

Hated More: Online Violence Targeting Women of Color Candidates in the 2024 US Election

Authored by

Dhanaraj Thakur, *Research Director at the Center for Democracy & Technology*

Müge Finkel, *Director, Ford Institute for Human Security, GSPIA, University of Pittsburgh*

With contributions by

Steven Finkel, *Daniel Wallace Professor of Political Science, University of Pittsburgh*

Firat Duruşan, *PhD, Center for Computational Social Sciences, Koç University*

Erdem Yörük, *Director, Center for Computational Social Sciences, Koç University*

Isik Topçu, *MA, Center for Computational Social Sciences, Koç University*

Melih Can Yardi, *MA, Center for Computational Social Sciences, Koç University*

Amanda Zaner, *MID, GSPIA, University of Pittsburgh*

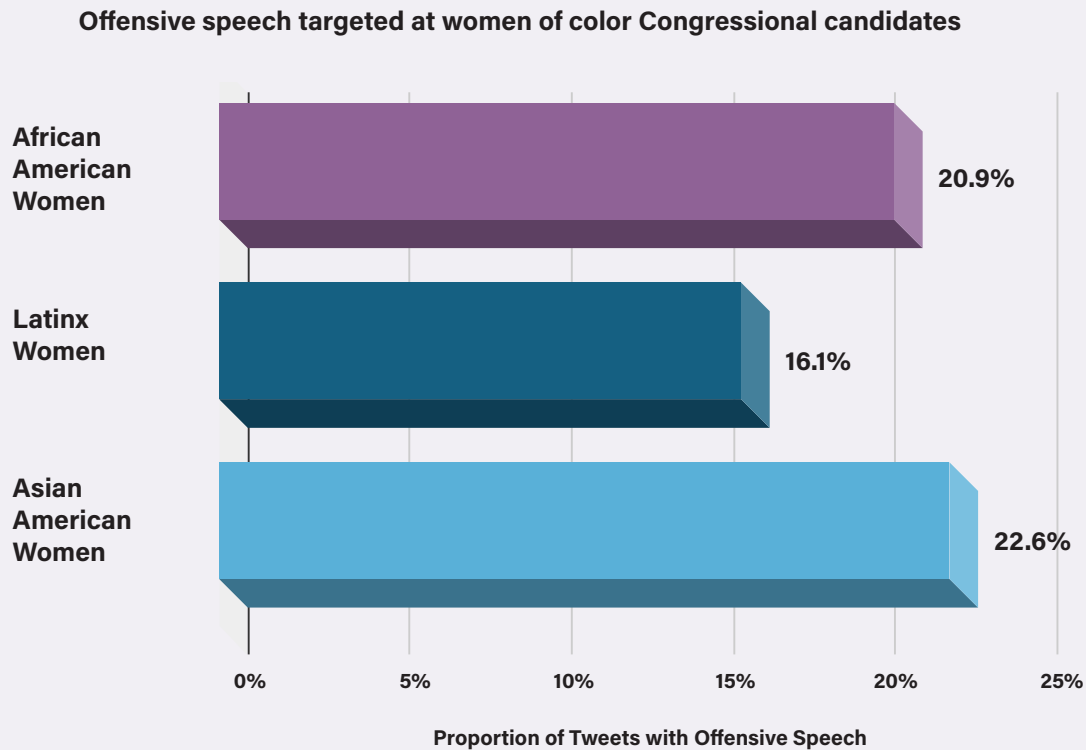
“We must never forget how easily hate speech can turn to hate crime; how ignorance or indifference can lead to intolerance; or how silence in the face of bigotry is complicity.”

- UN Secretary General, [July 2022 on X](#)

Women, and women of color in particular, face [numerous challenges](#) when running for political office in the U.S. These include attacks they are subject to in various online spaces that, like their peers, they must use to campaign and promote their work. [These attacks often aim](#) to undermine and prevent [women’s participation in politics](#). Previous [research by CDT](#) found that women of color Congressional candidates in the 2020 U.S. election were more likely to be subjected to violent and sexist abuse, and mis- and disinformation on X/Twitter compared to other candidates. These forms of abuse might contribute to the underrepresentation of women of color in politics, and may also undermine the effectiveness of the US democratic system in reflecting the interest and priorities of all voters in policy-making.

In this research brief, we turn to the 2024 U.S. elections to examine the nature of offensive speech and hate speech that candidates running for Congress are subject to on the social media platform X (formerly Twitter), which remains an important forum for political candidates. More specifically, we compare the levels of offensive speech and hate speech that different groups of Congressional candidates are targeted with based on race and gender, with a particular emphasis on women of color. We also examine these factors for U.S. Vice President Kamala Harris as a woman of color and presidential candidate.

For purposes of this research, we identified all tweets posted between May 20 and August 23, 2024 that mentioned any candidate running for Congress with one or more accounts on X (a total of 1031 candidates). This resulted in a dataset of over 800,000 tweets. Additionally, we examined tweets that mentioned Harris during this period. Using several fine-tuned language models, we identified tweets that contained offensive speech or hate speech (see the methods section for more details) about the candidate.



▲ **Figure 1. Proportion of Offensive Speech targeted at women of color candidates.**

We define offensive speech as words or phrases that demean, threaten, insult, or ridicule a candidate. We define hate speech as a subset of offensive speech where specific reference is made to someone’s identity including race, gender, sexual orientation, or religion. Our findings show that women of color and African American women candidates in particular are subject to more offensive speech overall, and specifically to more hate speech, than other candidates.

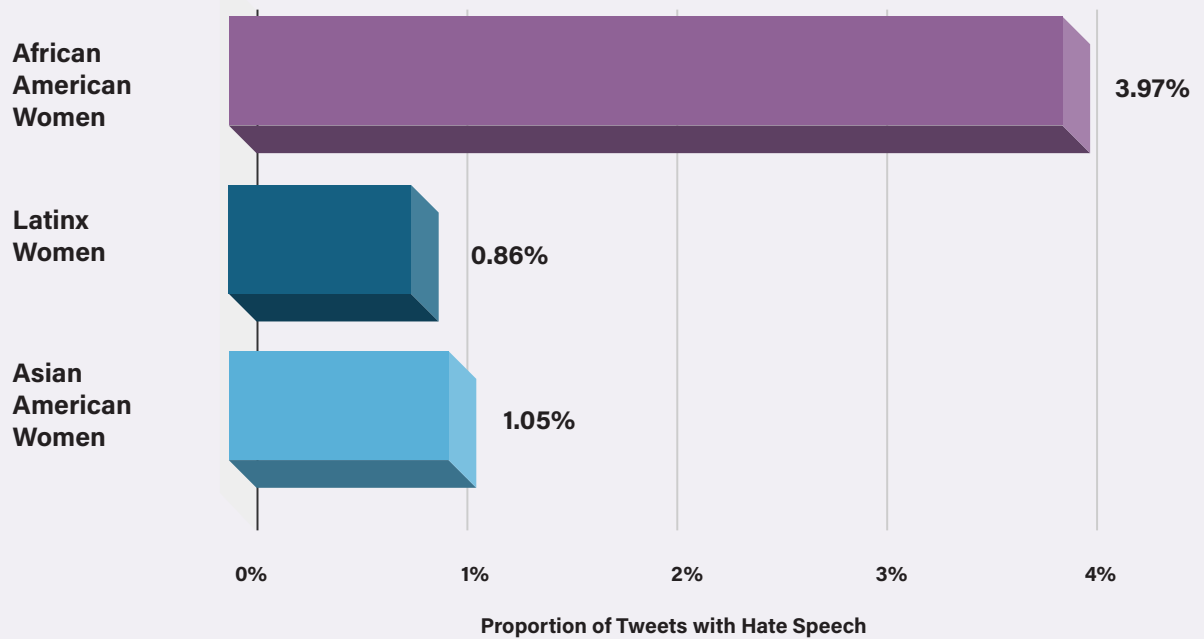
[CONTENT WARNING - Some of the examples to follow include profanity and threats that some may find offensive or triggering.]

Key Findings

- 1. More than 1 in 5 tweets targeted at Asian-American and African-American women candidates contained offensive language about the candidate.**

We found that all groups of candidates were subject to significant levels of offensive speech. These posts included insults and profanity (“@candidate f**k you, you’re a piece of s**t”), were demeaning (@candidate “You are dumb as a rock”), or threatening (@candidate “you deserve no safety or freedom”).

Hate speech targeted at women of color Congressional candidates



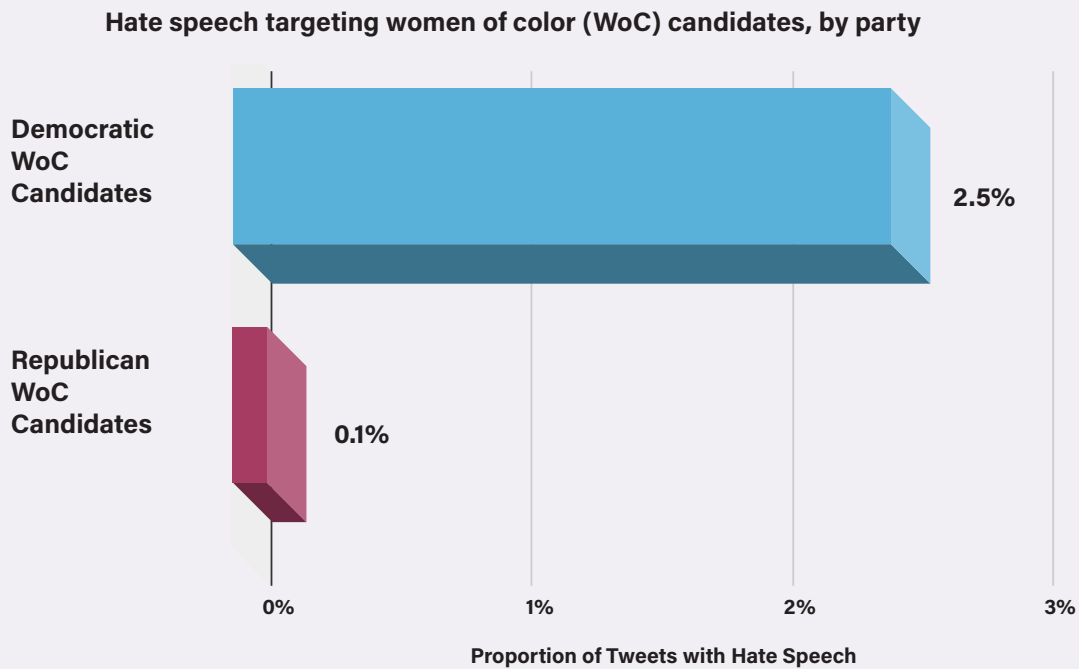
▲ Figure 2. Proportion of Hate Speech targeted at women of color candidates.

On average, 16% of all tweets that mentioned any candidate contained offensive speech. **Some groups were more likely than any other group of candidates to be targeted with this kind of offensive speech, in particular Asian-American (23%) and African-American (21%) women candidates.** White women were subject to less at 18%, and for Latinx women candidates, the proportion was closer to the overall average, similar to that of white men and men of color (16%).

In general women candidates (18%) were more likely to be targeted with offensive tweets than men (15%), although the distribution of different types of offensive tweets (whether they contained profanity, threats, or ridicule) was similar for all candidates, regardless of race or gender.

2. African-American women candidates are targeted with significantly more hate speech than other groups.

As noted previously, hate speech contains references to a candidate's identity such as race (@candidate "get ready to get deported in November"), gender ("We should replace @candidate with a man, see how she likes it."), religion (@candidate "I don't know what Bible you're reading but it's certainly not the Christian Bible. You're pathetic."), and sexual orientation. On average, less than 1% of all tweets that mentioned a candidate contained hate speech. However, we



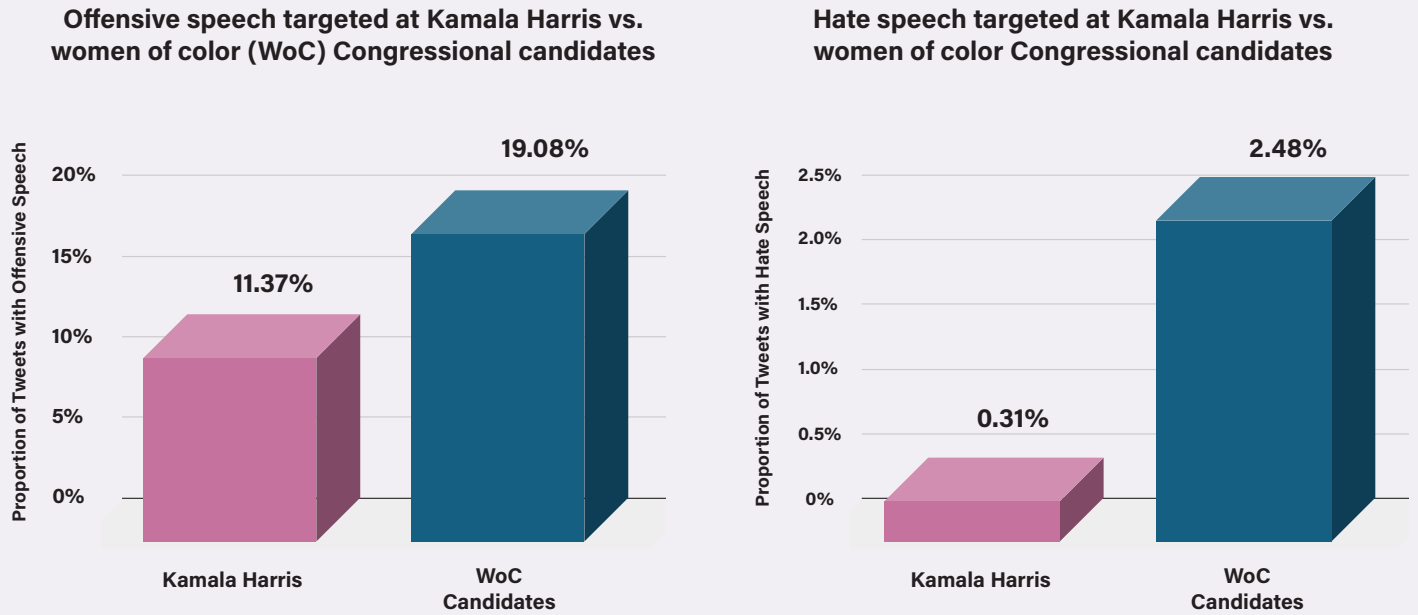
▲ Figure 3. Hate speech targeting women of color (WoC) candidates, by party.

found that African-American women candidates were more likely than any other candidate to be subject to this type of post (4%). In fact, they were more likely to be targeted with hate speech compared to other women of color candidates: 5 times more than Latinx women and 4 times more than Asian-American women.

The differences were also significant compared to other groups. African-American women were 7 times more likely to be targeted with hate speech than African-American men, more than three times as likely as white women, and more than 18 times as likely as white men. This finding is in keeping with prior research that points to the unique and [challenging experiences](#) of African-American women on social media.

3. The proportion of race and gender-based hate speech targeted at women candidates differs by party affiliation.

We found that the proportions of hate speech women candidates are targeted with depend on their party affiliation. While on average Republican women candidates (2.2%) are more likely to be targeted with hate speech than Democratic women candidates (1.4%), the effect is reversed and much higher for Democratic women of color candidates, who are subject to more than 19 times more hate speech than Republican women of color candidates. Hate speech based specifically on racial identity and gender identity is also more prevalent among Democratic women of color candidates than those who are Republican.



▲ Figures 4 and 5. Comparison of the proportion of hate speech and offensive speech targeted at Harris and women of color congressional candidates

Democratic women of color candidates were targeted with more than twice as much gender-based (1.2%) and race-based (.5%) hate speech than Republican women of color candidates. A similar pattern obtains for offensive speech as a whole.

4. The proportion of hate speech and offensive speech targeted toward Kamala Harris – the first woman Vice-President and current candidate for the presidency – is less than that for Congressional women of color candidates.

The proportion of offensive speech that Harris was targeted with was 40% less than women of color candidates overall, and 45% less than African American women candidates. In terms of hate speech, 2.5% of all tweets targeting women of color candidates contained hate speech, 4% for African-American women candidates, and 1% for Asian-American women candidates. **Only 0.3% of all tweets Harris was targeted with contained hate speech.** While further research is needed to understand this pattern, the low counts for Harris could be partially explained by factors such as the time period in which we collected our data, the number of tweets expressing excitement or other positive feelings about her entry into the presidential race during this period, and the possibility that Harris appears to have eschewed a political campaign foregrounding her identity and does not explicitly discuss her gender and race to most audiences.

Conclusion

Our findings suggest that offensive speech targeted at Congressional candidates is not distributed equally, with race and gender being factors in which groups are targeted. This is even more true with hate speech, where a significantly higher proportion of these attacks focus on African-American women candidates compared than any other group. The onus is on social media companies and other stakeholders who are part of the information ecosystem (e.g., AI companies, political parties, journalists, researchers, etc.) to act on the reality that identity mediates how people interact with candidates online.

Due to the disproportionate attacks that may deter women of color and other minorities from running for office, social media platforms should put [specific measures](#) in place as part of their trust and safety systems. This should include clear policies that prohibit attacks against someone based on race or gender, greater transparency into how their systems address these types of attacks, better reporting tools and means for accountability, regular risk assessments with an emphasis on race and gender, and privacy preserving mechanisms for independent researchers to conduct studies using their data. When studying the impact of the information environment on elections, [researchers should take an intersectional approach](#) and center race, and gender as a part of the analysis.

The consequences of the status-quo where women of color candidates are targeted with significant attacks online at much higher rates than other candidates creates an immense barrier to creating a truly inclusive democracy.

Methodology

We manually compiled a list of candidates running for Congressional seats in the 2024 elections, including candidate names, state, party affiliations, gender, race, incumbency status, and X (Twitter) account usernames. To do this we relied on sources such as Ballotpedia, and cross-referenced candidate lists from the Center for Democracy & Technology’s [“An Unrepresentative Democracy”](#) study, along with other public candidate information lists.

For each candidate, we used their X account username as a keyword and retrieved tweets mentioning them, specifying single days between May 20th and August 23rd as search parameters. We filtered out non-English tweets and those with more than three mentions to refine the dataset. No sampling strategy was applied to the Congressional candidates; we collected all relevant tweets for all candidates. We collected a total of 809,296 tweets by automating searches using X’s advanced search feature (<https://x.com/search-advanced>). The dataset includes all US Congressional Candidates with X accounts (1031 in all). For Kamala Harris, we collected 55,330 posts after filtering.

We developed our own offensive and hate speech models, building in part on existing ones. For binary hate speech, offensive speech and offensive sub-types, we developed our own models because existing ones either did not fit our research’s ontology or their performance was lower than needed. For hate speech subtypes, we used a model developed by researchers at UC Berkeley (<https://huggingface.co/wesleyacheng/hate-speech-multilabel-classification-with-bert>).

This process involved fine tuning the following language resources by using a human annotated corpus of training data which we prepared:

- Offensive Speech Binary: cardiffnlp/twitter-roberta-base-offensive ;
- Offensive Speech Types: Profanity: xlm-roberta-base; Threat: cardiffnlp/twitter-roberta-base-offensive ; Demeaning: xlm-roberta-base; Ridicule: xlm-roberta-base
- Hate Speech (Binary): FacebookAI/xlm-roberta-base;
- Hate Speech (Types): wesleyacheng/hate-speech-multilabel-classification-with-bert

Additional papers on the methods used:

- International Conference on Computational Social Science (IC2S2): <https://www.fordinstitute.pitt.edu/working-groups-0/online-violence-against-women-politics>.
- Validation of Politus Data: <https://arxiv.org/abs/2310.17220>.

Finally, we did not include the full tweet in examples provided in this report to protect the privacy of the poster and the candidate involved.



Read more about CDT Research's work at cdt.org/CDTResearch.

Read more about the University of Pittsburgh Ford Institute for Human Security's work at fordinstitute.pitt.edu.

Acknowledgements

We would like to thank Michal Luria, Mona Elswah, DeVan L. Hankerson, Samir Jain, Drew Courtney, Tim Hoagland, Jamal Magby, and Napath Good for their support and inputs throughout this project.

This work is made possible through a grant from the John S. and James L. Knight Foundation to CDT and the Kabak Grant through the Global Studies Center at the University of Pittsburgh to the Ford Institute for Human Security.

Suggested Citation

Thakur, D. and Finkel, M. (2024) Hated More: Online Violence Targeting Women of Color Candidates in the 2024 US Election. Center for Democracy & Technology and the Ford Institute for Human Security, University of Pittsburgh. <https://cdt.org/insights/report-hated-more-online-violence-targeting-women-of-color-candidates-in-the-2024-us-election/>

*The **Center for Democracy & Technology** (CDT) is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1994, CDT has been a trusted advocate for digital rights since the earliest days of the internet. The organization is headquartered in Washington, D.C. and has a Europe Office in Brussels, Belgium.*

*The **Ford Institute for Human Security at the Graduate School of Public and International Affairs at the University of Pittsburgh** was founded in 2003 as one of the first academic centers in the US devoted to the study of human security aligned with the United Nations resolution 66/290, which defines threats to human security to be “widespread and cross-cutting challenges to the survival, livelihood and dignity of people.” The **Online Violence Against Women in Politics Working Group** is an international working group in partnership with Politus Project researchers at the **Center of Computational Social Sciences at Koç University (Istanbul)**. The working group conducts research on a new and growing threat against women in politics on online spaces.*