-credibility content. One is that the higher exposure to misinformation can exacerbate or reinforce existing prejudices and biases, thereby fueling the propagation of hate speech targeting minority groups. In this scenario, the consumption of misinformation acts as a catalyst for the expression of hateful sentiments, creating a feedback loop where exposure to biased or inaccurate information reinforces and amplifies discriminatory attitudes. Another possible interpretation of our results is that hate speakers may choose to engage with low-credibility content. Because engagement with users on social media platforms reflects their active interest and preferences, hate speakers may

be selectively exposed or drawn to low-credibility news outlets due to their own inclinations or ideological alignment. Future research could focus on establishing a causal link between fake news and hate speech by conducting experiments that manipulate exposure to misinformation and observe its effect on the spread of hateful rhetoric over time.

Our analysis has several limitations. (i) The reliance on Twitter data may not represent the diversity of user behaviors across various social media platforms. (ii) Our control group might still include hate speakers, as we did not analyze all of their posts. (iii) We primarily focus on racism-related hate speech due to the data availability. Hate speech, however, encompasses diverse forms of discrimination based on other attributes. Further analysis is needed to generalize our results to all forms of hate speech. (iv) Our use of content posted by followed users as a proxy for exposure has limitations: on the one hand, not all content posted by followed accounts may be seen by users, and on the other hand, users may be exposed to content via other mechanisms, such as search and recommendations. While reconstructing actual news feeds would have been ideal, it was unfortunately impossible even when the Twitter API was still available to researchers. Finally, (v) our fake news analysis relies on credibility labels at the source level. Given that low-credibility news outlets can occasionally produce accurate content, additional analysis at the individual article level could provide a more comprehensive understanding of fake news.

Despite such limitations, this study contributes to our understanding of the interplay between hate speech and exposure to fake news, which can help inform interventions aiming to mitigate either hate speech or harmful misinformation. The finding that hate speakers are exposed to lower volumes of content from credible news sources emphasizes the importance of promoting reliable information, providing alternatives to interventions that focus solely on the prevalence of misinformation.

**Ethical Impact.** Due to privacy concerns, we limited our analysis to the tweets public profiles posted and were exposed to, avoiding the use of identifiable information. We only reported aggregate results. We also do not share the data due to privacy and changes in X/Twitter data collection and sharing policies. We acknowledge the sensitivity of the topic under study and strive to maintain inclusive language to the best of our ability. A potential negative impact of this study may be the alienation of people who used hate speech by labeling them as hate speakers. Our objective is not to target those people but to understand the factors associated with their harmful behavior, in order to mitigate such behavior. We emphasize that unsupported conclusions, such as claiming that anyone who is exposed to fake news is a hate speaker, are not in line with our research. We discourage any such misuse of our work.

**Acknowledgements.** This work was supported in part by the Knight Foundation.



## References

- Bovet, A.; and Makse, H. A. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature communications*, 10(1): 7.
- Cinelli, M.; Pelicon, A.; Mozetič, I.; Quattrociocchi, W.; Novak, P. K.; and Zollo, F. 2021. Dynamics of online hate and misinformation. *Scientific reports*, 11(1): 22083.
- Efron, B. 1987. Better bootstrap confidence intervals. *Journal of the American statistical Association*, 82(397): 171–185.
- Grinberg, N.; Joseph, K.; Friedland, L.; Swire-Thompson, B.; and Lazer, D. 2019. Fake news on Twitter during the 2016 US presidential election. *Science*, 363(6425): 374–378.
- Jikeli, G.; Karali, S.; Miehling, D.; and Soemer, K. 2023. Antisemitism on Twitter: A Dataset for Machine Learning and Text Analytics. Zenodo. Doi:10.5281/zenodo.7932888.
- Jikeli, G.; Karali, S.; and Soemer, K. 2023. Hate Speech and Bias against Asians, Blacks, Jews, Latines, and Muslims: A Dataset for Machine Learning and Text Analytics. Zenodo. DOI:10.5281/zenodo.8147308.
- Lazer, D. M.; Baum, M. A.; Benkler, Y.; Berinsky, A. J.; Greenhill, K. M.; Menczer, F.; Metzger, M. J.; Nyhan, B.; Pennycook, G.; Rothschild, D.; et al. 2018. The science of fake news. *Science*, 359(6380): 1094–1096.
- Mathew, B.; Dutt, R.; Goyal, P.; and Mukherjee, A. 2019. Spread of hate speech in online social media. In *Proceedings of the 10th ACM Conference on Web Science*, 173–182.
- Mosleh, M.; Cole, R.; and Rand, D. G. 2024. Misinformation and harmful language are interconnected, rather than distinct, challenges. *PNAS Nexus*, 3(3): pgae111.
- NewsGuard. 2022. Rating Process and Criteria. <https://www.newsguardtech.com/ratings/rating-process-criteria/>. Accessed: 2023-09-30.
- Noor, N. B.; Yousefi, N.; Spann, B.; and Agarwal, N. 2023. Comparing toxicity across social media platforms for COVID-19 discourse. Preprint 2302.14270, arXiv.
- Perera, S.; Meedin, N.; Caldera, M.; Perera, I.; and Ahangama, S. 2023. A comparative study of the characteristics of hate speech propagators and their behaviours over Twitter social media platform. *Heliyon*, 9(8): e19097.
- Vaidya, A.; Nagar, S.; and Nanavati, A. A. 2024. Analysing the Spread of Toxicity on Twitter. In *Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD)*, 118–126.
- Yousefi, N.; Noor, N. B.; Spann, B.; and Agarwal, N. 2023. Towards Developing a Measure to Assess Contagiousness of Toxic Tweets. In *Workshop Proceedings of the 17th International AAAI Conference on Web and Social Media*.

## Ethics Checklist

1. For most authors...
  - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes, please see Introduction, Discussion and Ethical Impact.**
  - (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes.**
  - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, please see Introduction.**
  - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes, we explicitly state that the distribution of user groups targeting different populations is skewed towards the group targeting Jews in the Data and Methodology section. So we break down our analysis to different target populations in RQ2 and RQ4.**
  - (e) Did you describe the limitations of your work? **Yes, please see Limitations at the end of Discussion.**
  - (f) Did you discuss any potential negative societal impacts of your work? **Yes, please see Ethical Impact section.**
  - (g) Did you discuss any potential misuse of your work? **Please see Ethical Impact Section.**
  - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, we opted out of sharing the data publicly and presenting only aggregate results as we discussed in Ethical Impact Section.**
  - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes.**
2. Additionally, if your study involves hypotheses testing...
  - (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
  - (b) Have you provided justifications for all theoretical results? **NA**
  - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
  - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
  - (e) Did you address potential biases or limitations in your theoretical framework? **NA**
  - (f) Have you related your theoretical results to the existing literature in social science? **NA**
  - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**

3. Additionally, if you are including theoretical proofs...
  - (a) Did you state the full set of assumptions of all theoretical results? NA
  - (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? NA
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? NA
  - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? NA
  - (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? Yes, we used a preexisting dataset and cited the creators.
  - (b) Did you mention the license of the assets? No, the dataset requires attributions and we cited the authors.
  - (c) Did you include any new assets in the supplemental material or as a URL? No.
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? No, the data only consist of public posts.
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, but we are not using user data, only the posts they authored or exposed to, which we state in Ethical Impact.
  - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? NA
  - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? NA
6. Additionally, if you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots? NA
  - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
  - (d) Did you discuss how data is stored, shared, and de-identified? NA