Toll of Trolls: Gendered Online Hostility and its Impact on Women's Political Ambition

Annie Jarman*

Abstract

Online hostility is an unavoidable part of contemporary political life and influences politicians' behavior. Yet, two important questions remain unresolved: whether women politicians face disproportionate abuse online, and whether such abuse decreases their political ambition. I shed new light on both questions by developing a novel method of distinguishing gendered hostility from general hostility in tweets directed to politicians. I pair this with hand-collected panel data on the career decisions of 1,247 U.S. state legislators across four election cycles. The results counter widespread expectations. Rather than experiencing similar online environments, women politicians encounter twice as much gendered hostility in addition to the general hostility both men and women receive. Even so, online hostility is not driving women incumbents from office. Instead, women are less likely than men to exit office in relation to online hostility. Women's resilience suggests the possibility of a gender-based selection effect, where only women willing to tolerate online hostility will pursue political office.

Keywords: violence against women in politics, ambition, social media, machine learning, representation

*PhD Candidate, Washington University in St. Louis. Email: a.a.jarman@wustl.edu. Web: https://anniejarman.github.io.

In recent years, many women politicians have reported high levels of online abuse and called on social media companies to protect them (Brown, 2021; CCDH, 2024; Frankel, 2020; IPU, 2018; Phillips, 2017). Politicians and political scientists alike theorize that the motive behind such abuse is to push women out of political office. In 2020, a hundred global women lawmakers signed a letter to Facebook where they stated that gendered hostility on the platform was intended to silence women and undermine democracy (Frankel, 2020). A large body of research echos this concern, theorizing that online hostility may be eroding women's representation by raising the costs of holding office (Carson et al., 2024; Erikson et al., 2023; Harmer and Southern, 2023; Herrick and Franklin, 2019; Herrick and Thomas, 2022; Håkansson, 2024; Håkansson and Lajevardi, 2024; Jankowicz, 2017; Kosiara-Pedersen, 2023; Krook, 2017, 2020; Sobieraj, 2018; Thomas et al., 2019; Wagner, 2022).

Despite the attention scholars have dedicated toward understanding online hostility against politicians, two central questions remain. First, do women politicians in fact receive disproportionate abuse on social media platforms? Observational studies and survey-based research yield contrasting results. While politicians' self-reported accounts consistently suggest that women experience more online hostility (Collignon and Rüdig, 2020; Herrick and Thomas, 2022; Herrick et al., 2019; Håkansson, 2021; Pathé et al., 2013; Zeiter et al., 2019), computational text analyses often conclude that men politicians receive a higher volume of online abuse (Fuchs and Schäfer, 2021; Gorrell et al., 2020; Greenwood et al., 2019; Theocharis et al., 2020; Ward and McLoughlin, 2020). I argue that this existing discrepancy stems from a mismatch in what each method captures. While self-reports often reflect the content and impact of abuse, large-scale text analyses generally measure only the volume of hostile messages. Drawing on gender role congruity theory, which frames online hostility as a form of backlash against women who violate traditional gender roles (Eagly and Karau, 2002; Felmlee et al., 2020; Krook and Restrepo Sanín, 2016; Wilhelm and Joeckel, 2019), I hypothesize that gender-motivated abuse toward women politicians will manifest in specifically gendered *content* rather than simply higher *volumes* of abuse (Bardall et al., 2020; Erikson et al., 2023; Holm et al., 2024).

Second, are women politicians disproportionately curtailing their political careers in response to gendered hostility? Survey and interview data suggest that women politicians are more likely than men to reduce their public engagement and even consider exiting office after experiencing online hostility (Erikson et al., 2023; Herrick and Franklin, 2019; Håkansson and Lajevardi, 2024; Ramachandran et al., 2024). Yet a growing body of work suggests that women who enter politics correctly anticipate gendered hostility and discrimination (Fox and Lawless, 2024; Kjøller and Pedersen, 2025; Shames, 2017). Gender-based selection theory suggests the women who do make it into office will be particularly resilient to abuse as a result, treating online hostility as one more hurdle alongside other gendered costs of political careers (Anzia and Berry, 2011b; Bauer, 2020b; Bauer and Cargile, 2023; Butler et al., 2022; Eatough and Preece, 2025; Fulton, 2012b; Lazarus and Steigerwalt, 2018; McGrath et al., 2025; Schmitt and Brant, 2019; Thomsen and Sanders, 2020).

To address these questions, I collect all the public tweets referencing all lower-house state representatives on Twitter between 2015 – 2018. I identify hostile and gendered content using a replicable and cost-effective approach that combines human coding, GPT4 zero-shot classification, and fine-tuning a pre-trained BERT model. I examine the influence of online hostility on ambition by linking measures of exposure to panel data on 1,247 politicians (370 women, 913 men) across four election cycles (2015–2023). Specifically, I measure how much – and what type of – online hostility each politician encountered in the six months prior to the campaign filing deadlines, and whether this influenced their decisions to campaign for reelection, a different office, or not at all.

The results bring clarity to both puzzles. When computational text analysis incorporates gendered language, it produces findings that closely align with women's self-reported experiences (Bjarnegård et al., 2022; Erikson et al., 2023; Kosiara-Pedersen, 2023). I find that although men and women experience similar levels of general hostility, women experience nearly twice as much gender-based hostility. Notably, gender-based hostility intensifies when women politicians deviate further from expected gender norms of behavior, consistent with gender role congruity theory. Yet despite this disproportionate burden, women are significantly less likely than men to exit office in response to either gendered or general hostility. This result is robust across a variety of specifications. Instead, women respond to increased exposure by campaigning for a different office than they currently hold, which may indicate an adaptive strategy to mitigating the effects of online abuse.

This paper makes two primary contributions. First, I revise the widespread view that men and women politicians encounter similar online environments by accounting for the gendered content of abuse. Women receive nearly twice as much gender-based hostility as men, aligning computational analyses with women's self-reported experiences. Second, I show that women's political ambition is more resilient to both gendered and general hostility than is often assumed. Taken together, these findings revise two influential narratives in the literature and offer a more unified account of gendered online violence and ambition: women politicians do face more severe and gender-based hostility than men, but they are less likely than men to exit in response, often adapting instead by seeking new offices.

Visibility and Gender Role Congruity

Although self-reported and observational studies of gendered online hostility often diverge in their findings, they do show consistent agreement that high-profile women receive more abuse than similarly prominent men (Guerin and Maharasingam-Shah, 2020; Holm et al., 2024; Håkansson, 2021; Oates et al., 2019; Rheault et al., 2019). As visibility increases – whether in media appearances, office status, or social media followers – a gender gap in online abuse emerges. This recurring pattern suggests a rare point of alignment between methods, with visibility helping to explain when and why women politicians face disproportionate online hostility.

Role congruity theory offers a valuable framework for understanding why visible women politicians receive more online hostility than their less visible peers (Rheault et al., 2019). In societies that excluded women from public life, the traits linked to competent leadership - assertive, confident, dominant, ambitious – have been historically coded as masculine (Eagly and Karau, 2002; Eagly et al., 1992; Heldman et al., 2018; Katz, 2016). Women who exhibit the desirable "masculine" leadership traits – including political ambition – may be perceived as threats to the standing social order and penalized (Brescoll et al., 2018; Eagly and Karau, 2002; Fulton, 2012a).This role-incongruity helps explain why women politicians experience unique harassment. Women's presence in public life, particularly in positions of power, can be perceived as violating gender norms (Eagly and Karau, 2002). Researchers conceptualize online and offline violence against women in politics (VAWiP) as a form of backlash against women who violate traditional gender roles (Eagly and Karau, 2002; Felmlee et al., 2020; Krook and Restrepo Sanín, 2016; Wilhelm and Joeckel, 2019).

There are at least two mechanisms by which visibility may intensify the perception that women in positions of power are violating gender norms. First, high-profile women may attract more abuse because increased visibility exposes them to a larger audience, heightening the number of observers who choose to act on gender biases (Rheault et al., 2019). In this view, all women politicians engage in role-incongruent behavior, but visibility increases the salience of this deviation and the likelihood that others will respond to it (Krook and Restrepo Sanín, 2016; Mansbridge and Shames, 2008).

Second, some observers may infer that women who attain public notability are substantially more role-incongruent than their less visible peers. Achieving prominence in politics often requires demonstrating high levels of ambition, assertiveness, and confidence: traits historically associated with masculinity (Eagly and Carli, 2003; Folke and Rickne, 2016; Okimoto and Brescoll, 2010). Whether or not prominent women exhibit these traits, their visibility may lead observers to perceive them as more agentic and, therefore, as more norm-violating than less visible women politicians (Håkansson, 2021). In this view, visibility does not merely amplify awareness of a norm violation; it intensifies perceptions of that violation by signaling a deeper departure from stereotypical femininity.

These two mechanisms are observationally equivalent and likely mutually reinforcing. Recognizing that visibility amplifies the gender role-incongruity that triggers backlash helps explain why text-based approaches tend to detect greater hostility toward women only among the most visible politicians (Guerin and Maharasingam-Shah, 2020; Oates et al., 2019; Rheault et al., 2019). Understanding the link between visibility and gender role-incongruity underscores the need for text-based research to engage more fully with online abuse as a form of gendered violence.

Gendered Backlash Takes Gendered Forms

I argue that applying the gender role congruity framework can help resolve the discrepancy between self-reported and text-based findings on whether online hostility toward politicians is gendered. In particular, role congruity theory would suggest the need to look at gendered forms of hostility in abuse targeting women politicians. Violence against women in politics differs from general political violence through gendered motives, forms, or impacts (Bardall et al., 2020). Bardall et al.(2020) argue that a higher volume of online hostility directed at women may signal gendered motives. As discussed above, research comparing the frequency of abuse toward men and women politicians produces divergent findings depending on the methodological approach. Self-reported data often indicate greater volumes of abuse toward women, whereas observational studies frequently find the opposite, with the important exception that such studies consistently detect disproportionate hostility directed at the most visible women in politics (Holm et al., 2024).

In contrast, self-reported data provides strong evidence for gendered forms of abuse (Bardall et al., 2020). Women politicians self-report higher rates of sexualized harassment, gender-based hate speech, and sexist abuse than men (Bjarnegård, 2018; Erikson et al., 2023; Esposito and Breeze, 2022; Herrick et al., 2022; Holm et al., 2024; Kosiara-Pedersen, 2023). Computational text analyses that examine gendered content find the same pattern (Gorrell et al., 2020; Southern and Harmer, 2019; Ward and McLoughlin, 2020). Understanding visibility as a proxy for gender role-congruity helps explain why observational researchers find support for gendered motives in the backlash against high-profile women (Bardall et al., 2020; Rheault et al., 2019). The same logic suggests that gendered forms of abuse are also more likely to appear when women deviate from traditional gender expectations (Bardall et al., 2020; Erikson et al., 2023).

The choice to center criticisms of women politicians on their sex and gender rather than their behavior reflects a deep discomfort with women in power. Gendered hostility is one way to penalize women who deviate from their circumscribed roles by occupying positions of power (Brescoll et al., 2018; Eagly and Karau, 2002). Online violence against women in politics (VAWiP) degrades and dehumanizes women through sexist tropes. Linguists observe that English speakers reduce women to edible, nonhuman, and sexual entities by comparing women to animals ("bitch"), foods ("peach"), or sex workers ("whore") (Montell, 2019). Dovi (2025) categorizes the depictions of women politicians as physically disgusting, morally bankrupt, or contaminated by their presence in politics as "nasty claims". Similarly, Skubic (2025) identifies recurring VAWiP tropes that domesticate, objectify, sexualize, or dehumanize women as animals. Despite the variation in terminology, the central theme to emerge across a plethora of schemas is the incorporation of gender and sex into hostility targeting women. A woman political activist described the abuse as "it's always gender and it's always sexualized, whether or not the thing they're mad about is about gender and sex" (Sobieraj, 2020).

Although gender-based hostility targets individual politicians, the message is about the role of women in politics broadly. This is illustrated clearly in Jane (2018), which showcases the formulaic and impersonal components of online hate speech toward women. By emphasizing a woman's gender, trolls imply her flaws are tied to her sex classification – and therefore activate role congruity biases that all women are unfit for political office (Dovi, 2025; Krook and Restrepo, 2019; Krook, 2020; Montell, 2019; Sobieraj, 2020).[1]

I extend existing research on online VAWiP by examining gendered language outside hostile contexts. Some messages use forms of gendered language in explicitly supportive messaging, such as the hashtag #Vote4Women. Nevertheless, by heightening the salience of a politicians gender, even well-intentioned gendered language may inadvertently activate stereotypic expectations associated with gender roles (Bigler and Leaper, 2015). For example, calls to elect more women, as expressed in #Vote4Women, may subtly reinforce the perception of women candidates as interlopers in a man's domain (Puwar, 2004). In light of this, I analyze how non-hostile gendered language may still communicate gender norms online.

Drawing on the literature on gendered violence, I expect that backlash to perceived gender role violations will manifest in gendered forms (Bardall et al., 2020; Erikson et al., 2023). I focus on visibility as an existing proxy for role-incongruity (Håkansson, 2021; Rheault et al.,

[1]This is not the case for gendered hostility against men because the presence of men in politics is congruent with gender expectations (Clayton et al., 2024). Thus, gendered attacks against an individual man do not transfer to taint other men's perceived competence (Dovi, 2025).

2019) and introduce two novel indicators - legislator tone and the presence of women in the chamber - to further validate the relationship between gendered behavior and content.

I expect the following:

H1: Visible women receive more hostile and non-hostile-gendered content than less visible women and similarly visible men.

H2: Women who deviate more from gender roles receive more hostile and non-hostile-gendered content than conforming women.

H3: The online environment between men and women will differ more in gendered hostility than in general hostility.

Case Study: State Legislators

This study is the first to apply computational text analysis to social media data to examine online hostility directed at subnational politicians in the US. The comparatively low visibility of state legislators relative to members of Congress makes them a valuable test case for the relationship between political visibility and gender-based abuse. The identification of systematic gendered patterns in this setting strengthens the claim that such abuse is widespread and structurally embedded. As women ascend to higher office, they typically become more visible and less role congruent, suggesting that abuse may intensify with political power.

I scraped all the public tweets [2] toward 3,399 lower-house state representatives that were available in August 2022 and originated between October 2015 and July 2018. This time frame covers the campaign period and the full legislative term. The Twitter handles for the representatives were collected by Butler et al. (2023). Butler et al. (2023) found that 71% of women state representatives in 2017 had a Twitter account compared to 60% of men. Democratic representatives in 2017 were also slightly more likely to have a Twitter account than Republican representatives. The resulting corpus has over three million tweets. Butler et al. analyzed the tweets these state representatives posted during the timeframe, allowing me to control for the legislator's behavior on Twitter.

Measuring Hostile and Gendered Content

I utilized hand coding, zero-shot classification, and BERT feature representation to label the Twitter corpus for hostile and gendered content. I first classify the mentions with the

[2] Mentions are public tweets which tag the legislator by including their handle.

'gpt-4' model on the OpenAI API. [3] A team of expert coders given identical instructions to GPT-4 validate the zero-shot classification labels. [4] Next, I used the training corpus to finetune a pre-trained BERT model. [5]Validation of these labels indicated high accuracy and convergence.[6] A full description of this process is available in the Supplementary Materials.

Briefly, I consider a mention "hostile" if it contains impolite, rude, derogatory comments, vulgarities, threats, hate speech, dismissive tones, sexual harassment, racism, ad hominem attacks, or calls for a legislator's resignation. A mention is "gendered" if it references a legislator's physical appearance, relationship status, parental role, competence based on gender, utilizes gendered pejoratives, comments on stereotyped gender traits, or directly mentions the legislator's gender. Table 1 shows a randomly chosen tweet from each combination of labels. These confirm that GPT-4 accurately picks up on hostile and gendered content.

Table 1: Zero-shot Labels Show High Face Validity

Tweet	**Gendered?**	**Hostile?**
"Representative @legislator calls for an in-depth review of our transportation funding. #msleg"	0	0
"I proudly donated to and @legislator yesterday! Why? Because I know these phenomenal women will get the job done! #BlackWomenLead #ElectBlackWomen"	1	0
"@legislator You had a chance to help make the plan better and you blew up negotiations. Was it so you could complain for political reasons?"	0	1
" hey hey ho ho this bitch @legislator needs to go!"	1	1

[3]Before labeling the data, I removed all handles except the target legislator's in the text and replaced the legislator's handle with @legislator. The OpenAI annotation was conducted on August 15, 2023, with a model temperature of 0 to facilitate deterministic output. The expert coders annotated between August 6 and 9, 2023.

[4]The ICR between the zero-shot and expert classification for the hostile label is 0.89, and the F1 score is 0.75. The ICR for the gendered label is 0.93, and the F1 score is 0.74. These scores indicate that the zero-shot model has a relatively good balance between precision and recall.

[5]The HateBERT model was pre-trained on banned Reddit posts, making it ideal for detecting hostile content in Twitter mentions (Caselli et al., 2021). A pre-trained BERT model for gendered content does not exist. Instead, I use BERTweet to label the corpus for gendered content (Nguyen et al., 2020).

[6]A team of expert coders validated the BERT results. Both models are highly accurate, with an F1 score of 0.77 for hostility and 0.81 for gendered content.

Do Women Receive More Online Hostility? Yes.

I test my hypotheses using three dependent variables: the percentage of received mentions that are hostile (*% Hostile*), the percentage of hostile mentions that are gendered (*% Hostile Gendered*), and the percentage of mentions that are gendered but not hostile(*% Gendered, Not-Hostile*). My independent variable of interest is the interaction between *Woman* and *Visibility*.

I follow the literature and measure a legislator's *Visibility* by the number of mentions she receives (Cha et al., 2010; Gorrell et al., 2020; Rheault et al., 2019; Theocharis et al., 2020; Ward and McLoughlin, 2020). The number of mentions a legislator receives in this time frame correlates strongly with her legislative leadership position, how often the news references her, and how many times she is searched for on Google (see Supplementary Materials, part C). I control for the legislator's partisanship, ideology, and behavior on Twitter. I model the relationship with linear regression and include state-fixed effects with robust errors. [7]

Women Receive Disproportionate Gender-Based Content

Figure 1 illustrates the importance of considering gendered content when evaluating the online environments of men and women politicians. The leftmost plot depicts the positive relationship between *Visibility* and *% Hostile*. Had I followed much of the computational text analysis literature in examining solely the frequency of online abuse, I would conclude that trolls do not disproportionately target women in politics. The central plot in Figure 1 belies the error in this conclusion.

Going from low to high *Visibility* increases a woman's percentage of hostile mentions with gendered content by $\sim$ 22 percentage points and a man's by $\sim$10. At the highest levels of *Visibility*, a full quarter of the hostility women receive relates to their gender, compared to 15% for men. Visible women also receive more non-hostile gendered content than less visible women and similarly visible men. The much smaller y-axis in the rightmost panel suggests that gendered content is more likely to appear in hostile than non-hostile messaging. These results support my hypothesis that visible women receive more hostile and non-hostile gendered content than less visible women and similarly visible men **(H1)**.

These findings illuminate the unique challenges women in politics endure because of gender role expectations. Gender-based hostility is a form of semiotic violence that challenges the right of women to participate in politics (Krook and Restrepo, 2019). As such, the damaging effects of gendered-hostility extend far beyond the individual targets to undermine

[7] The proportion of women in each state's legislative chamber varies across states but remains constant over the study period. As a result, state-fixed effects capture this cross-sectional variation.

Figure 1: Visibility Impacts Gender-Based Content for Women

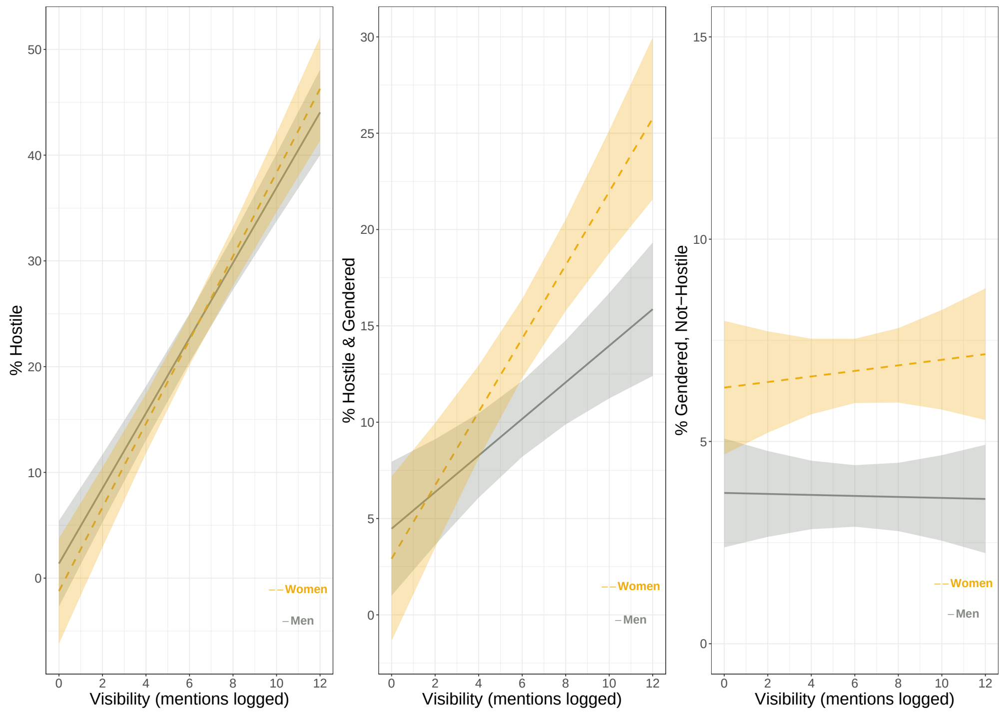

Note:The x-axis shows *Visibility*, measured as the logged number of mentions a legislator received (e.g., a score of x equals 10^{x} mentions). From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile. I obtained the predicted values from Table A.1 in the Supplementary Materials.

democracy itself (Frankel, 2020; Holm et al., 2024; Sobieraj, 2020).

Role-Incongruity Correlates with Gendered Content

I theorized that if a backlash against visible women took the form of gendered content, this would further support the conceptualization of public visibility as violating gender role expectations for women politicians (Rheault et al., 2019). The results support this theory by indicating that online hostility toward visible women state representatives takes the gendered form of disparate content rather than volume. To lend further support to the centrality of gender role congruity in online VAWiP, I demonstrate that backlash toward two other forms of gender role-incongruity also takes the form of gendered-hostility rather than an increased frequency of generic hostility.

The first alternate measure of gender role congruity is the positive or negative tone of the legislator's tweets during this time frame. [8]Many citizens still expect women in

[8]Research assistants for Butler et al. hand-labeled ten thousand tweets from legislators as being negative, neutral, or positive in sentiment. They calculate tone by multiplying a legislator's negative tweets by negative one, neutral tweets by zero, and positive tweets by one, then summing the result.

politics to express a more positive viewpoint than men (Bauer, 2020a). Consistent with this expectation, women legislators in the dataset tweeted more positively than their partisan counterparts (Butler et al., 2023), and women members of Congress posted more images of themselves as happy than neutral or upset (Boussalis et al., 2022).

Figure 2: Women's Tone Corresponds with the Gendered Content They Receive

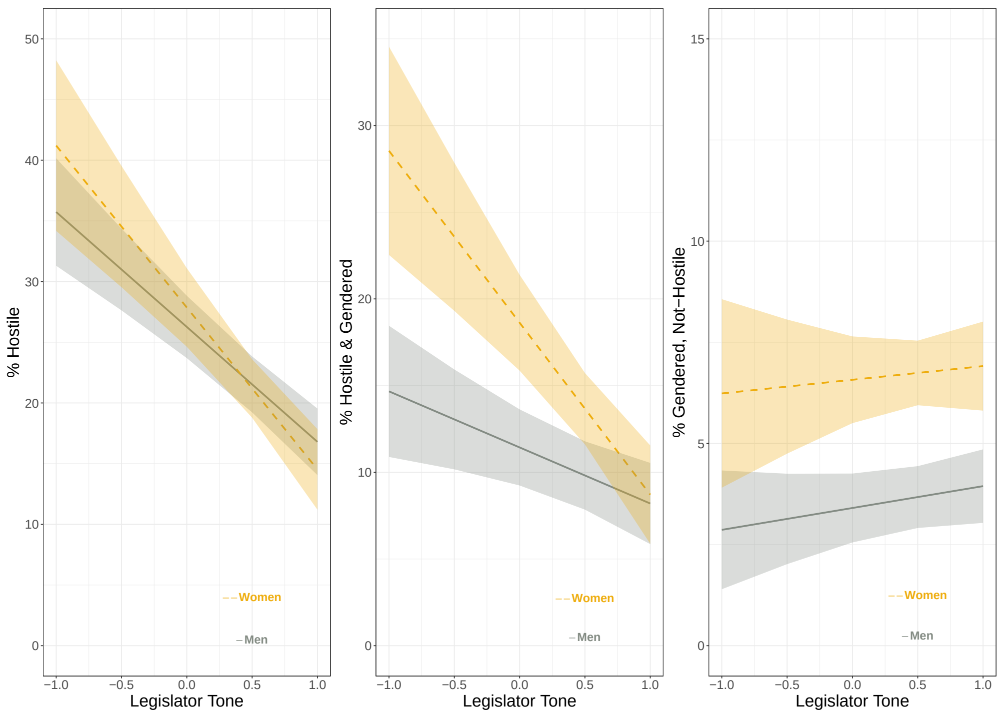

Note:The x-axis shows Legislator Tone, ranging from −1 (entirely negative) to 1 (entirely positive), with 0 indicating neutrality. From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile. I obtained the predicted values from Table A.2 in the Supplementary Materials.

Women tweeting in a non-positive tone are thus using a communication style that is incongruent with traditional gender roles. As shown in Figure 2, the corresponding backlash is evident in the content rather than the volume of their hostile mentions. Legislators who express positive emotions tend to receive less hostility. However, women who tweet in a negative – or even neutral– tone receive *double* the percentage of hostile-gendered content that negative men do. Again, gendered content is expressed more frequently in hostile than non-hostile terms.

The second alternate measure of gender role congruity is the growth rate of women's representation in the state legislature. [9] The proportion of women in office influences

[9]I calculate the growth rate of each state legislature as the difference in the number of women between 2017 and 2007 divided by the number of women in 2007. To capture cultural shifts, I use a growth rate rather than a static proportion of women officeholders.

whether observers perceive women politicians as interlopers or insiders (Puwar, 2004). In legislatures with low proportions and limited growth rates of women officeholders, women politicians may stand out as more visible deviations from the male norm. Alternatively, a sudden increase in women's representation may prompt a backlash, as suggested by research on violence against women in politics (Herrick et al., 2022; Sanín, 2022).

Once again, women who more strongly violate gender role expectations experience backlash in the form of gendered content rather than overall volume. In states where the proportion of women is declining, women legislators receive approximately twice as much gender-based vitriol as men legislators. At the same time, the data reveals a more hopeful trend. As women's representation increases, gender remains a salient attribute - but it becomes less frequently associated with hostility and more frequently linked to non-hostile messaging. This pattern persists even in states where the proportion of women representatives remains below the national median. The model in Figure 3 suggests that as the proportion of women officeholders increases, their presence becomes normalized, triggering less backlash for gender role violations.

Figure 3: Increasing Women's Presence Decreases Gendered Hostility

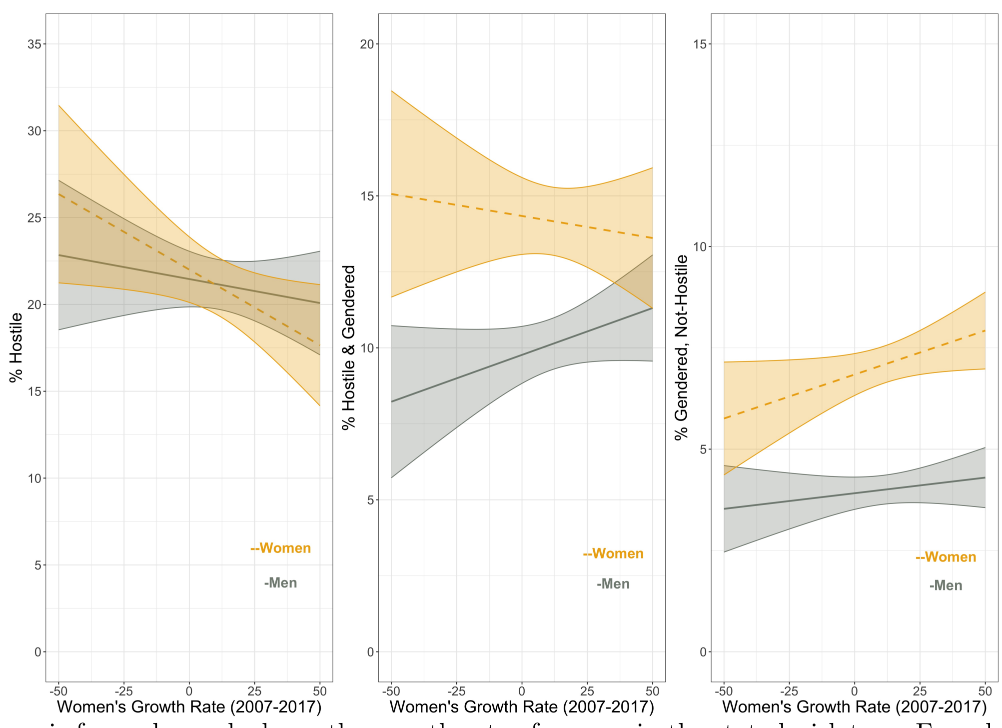

Note:The x-axis for each graph shows the growth rate of women in the state legislature. From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile. I obtained the predicted values from Table A.3 in the Supplementary Information.

The results support my argument that the divergent results between self-reports and observational data stem from a mismatch in what each method captures. Incorporating gendered language into observational analyses yields findings that align with women's self-reported experiences (Bjarnegård et al., 2022; Erikson et al., 2023; Kosiara-Pedersen, 2023). Backlash to women politicians who behave in gender role-incongruent ways often takes the form of distinctly gendered content. Patterns across visibility and other congruity measures reinforce the framing of public visibility as gender roles incongruent for women (Rheault et al., 2019). Across the three measures of gender role congruity, general hostility does not vary systematically by legislator gender.[10] In contrast, gendered hostility is closely tied to behavior and disproportionately directed at women. As women's behavior becomes more incongruous with traditional gender norms, they face significantly steeper increases in gendered hostility than men exhibiting similar behaviors.

These findings support the expectation that women who defy gendered norms are more likely to be targeted with gender-specific abuse **(H2)** and affirm that online environments diverge more starkly for men and women in the volume of gendered hostility than general hostility **(H3)**. While gendered language does appear in non-hostile contexts, it is nearly three times more likely to occur alongside hostility, and non-hostile gendered language shows little systematic relationship to legislator behavior. These patterns underscore the need for computational text analysis research to move beyond volume-based measures and account for both gender roles and gendered content to capture the full scope of online political violence against women.

Thus far in the paper I have shown that women politicians *do* receive disproportionate abuse on social media platforms. Although men and women experience similar levels of general hostility, women experience nearly twice as much gender-based hostility. Scholars theorize that perpetrators who criticize on the basis of gender wish to punish women politicians' gender role violations and reduce their political participation (Chadha et al., 2020; Frankel, 2020; Holm et al., 2024; Kjøller and Pedersen, 2025; Krook and Restrepo, 2019; Krook, 2020; Pedersen et al., 2024; Ramachandran et al., 2024; Womens Media Center, 2017; Yan and Bernhard, 2024). Are perpetrators of gendered hostility successful in causing women incumbents to exit politics?

[10] While the overall volume of abuse directed at men and women may be comparable, its impact on women's political engagement may be more severe given their continued global underrepresentation in formal office, political influence, leadership positions, and decision-making roles (Bochel and Bochel, 2008; National Press Foundation, 2024; **?**).

The Gendered Impact of Online Hostility

The literature on gendered violence, violence against women in politics, and online hostility would lead us to believe that women politicians do exit politics in response to online abuse. Research within these fields consistently shows that abuse disproportionately impacts women politicians' behaviors. Surveys and interviews find that women are more likely than men to reduce their public presence following online attacks (Chadha et al., 2020; Holm et al., 2024; Håkansson and Lajevardi, 2024; Pedersen et al., 2024). Even those who are not currently targeted often self-censor, either because of past experiences of after witnessing attacks toward a peer (Holm et al., 2024; Jankowicz, 2017; Ramachandran et al., 2024; Sobieraj, 2020). Among state legislators, targeted women report greater reluctance to use social media, make traditional media appearances, give interviews, and hold public events (Ramachandran et al., 2024). Most compellingly, larger shares of women than men report considering exiting politics after experiencing online attacks (Erikson et al., 2023; Herrick and Franklin, 2019; Håkansson and Lajevardi, 2024; Ramachandran et al., 2024).

The broader literature on women's experiences in office suggest an opposite expectation of women being resilient to gendered costs. Compared to men, women politicians report greater difficulty fundraising (Barber et al., 2016; Jenkins, 2007; Thomsen and Swers, 2017), receive less valuable committee seats than men (McGrath et al., 2025), and are more likely to be targeted by challengers (Lazarus and Steigerwalt, 2018). At the same time, women raise just as much money as men (Hayes and Lawless, 2016), are more effective at passing legislation (Eatough and Preece, 2025), and demonstrate comparable rates of reelection as men (Butcher, 2021, 2022, 2023; Butcher and Haynes, 2024; Thomsen and Swers, 2017). In addition, women are more responsive to constituents (Bauer and Cargile, 2023; Butler et al., 2022; Dickson, 2025; Thomsen and Sanders, 2020) and cultivate more comprehensive legislative agendas (Butler et al., 2022; Lazarus and Steigerwalt, 2018; Schmitt and Brant, 2019). Collectively, the research indicates that although women experience different hurdles than men politicians, they find ways to overcome them and continue serving in elective office.

Gender-based selection theory helps explains why women exhibit higher performance than men despite confronting greater challenges. If women either perceive or experience sex discrimination in the electoral process, then only the most exceptionally capable and ambitious women will emerge as candidates (Anzia and Berry, 2011a). As a result, the pattern of "when women run, women win" masks a reality where women must often exceed their male counterparts in demonstrable quality to be perceived as equally competitive (Bauer, 2020b; Fulton, 2012b; Lazarus and Steigerwalt, 2018; Pearson and McGhee, 2013).

Applied to the issue of gendered online hostility, gender-based selection theory suggests that the women who pursue political office will exhibit higher tolerance for abuse. Women who pursue political careers correctly anticipate gendered abuse and discrimination (Fox and Lawless, 2024; Kjøller and Pedersen, 2025; Shames, 2017), meaning these costs were baked into their decision to enter politics. Their presence in office suggests they decided prior to campaigning that the gendered costs of online hostility would be worth bearing (Black, 1972).

Indeed, women politicians have developed preventative strategies to mitigate or manage online abuse. There are three courses of action encouraged by candidate training groups. The first strategy is focused on prevention, where women strengthen their digital security using a suite of tools that prevent hacks and doxxing (Jankowicz, 2022; SheShouldRun, 2022). The second strategy is offloading social media accounts to staff (Schriock and Reynolds, 2021; SheShouldRun, 2022; Tenove et al., 2023). Finally, women are proactive in building communities with other women who can empathize, help report offensive comments to platforms, and advise on legal action (Jankowicz, 2022; Schriock and Reynolds, 2021; Sobieraj, 2020).

Women demonstrate additional coping mechanisms beyond these formal actions. Some women normalize the abuse and report being unbothered by the attacks, even as they describe their tactics for avoiding future hostility (Astor, 2018; Sobieraj, 2020). Others fight fire with fire by "outing" their attackers or using insulting language as a means of self-defense (Boukemia et al.; Phillips, 2017; Sobieraj, 2020). Many women reduce attackers' ability to discredit them by extensively citing official sources in their online communication (Jankowicz, 2022; Schriock and Reynolds, 2021; Sobieraj, 2020). Even the decision to avoid speaking on volatile issues - while it has harmful implications for democracy - is evidence of adaptive strategies to manage abuse (Celuch et al., 2024; Sobieraj, 2020).

Competing Expectations of The Toll of Trolls

Thus, we infer two competing theories within the literature for how women will respond to gendered hostility. The online hostility and violence against women in politics literature argues that the high costs of online abuse will reduce the utility of serving in office and thus spur women's disproportionate exit. This is the dominant expectation in the discipline, with over a dozen publications voicing it (Carson et al., 2024; CCDH, 2024; Erikson et al., 2023; Harmer and Southern, 2023; Herrick and Franklin, 2019; Herrick and Thomas, 2022; Håkansson, 2024; Håkansson and Lajevardi, 2024; Jankowicz, 2017, 2022; Kosiara-Pedersen, 2023; Krook, 2017, 2020; Ramachandran et al., 2024; Sobieraj, 2018, 2020; Thomas et al., 2019; Wagner, 2022). The gender-based selection literature has not directly engaged with the question of how women politicians will respond to online hostility, however, it suggests the opposite. Because women in the candidate pool are aware of the gendered hostility women politicians experience, those who choose to pursue office will be exceptionally resilient to this abuse - perhaps more so than men - and develop adaptive strategies to overcome it.

H4a: Exposure to hostility will disproportionately increase women's likelihood of choosing to exit politics relative to men.

H4b: Women will demonstrate approximate or greater likelihood of campaigning despite exposure to hostility, relative to men.

Identifying Gendered Hostility

In determining the influence of exposure on ambition, I again choose to use state legislators as my case study for two reasons. First, state legislators have fewer resources to screen abuse, which increases the likelihood of violence making direct contact with the targeted politicians (Butler et al., 2023; Wood, 2021).[11] Second, most national politicians begin their career in state legislatures (Holman, 2017; Manning, 2024; McCrain and O'Connell, 2023) State legislators have already demonstrated political ambition by entering office, yet many are still early in their careers. As a result, state legislators are an ideal test case for examining how online hostility influences a politician's decision to exit.

I extend the Twitter data collected above by additionally scraping all the public tweets [12] toward 1,247 state representatives across four legislative election cycles, from 2015 to 2023. [13] The Twitter handles for the representatives were collected by Butler et al. (2023).[14] At the time the handles were identified, women and Democratic representatives were more likely to have a Twitter account than men or Republican representatives. The resulting corpus has approximately 14 million tweets directed at 370 women and 913 men politicians over four election cycles.[15] Approximately half of the politicians are Republican, although women are

[11]National politicians often hire staff whose job is to curate a daily "social media scrap" of relevant activity for their employer. As a result, well-resourced politicians are able to avoid direct engagement with abusive content online.

[12]Mentions are public tweets which tag the legislator by including their handle.

[13]These tweets were collected in February 2023 and originated after June 2015.

[14]Note: The original intention was to scrape Twitter mentions for the 3,399 state representatives identified by Butler et al. (2023) who were active on Twitter in 2017. As this process occurred, Twitter discontinued its academic API, cutting off the data collection at 1300 state representatives. Because the politicians were ordered alphabetically by first name, the dataset only contains politicians whose first names begin with the letter 'a' through 'i'. Balance tests (available in Appendix B) between the scraped and un-scraped politicians do not reveal proportional differences across gender, party, or race.

[15]Most state legislatures have two-year terms.

nearly twice as likely to be members of the Democratic party. Most of the politicians (79%) are white.

As before, I use hand-coding, zero-shot classification, and fine-tuning a pre-trained BERT model to label the text content of each tweet. This process, label instructions, and comparison between label scores is available in full in the Supplementary Materials. In addition to identifying hostile and gendered content, I additionally label for supportive text. A mention is "supportive" if it explicitly supports the legislator, encourages others to support the legislator, or praises the legislator. Table 2 shows a randomly chosen tweet from each combination of labels. These confirm that the LLM accurately detects hostile, gendered, and supportive content. [16] I applied design-based supervised learning to correct errors in the BERT predictions to match the hand-coded labels (Egami et al., 2024).

Table 2: Labels show high face validity

Example Tweet	**Hostile**	**Gendered**	**Supportive**
"`@legislator` Really have to be the stupidest person on Twitter don't you?"	1	0	0
"`@legislator` You are in way over your head little girl!!"	1	1	0
"congressman `@legislator` gives his take on HB112"	0	1	0
"Women like `@legislator` are authentically themselves in representing the American people"	0	1	1
"`@legislator` Thank you for all that you're doing!!"	0	0	1

Identifying Movement Within or Out of Office

To establish whether online hostility affects political ambition, I collected the career paths of 1,247 state representatives campaign decisions and outcomes from 2017 through 2023 using CAWP and Ballotpedia. At the end of each politician's term, I document whether they choose to campaign for reelection, a different office, or not at all. If they choose to campaign, I track which campaigns the politician entered, the state of the primary and general race, and the election result. I also account for whether the politician received an appointed position or was term-limited out of their office. If a politician reaches the end of their term and does not campaign nor assume an appointed position, she is considered

[16] The F1 scores between the BERT predictions and the hand-coded tweets are 0.88 for gendered, 0.78 for hostile, and 0.83 for supportive.

to exit. I continue to collect campaign decisions for politicians who have exited to account for political rebirths. The resulting dataset has 4,864 observations of politician-election year pairs. The demographic data of each politician—including race, party, gender, and district—were collected by Butler, Kousser, and Oklobdzija (2023).

Candidates must file paperwork by a specific deadline to appear on the ballot in their state. Deadlines vary by state and year. I collected filing deadlines from Ballotpedia and National Conference of State Legislatures. Politicians may determine which races to enter long before the filing deadline, however, the filing deadline represents a hard cut-off for the decision to campaign or exit. By adding the filing deadlines to the career path dataset, I can approximate a decision timeline for each politician – election year pair. I use the filing deadlines to form summary measures of a politician's online environment in the six months leading up to the decision cut-off. [17]

As all the individuals held a state representative office in 2017, I filtered the dataset to begin in 2018. For legislatures with a two-year term, the data thus encompasses three election cycles. For individuals who exit and do not campaign again, only the first cycle they sat out is labeled as an exit. Subsequent inactive years are not considered in the model. Similarly, politician - election year pairs where the individual is term-limited out of office are not included in the dataset.

Table 3 shows the breakdown of politician campaign choices by gender across all politician-election year pairs. While the most common choice is to campaign for reelection to the same office, more than 20% of the decisions are to exit. Movements toward a different office are comparatively infrequent.

Table 3: Politicians prefer to campaign for reelection

	Reelection	**Exit**	**Different Office**
Women	756 (68.2%)	207 (20.8%)	108 (10.91%)
Men	1581 (67.9%)	562 (24.2%)	184 (7.9%)

Are Women Exiting Office? No.

The literature yields competing theories on whether women politicians will respond to online hostility by exiting office. To test these alternate explanations, I use a multinomial logit to estimate the correlation between online hostility and the decision to exit politics.

[17] In the appendix, I discuss the decision to use six months prior to a filing deadline. I compare it to two alternate timeframes of two years and political lifetime.

I approximate politician quality using five indicators: legislator effectiveness,[18] tenure in office,[19] lagged fundraising and donors,[20] and visibility.[21] Standard errors are bootstrapped and clustered at the individual level.

Women Are Less Likely than Men to Exit Office in Response to Hostility

The baseline category in the multinomial logit is the decision to not campaign. Model 1 shows the likelihood of an individual campaigning for a different office than they currently hold versus exit. Model 2 indicates the likelihood of campaigning for reelection versus exit. Thus, a positive coefficient indicates that the predictor is associated with a higher likelihood of being in the outcome category (e.g., campaign for reelection) rather than exit politics. A negative coefficient suggests that increasing the value of the predictor lowers the likelihood of a politician being in the outcome category rather than exiting.

Contrary to the widely published theory that online hostility is pushing women out of office **(H4a)**, the results in Table 4 indicate that women are significantly *less* likely than men to exit office in response to gendered and general hostility **(H4b)**. Instead, women respond to increased exposure by campaigning for a different office than they currently hold. This is strong support for the theory of gender-based selection, where women enter office expecting gendered discrimination (Fox and Lawless, 2024; Kjøller and Pedersen, 2025; Shames, 2017) and develop adaptive strategies to manage and mitigate this abuse (Boukemia et al.; Jankowicz, 2022; Phillips, 2017; Ramachandran et al., 2024; Schriock and Reynolds, 2021; SheShouldRun, 2022; Sobieraj, 2020).

I strengthen the finding that women are less likely than men to exit with a series of robustness checks. First, I test whether the result in Table 4 holds when modeled as a binomial with a binary dependent variable of exit or campaign rather than a multinomial regression. The resulting table confirms the primary effect in Table 4 and is displayed in the Supplementary Materials.

Second, I examine the effect of women moving toward a different office. An individual's

[18] The State Legislative Effectiveness Score captures the number of bills a legislator sponsored, the extent to which those bills advanced through the legislative process, and their substantive importance (Bucchianeri et al., 2025). I use the 2017 SLES score because all the legislators in the dataset held office during this year.

[19] I collected the first-year an individual was elected to any office using Ballotpedia, Wikipedia, personal websites, and newspaper articles.

[20] I incorporate fundraising information from Follow The Money https://www.followthemoney.org/our-data/about-our-data. This lagged variable provides the number of unique donors and the total sum raised by a candidate in their previous campaign. Where possible, I supplement fundraising data with DIME (Bonica, 2024).

[21] I use the number of mentions a politician received on Twitter as a measure of their name recognition and visibility to the public (Cha et al., 2010; Gorrell et al., 2020; Jarman; Rheault et al., 2019; Theocharis et al., 2020; Ward and McLoughlin, 2020).

Table 4: Multinomial Logit with Cluster-Robust SEs (Bootstrapped)

	Campaign for a Different Office	Campaign for Reelection
Constant	-1.760^{***}	-0.980^{***}
	(0.290)	(0.100)
Log Mentions	0.170	0.050
	(0.200)	(0.060)
Fundraising (Scaled)	-0.040	0.030
	(0.050)	(0.230)
Democrat	-0.220^{***}	0.040
	(0.040)	(0.180)
Legislative Effectiveness	0.050	0.100
	(0.130)	(0.640)
Non-white	0.320^{***}	0.270
	(0.060)	(0.450)
Count Donors (Scaled)	0.460^{**}	0.440^{*}
	(0.210)	(0.260)
Year First Elected (Scaled)	0.440^{***}	0.250
	(0.140)	(0.190)
In Office	0.880^{***}	2.830^{**}
	(0.100)	(1.180)
Woman	0.070	-0.220
	(0.250)	(1.300)
General Hostility	-0.670^{***}	-1.160
	(0.070)	(1.010)
Gendered Hostility	-0.600^{***}	-1.010
	(0.180)	(1.000)
General Hostility × Woman	1.310^{***}	2.160
	(0.130)	(3.600)
Gendered Hostility × Woman	2.590^{***}	-0.440
	(0.100)	(3.480)
N	2619	
LogLik	-1650.49	
AIC	3357.00	
BIC	3521.40	
McFadden R^2	0.12	
Cox–Snell R^2	0.15	
Nagelkerke R^2	0.20	

Note: Cluster-robust (bootstrapped) standard errors shown in parentheses. Baseline category (Exit) omitted. *p<0.1; **p<0.05; ***p<0.01.

place in the pipeline strongly correlates with how visible they are to the public, which in turn influences how much online hostility they receive (Bjørgo, 2022; Håkansson, 2021; Jarman; Rheault et al., 2019; Theocharis et al., 2020). This suggests that a politician who wished to reduce their exposure to online abuse could move to a lower profile office. There is some support for this theory in the literature. In a survey of politicians, women indicated twice as much willingness as men to decrease their pay in exchange for reduced harassment (Kjøller and Pedersen, 2025).

Additionally, there is reason to expect that some women respond to gendered hostility with increased ambition. Sexually-based violence and gendered policy threats can be politically mobilizing for women because staying out of politics becomes perceived as too politically costly (Agerberg and Kreft, 2023; Clayton et al., 2023; Dittmar, 2020). Mobilization in response to marginalization is most pronounced among individuals with experiences of discrimination, suggesting that women politicians with higher exposure to gendered hostility will be more likely to campaign for more influential offices (Kjøller and Pedersen, 2025; Pedersen et al., 2025).

I repeat the multinomial model, this time with separate categories for movement toward a higher or lower office. The results indicate that women are significantly more likely than men to campaign for both higher and lower offices rather than exit. I rank offices as higher or lower by considering whether the seat has more power within that same state's context. For instance, national positions and statewide executive seats are considered higher than a state legislative seat, whereas local or county level roles are considered lower. A full documentation of which offices appear in the dataset, at what frequency, and how they are coded, as well as the regression results, are available in the Supplementary Materials.

Third, I check whether men may be receiving hostility of a worse kind. To test this possibility, I use the open-source model Detoxify, which calculates the probability that a tweet includes the following: toxicity, severe toxicity, obscene, threat, insult, identity attack, sexually explicit language. [22] [23] [24] I sum these scores to create a predicted abuse score for each tweet. Next, I compare mean abuse scores per tweet for men and women using a Welch two-sample t-test. Figure 4 is a dot-and-whisker plot illustrating the mean abuse score per message by target gender across four categories of tweet labels: Gendered Hostility, General

[22] https://github.com/unitaryai/detoxify?tab=readme-ov-file

[23] The training data for this model was annotated in 2018 on the Figure Eight crowd rating platform, with each tweet being annotated by up to ten workers. The 'unbiased' model corrects for a tendency of text-classification models to assume all mentions of identities are intended negatively.

[24] Toxicity is defined as "a rude, disrespectful, or unreasonable comment that is somewhat likely to make you leave a discussion or give up on sharing your perspective". Severe toxicity follows the same schema, with the difference of *very* likely as opposed to *somewhat*. The category of identity attack specifically considers gender (male/female), sexuality, religion, race, and psychiatric or mental illness.

Hostility, Gendered Other,[25] and General Other.[26] The dots show the mean and the whiskers show 95% confidence intervals.

Figure 4: Hostility toward women is more abusive

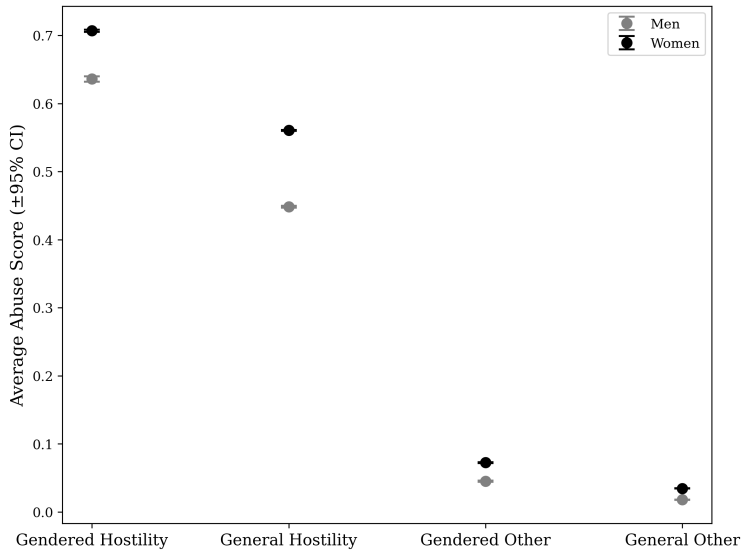

Note: Women are consistently higher than men in both hostility types. General Hostility: 0.561 vs. 0.448 (Δ=0.112; Welch t-test $p < 0.001$; Cohen's $d = 0.13$), Gendered Hostility: 0.707 vs. 0.636 (Δ=0.071; $p < 0.001$; $d = 0.07$). Women are also slightly higher in the "Other" categories (Gendered Other: 0.072 vs. 0.045, $p < 0.001$; General Other: 0.034 vs. 0.018, $p < 0.001$).

The results in Figure 4 show several important findings. Gendered hostility is more abusive than general hostility, particularly when it is directed toward women. Even when hostility is not gender-based, it is still more abusive when targeting a woman than a man. This finding adds weight to the result in Table 4. Despite receiving more severe abuse than men, women's ambition is more resilient.

Lastly, I check whether women receive greater proportions of supportive messages online, perhaps mitigating the effect of hostility on ambition. To test this, I computed the proportion of messages men and women received with gendered hostility, general hostility, and supportive content (both gendered and general). I ran a two-proportion z-test with pooled variance to compare the online environments of men and women. Figure 5 displays the results. On average, women receive nearly twice as much gendered hostility, one-and-a-half times the general hostility, and only one-third the support that men do. Paired with Table 4, these results further support applying the theory of gender-based selection by strengthening our understanding of women's exceptionalism.

[25] Tweets which were labeled as gendered and supportive, or gendered and non-hostile.

[26] Tweets which were labeled as non-gendered and supportive, or non-gendered and non-hostile.

Figure 5: Women receive more hostility and less support

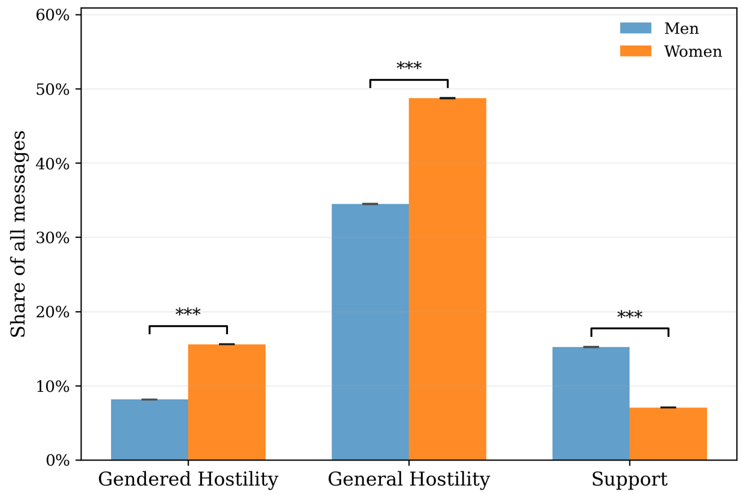

Note: These results come from a two-proportion z-test with pooled variance. Differences are statistically significant at the $p < 0.001$ level.

Discussion and Conclusion

This article addresses two of the most pressing questions in modern politics - how politicians experience online hostility, and what consequences exposure to this hostility has on women in office. My research sheds new light on two important pieces of conventional wisdom in the literature. First, I demonstrate that the predominant narrative that women and men politicians experience similar online environments does not hold after accounting for the gendered content of abuse. Second, contrary to the extensive theorization that online hostility is leading women politicians to exit office, I use observational data to showcase that women's political ambition is resilient to gendered hostility.

This study makes two central contributions. First, it explains why some observational research has underestimated the prevalence of gendered hostility online. When computational text approaches to measuring online hostility fail to examine the gendered content of abuse, they risk missing the unique form of backlash directed at women who violate gender norms (Bardall et al., 2020). Incorporating gendered language into observational analyses yields findings that align with women's self-reported experiences (Bjarnegård et al., 2022; Erikson et al., 2023; Kosiara-Pedersen, 2023). I show that visible women receive nearly twice as much gender-specific abuse as men despite experiencing similar overall volumes of hostility. Importantly, I show that backlash against role-incongruent women is not only more frequent but also more gendered in content.

Second, the study pairs observational data of politicians' online environments with their career choices across four campaign cycles to demonstrate that online hostility is not pushing women out of office. Although women politicians receive more severe and gender-based online hostility than men, they are less likely than men to exit office in response to online abuse. This finding suggests a gender-based selection, where women in the candidate pool correctly anticipate gendered abuse and discrimination (Fox and Lawless, 2024; Kjøller and Pedersen, 2025; Shames, 2017). As a result, the women who choose to pursue political careers are exceptionally resilient to this abuse and develop adaptive strategies to overcome it (Anzia and Berry, 2011b).

Women's resilience to gender-based hostility does not negate the importance of countering the bias which inspired it. If policymakers were to take these findings seriously, we should observe more effort to understand the causal mechanisms behind differential responses to hostility within and across gender groups. Future research should consider the comparative weight of personality and structural supports, such as staff to handle social media accounts or harassment training, on whether politicians exit office.

Future research should additionally consider the broader impacts of online hostility on ambition and democratic representation. What, if any, ripple effects on policy and representation result from restricting political candidacies to only the individuals most willing to tolerate sustained abuse? Politicians believe that the intention behind gendered hostility is to silence women and push them out of the political domain (Frankel, 2020; Jankowicz, 2022; Phillips, 2017; Sobieraj, 2020). The demonstrated resilience of the women in this study suggests that the silencing effect of abuse may be most powerful in winnowing the candidate pool by making political office less attractive to all but the most thick-skinned individuals.

References

Mattias Agerberg and Anne-Kathrin Kreft. Sexual violence, gendered protection and support for intervention. *Journal of Peace Research*, 60(5):853–867, September 2023. ISSN 0022-3433. doi: 10.1177/00223433221092960. URL https://doi.org/10.1177/00223433221092960. Publisher: SAGE Publications Ltd.

Sarah F. Anzia and Christopher R. Berry. The Jackie (and Jill) Robinson Effect: Why Do Congresswomen Outperform Congressmen? *American Journal of Political Science*, 55 (3):478–493, 2011a. ISSN 1540-5907. doi: 10.1111/j.1540-5907.2011.00512.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-5907.2011.00512.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-5907.2011.00512.x.

Sarah F. Anzia and Christopher R. Berry. The Jackie (and Jill) Robinson Effect: Why Do Congresswomen Outperform Congressmen? *American Journal of Political Science*, 55 (3):478–493, 2011b. ISSN 1540-5907. doi: 10.1111/j.1540-5907.2011.00512.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-5907.2011.00512.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-5907.2011.00512.x.

Maggie Astor. For Female Candidates, Harassment and Threats Come Every Day. *The New York Times*, August 2018. ISSN 0362-4331. URL https://www.nytimes.com/2018/08/24/us/politics/women-harassment-elections.html.

Michael Barber, Daniel M. Butler, and Jessica Preece. Gender inequalities in campaign finance. *Quarterly Journal of Political Science*, 11(2):219–48, 2016. URL https://brightspotcdn.byu.edu/f8/76/167206e44e93b9087a116f9e41b3/qjps-manuscript-campaign-finance.pdf. Publisher: Now Publishers, Inc.

Gabrielle Bardall, Elin Bjarnegård, and Jennifer M Piscopo. How is Political Violence Gendered? Disentangling Motives, Forms, and Impacts. *Political Studies*, 68(4):916–935, November 2020. ISSN 0032-3217. doi: 10.1177/0032321719881812. URL https://doi.org/10.1177/0032321719881812. Publisher: SAGE Publications Ltd.

Nichole M. Bauer. A Feminine Advantage? Delineating the Effects of Feminine Trait and Feminine Issue Messages on Evaluations of Female Candidates. *Politics & Gender*, 16(3): 660–680, September 2020a. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X19000084. URL https://www.cambridge.org/core/journals/politics-and-gender/article/feminine-advantage-delineating-the-effects-of-feminine-trait-and-feminine-issue-messages-on-evaluations-of-female-candidates/643BD6018FD064B110C1355B8B3A525B.

Nichole M. Bauer. Shifting Standards: How Voters Evaluate the Qualifications of Female and Male Candidates. *The Journal of Politics*, 82(1):1–12, January 2020b. ISSN 0022-3816. doi: 10.1086/705817. URL https://www.journals.uchicago.edu/doi/full/10.1086/705817. Publisher: The University of Chicago Press.

Nichole M. Bauer and Ivy A. M. Cargile. Women Get the Job Done: Differences in Constituent Communication from Female and Male Lawmakers. *Politics & Gender*, 19(4): 1110–1133, December 2023. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X23000259. URL https://www.cambridge.org/core/journals/politics-and-gender/article/women-get-the-job-done-differences-in-constituent-communication-from-female-and-male-lawmakers/826A29E07C6AB3C0692F07E9828D0B94.

Rebecca S. Bigler and Campbell Leaper. Gendered Language: Psychological Principles, Evolving Practices, and Inclusive Policies. *Policy Insights from the Behavioral and Brain Sciences*, 2(1):187–194, October 2015. ISSN 2372-7322. doi: 10.1177/2372732215600452. URL https://doi.org/10.1177/2372732215600452. Publisher: SAGE Publications.

Elin Bjarnegård. Making Gender Visible in Election Violence: Strategies for Data Collection. *Politics & Gender*, 14(4):690–695, December 2018. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X18000624. URL https://www.cambridge.org/core/journals/politics-and-gender/article/making-gender-visible-in-election-violence-strategies-for-data-collection/BE1D388964BB2CBA760993F5844D1D6E. Publisher: Cambridge University Press.

Elin Bjarnegård, Sandra Håkansson, and Pär Zetterberg. Gender and Violence against Political Candidates: Lessons from Sri Lanka. *Politics & Gender*, 18(1):33–61, March 2022. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X20000471. URL https://www.cambridge.org/core/journals/politics-and-gender/article/gender-and-violence-against-political-candidates-lessons-from-sri-lanka/988B1E4E8D480C82C3D6D3CBDEF05289.

Tore Bjørgo. Patterns and Consequences of Threats Against Politicians. Aarhus, Denmark, June 2022.

Gordon S. Black. A Theory of Political Ambition: Career Choices and the Role of Structural Incentives. *American Political Science Review*, 66(1):144–159, March 1972. ISSN 0003-0554, 1537-5943. doi: 10.2307/1959283. URL https://www.cambridge.org/core/journals/american-political-science-review/article/abs/theory-of-political-ambition-career-choices-and-the-role-of-structural-incentives/491E20642B50572F651F2F12C5D9AACC.

Catherine Bochel and Hugh Bochel. Women 'Leaders' in Local Government in the UK. *Parliamentary Affairs*, 61(3):426–441, July 2008. ISSN 0031-2290. doi: 10.1093/pa/gsn010. URL https://doi.org/10.1093/pa/gsn010.

Bonica. Database on Ideology, Money in Politics, and Elections, 2024.

Jana Boukemia, Marius Sältzer, and Sébastien Boyer. Misogyny, politics, and social media determinants of hostile engagement against women parliamentarians on Twitter. *Legislative Studies Quarterly*, n/a(n/a). ISSN 1939-9162. doi: 10.1111/lsq.12486. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/lsq.12486. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/lsq.12486.

Constantine Boussalis, Travis Coan, and Mirya R. Holman. Mad, sad, but mostly glad: how men and women in politics communicate using emotions in images. *SocArXiv Papers*, June 2022. URL https://osf.io/preprints/socarxiv/w4adu/.

Victoria L. Brescoll, Tyler G. Okimoto, and Andrea C. Vial. You've Come a Long Way…Maybe: How Moral Emotions Trigger Backlash Against Women Leaders. *Journal of Social Issues*, 74(1):144–164, March 2018. ISSN 1540-4560. doi: 10.1111/josi.12261. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/josi.12261. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/josi.12261.

Abram Brown. Facebook Can Be Toxic For Female Politicians, Company Documents Show, October 2021. URL https://www.forbes.com/sites/abrambrown/2021/10/27/facebook-can-be-toxic-for-female-politicians-company-documents-show/. Section: Social Media.

Peter Bucchianeri, Craig Volden, and Alan E. Wiseman. Legislative Effectiveness in the American States. *American Political Science Review*, 119(1):21–39, February 2025. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055424000042. URL https://www.cambridge.org/core/journals/american-political-science-review/article/legislative-effectiveness-in-the-american-states/E235B546F05F29CB5FB2AFC52EAD206C.

Jordan Butcher. Rethinking retirement: ambition in term-limited legislatures. *The Journal of Legislative Studies*, 27(2):285–299, April 2021. ISSN 1357-2334, 1743-9337. doi: 10.1080/13572334.2020.1812199. URL https://www.tandfonline.com/doi/full/10.1080/13572334.2020.1812199.

Jordan Butcher. Be Careful what You Count: Updating Legislative Turnover in the 50 States. *American Politics Research*, 50(4):503–510, July 2022. ISSN 1532-673X. doi: 10.1177/1532673X221082319. URL https://doi.org/10.1177/1532673X221082319. Publisher: SAGE Publications Inc.

Jordan Butcher. Parties and professionals, an exploration of turnover in U.S. state legislatures. *The Journal of Legislative Studies*, 0(0):1–24, June 2023. ISSN 1357-2334. doi: 10.1080/13572334.2023.2225255. URL https://doi.org/10.1080/13572334.2023.2225255. Publisher: Routledge _eprint: https://doi.org/10.1080/13572334.2023.2225255.

Jordan Butcher and Noah Haynes. Already Gone: Retaining Women in State Legislatures. *Legislative Studies Quarterly*, 49(1):33–73, 2024. ISSN 1939-9162. doi: 10.1111/lsq.12435. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/lsq.12435. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/lsq.12435.

Daniel M. Butler, Elin Naurin, and Patrik Öhberg. Constituents Ask Female Legislators to Do More. *The Journal of Politics*, 84(4):2278–2282, October 2022. ISSN 0022-3816. doi: 10.1086/719631. URL https://www.journals.uchicago.edu/doi/full/10.1086/719631. Publisher: The University of Chicago Press.

Daniel M. Butler, Thad Kousser, and Stan Oklobdzija. Do Male and Female Legislators Have Different Twitter Communication Styles? *State Politics & Policy Quarterly*, 23

(2):117–139, June 2023. ISSN 1532-4400, 1946-1607. doi: 10.1017/spq.2022.16. URL https://www.cambridge.org/core/journals/state-politics-and-policy-quarterly/article/do-male-and-female-legislators-have-different-twitter-communication-styles/7DA632E453AD9A924FE9DD42A315B112. Publisher: Cambridge University Press.

Andrea Carson, Mikolajczak , Gosia, Ruppanner , Leah, , and Emily Foley. From online trolls to 'Slut Shaming': understanding the role of incivility and gender abuse in local government. *Local Government Studies*, 50(2):427–450, March 2024. ISSN 0300-3930. doi: 10.1080/03003930.2023.2228237. URL https://doi.org/10.1080/03003930.2023.2228237. Publisher: Routledge _eprint: https://doi.org/10.1080/03003930.2023.2228237.

Tommaso Caselli, Valerio Basile, Jelena Mitrović, and Michael Granitzer. HateBERT: Retraining BERT for Abusive Language Detection in English. In Aida Mostafazadeh Davani, Douwe Kiela, Mathias Lambert, Bertie Vidgen, Vinodkumar Prabhakaran, and Zeerak Waseem, editors, *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH 2021)*, pages 17–25, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.woah-1.3. URL https://aclanthology.org/2021.woah-1.3.

CCDH. Abusing Women in Politics: How Instagram is failing women and public officials. Technical report, Center for Countering Digital Hate, August 2024.

Magdalena Celuch, Oksa , Reetta, Ellonen , Noora, , and Atte Oksanen. Self-censorship among online harassment targets: the role of support at work, harassment characteristics, and the target's public visibility. *Information, Communication & Society*, 27(11):2171–2190, August 2024. ISSN 1369-118X. doi: 10.1080/1369118X.2023.2289978. URL https://doi.org/10.1080/1369118X.2023.2289978. Publisher: Routledge _eprint: https://doi.org/10.1080/1369118X.2023.2289978.

Meeyoung Cha, Hamed Haddadi, Fabricio Benevenuto, and Krishna Gummadi. Measuring User Influence in Twitter: The Million Follower Fallacy. *Proceedings of the International AAAI Conference on Web and Social Media*, 4(1):10–17, May 2010. ISSN 2334-0770. doi: 10.1609/icwsm.v4i1.14033. URL https://ojs.aaai.org/index.php/ICWSM/article/view/14033. Number: 1.

Kalyani Chadha, Linda Steiner, Jessica Vitak, and Zahra Ashktorab. Women's Responses to Online Harassment. *International Journal of Communication*, 14(0):19, January 2020. ISSN 1932-8036. URL https://ijoc.org/index.php/ijoc/article/view/11683. Number: 0.

Amanda Clayton, Diana Z. O'Brien, and Jennifer M. Piscopo. Women Grab Back: Exclusion, Policy Threat, and Women's Political Ambition. *American Political Science Review*, 117 (4):1465–1485, November 2023. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055422001344. URL https://www.cambridge.org/core/journals/american-political-science-review/article/women-grab-back-exclusion-policy-threat-and-womens-political-ambition/0B32347849AA04403C336CF9060A4CCC?utm_source=SFMC&utm_medium=email&utm_content=Article&utm_campaign=New%20Cambridge%20Alert%20-%20

Issues%20Society&WT.mc_id=New%20Cambridge%20Alert%20-%20Issues%20Society. Publisher: Cambridge University Press.

Amanda Clayton, Diana Z. O'Brien, and Jennifer M. Piscopo. Founding Narratives and Men's Political Ambition: Experimental Evidence from US Civics Lessons. *British Journal of Political Science*, 54(1):129–151, January 2024. ISSN 0007-1234, 1469-2112. doi: 10.1017/S0007123423000340. URL https://www.cambridge.org/core/journals/british-journal-of-political-science/article/founding-narratives-and-mens-political-ambition-experimental-evidence-from-us-civics-lessons/F8E9E37218C081D35D62CFBEEBD59C3C.

Sofia Collignon and Wolfgang Rüdig. Harassment and Intimidation of Parliamentary Candidates in the United Kingdom. *The Political Quarterly*, 91(2):422–429, 2020. ISSN 1467-923X. doi: 10.1111/1467-923X.12855. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-923X.12855. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-923X.12855.

Zachary P. Dickson. The Gender Gap in Elite-Voter Responsiveness Online. *Perspectives on Politics*, 23(2):477–493, June 2025. ISSN 1537-5927, 1541-0986. doi: 10.1017/S153759272400104X. URL https://www.cambridge.org/core/journals/perspectives-on-politics/article/gender-gap-in-elitevoter-responsiveness-online/0D87B082135F0C261D848D32B8A7639D?utm_source=SFMC&utm_medium=email&utm_content=Article&utm_campaign=New%20Cambridge%20Alert%20-%20Issues%20Society&WT.mc_id=New%20Cambridge%20Alert%20-%20Issues%20Society.

Kelly Dittmar. Urgency and ambition: the influence of political environment and emotion in spurring US women's candidacies in 2018. *European Journal of Politics and Gender*, 3(1):143–160, February 2020. ISSN 2515-1088, 2515-1096. doi: 10.1332/251510819X15728693158427. URL https://bristoluniversitypressdigital.com/view/journals/ejpg/3/1/article-p143.xml. Publisher: Bristol University Press Section: European Journal of Politics and Gender.

Suzanne Dovi. On Political Misogyny. *American Political Science Review*, 119(2):1054–1067, May 2025. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055424000777. URL https://www.cambridge.org/core/journals/american-political-science-review/article/on-political-misogyny/8DB92ABE37FBAEC4018D35488DC3B449.

Alice H Eagly and Linda L Carli. The female leadership advantage: An evaluation of the evidence. *The Leadership Quarterly*, 14(6):807–834, December 2003. ISSN 1048-9843. doi: 10.1016/j.leaqua.2003.09.004. URL https://www.sciencedirect.com/science/article/pii/S1048984303000584.

Alice H. Eagly and Steven J. Karau. Role congruity theory of prejudice toward female leaders. *Psychological Review*, 109:573–598, 2002. ISSN 1939-1471. doi: 10.1037/0033-295X.109.3.573. Place: US Publisher: American Psychological Association.

Alice H. Eagly, Mona G. Makhijani, and Bruce G. Klonsky. Gender and the evaluation of leaders: A meta-analysis. *Psychological Bulletin*, 111:3–22, 1992. ISSN 1939-1455. doi: 10.1037/0033-2909.111.1.3. Place: US Publisher: American Psychological Association.

Mandi Eatough and Jessica R. Preece. Crediting invisible work: congress and the lawmaking productivity metric (lawprom). *American Political Science Review*, 119(2):566–584, 2025. URL https://www.cambridge.org/core/journals/american-political-science-review/article/crediting-invisible-work-congress-and-the-lawmaking-productivity-metric-lawprom/370725CAE441F969DC152C66484F0427. Publisher: Cambridge University Press.

Naoki Egami, Musashi Hinck, Brandon M. Stewart, and Hanying Wei. Using Imperfect Surrogates for Downstream Inference: Design-based Supervised Learning for Social Science Applications of Large Language Models, January 2024. URL http://arxiv.org/abs/2306.04746. arXiv:2306.04746 [cs, stat].

Josefina Erikson, Sandra Håkansson, and Cecilia Josefsson. Three Dimensions of Gendered Online Abuse: Analyzing Swedish MPs' Experiences of Social Media. *Perspectives on Politics*, 21(3):896–912, September 2023. ISSN 1537-5927, 1541-0986. doi: 10.1017/S1537592721002048. URL https://www.cambridge.org/core/journals/perspectives-on-politics/article/three-dimensions-of-gendered-online-abuse-analyzing-swedish-mps-experiences-of-social-media/F52E7389E355C1C78335B44B9E66811E.

Eleonora Esposito and Ruth Breeze. Gender and politics in a digitalised world: Investigating online hostility against UK female MPs. *Discourse & Society*, 33(3):303–323, May 2022. ISSN 0957-9265. doi: 10.1177/09579265221076608. URL https://doi.org/10.1177/09579265221076608. Publisher: SAGE Publications Ltd.

Diane Felmlee, Paulina Inara Rodis, and Amy Zhang. Sexist Slurs: Reinforcing Feminine Stereotypes Online. *Sex Roles*, 83(1):16–28, July 2020. ISSN 1573-2762. doi: 10.1007/s11199-019-01095-z. URL https://doi.org/10.1007/s11199-019-01095-z.

Olle Folke and Johanna Rickne. The Glass Ceiling in Politics: Formalization and Empirical Tests. *Comparative Political Studies*, 49(5):567–599, April 2016. ISSN 0010-4140. doi: 10.1177/0010414015621073. URL https://doi.org/10.1177/0010414015621073. Publisher: SAGE Publications Inc.

Richard L. Fox and Jennifer L. Lawless. The Invincible Gender Gap in Political Ambition. *PS: Political Science & Politics*, 57(2):231–237, April 2024. ISSN 1049-0965, 1537-5935. doi: 10.1017/S1049096523000926. URL https://www.cambridge.org/core/journals/ps-political-science-and-politics/article/invincible-gender-gap-in-political-ambition/547D8E289D82299856644844269E2FC9?utm_source=SFMC&utm_medium=email&utm_content=Article&utm_campaign=New%20Cambridge%20Alert%20-%20Issues%20Society&WT.mc_id=New%20Cambridge%20Alert%20-%20Issues%20Society.

Lois Frankel. Democratic Women's Caucus, Speaker Pelosi Send Letter to Facebook Demanding it Stop the Spread of Gendered Disinformation and Misogynistic Attacks Against

Women Leaders, August 2020. URL https://frankel.house.gov/news/documentsingle.aspx?DocumentID=3153. Section: Press Releases.

Tamara Fuchs and Fabian Schäfer. Normalizing misogyny: hate speech and verbal abuse of female politicians on Japanese Twitter. *Japan Forum*, 33(4):553–579, October 2021. ISSN 0955-5803. doi: 10.1080/09555803.2019.1687564. URL https://doi.org/10.1080/09555803.2019.1687564. Publisher: Routledge _eprint: https://doi.org/10.1080/09555803.2019.1687564.

Sarah A. Fulton. Running Backwards and in High Heels: The Gendered Quality Gap and Incumbent Electoral Success. *Political Research Quarterly*, 65(2):303–314, June 2012a. ISSN 1065-9129. doi: 10.1177/1065912911401419. URL https://doi.org/10.1177/1065912911401419. Publisher: SAGE Publications Inc.

Sarah A. Fulton. Running Backwards and in High Heels: The Gendered Quality Gap and Incumbent Electoral Success. *Political Research Quarterly*, 65(2):303–314, June 2012b. ISSN 1065-9129. doi: 10.1177/1065912911401419. URL https://doi.org/10.1177/1065912911401419. Publisher: SAGE Publications Inc.

Genevieve Gorrell, Mehmet E. Bakir, Ian Roberts, Mark A. Greenwood, and Kalina Bontcheva. Which politicians receive abuse? Four factors illuminated in the UK general election 2019. *EPJ Data Science*, 9(1):1–20, December 2020. ISSN 2193-1127. doi: 10.1140/epjds/s13688-020-00236-9. URL https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-020-00236-9. Number: 1 Publisher: SpringerOpen.

Mark A. Greenwood, Mehmet E. Bakir, Genevieve Gorrell, Xingyi Song, Ian Roberts, and Kalina Bontcheva. Online Abuse of UK MPs from 2015 to 2019: Working Paper, April 2019. URL http://arxiv.org/abs/1904.11230. arXiv:1904.11230 [cs].

Cécile Guerin and Eisha Maharasingam-Shah. Public Figures, Public Rage: Candidate abuse on social media. Technical report, October 2020. URL https://www.isdglobal.org/isd-publications/public-figures-public-rage-candidate-abuse-on-social-media/.

Emily Harmer and Rosalynd Southern. Digital microaggressions and everyday othering: an analysis of tweets sent to women members of Parliament in the UK. In *Women in the Digital World*. Routledge, 2023. ISBN 978-1-00-337593-7. Num Pages: 18.

Danny Hayes and Jennifer L. Lawless. *Women on the Run: Gender, Media, and Political Campaigns in a Polarized Era*. Cambridge University Press, May 2016. ISBN 978-1-316-72077-6. Google-Books-ID: fXNNDAAAQBAJ.

Caroline Heldman, Meredith Conroy, and Alissa R. Ackerman. *Sex and Gender in the 2016 Presidential Election*. 2018. URL https://www.abc-clio.com/products/a5566c/.

Rebekah Herrick and Lori D. Franklin. Is it Safe to Keep This Job? The Costs of Violence on the Psychological Health and Careers of U.S. Mayors. *Social Science Quarterly*, 100

(6):2047–2058, 2019. ISSN 0038-4941. URL https://www.jstor.org/stable/26788620. Publisher: [University of Texas Press, Wiley].

Rebekah Herrick and Sue Thomas. Not Just Sticks and Stones: Psychological Abuse and Physical Violence among U.S. State Senators. *Politics & Gender*, 18(2):422–447, June 2022. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X2000063X. URL https://www.cambridge.org/core/journals/politics-and-gender/article/not-just-sticks-and-stones-psychological-abuse-and-physical-violence-among-us-state-senators/0546394A409652DDEE619E47481DC216. Publisher: Cambridge University Press.

Rebekah Herrick, Sue Thomas, Lori Franklin, Marcia L. Godwin, Eveline Gnabasik, and Jean Reith Schroedel. Physical violence and psychological abuse against female and male mayors in the United States. *Politics, Groups, and Identities*, 9(4):681–698, June 2019. ISSN 2156-5503. doi: 10.1080/21565503.2019.1629321. URL https://doi.org/10.1080/21565503.2019.1629321. Publisher: Routledge _eprint: https://doi.org/10.1080/21565503.2019.1629321.

Rebekah Herrick, Sue Thomas, and Kate Bartholomy. Gender, Power, and Colleague Aggression in U.S. State Senates. *Political Research Quarterly*, 75(1):134–146, March 2022. ISSN 1065-9129. doi: 10.1177/1065912920985313. URL https://doi.org/10.1177/1065912920985313. Publisher: SAGE Publications Inc.

Malin Holm, Elin Bjarnegård, and Pär Zetterberg. Comparing Gendered Exposure and Impact in Online Election Violence: Tunisian Political Candidates Targeted on Facebook. *Politics & Gender*, 20(4):879–902, December 2024. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X24000163. URL https://www.cambridge.org/core/journals/politics-and-gender/article/comparing-gendered-exposure-and-impact-in-online-election-violence-tunisian-political-candidates-targeted-on-facebook/2922D344C52536D71C63AD7E3393200C?utm_source=SFMC&utm_medium=email&utm_content=Article&utm_campaign=New%20Cambridge%20Alert%20-%20Issues%20Society&WT.mc_id=New%20Cambridge%20Alert%20-%20Issues%20Society.

Mirya R. Holman. Women in Local Government: What We Know and Where We Go from Here. *State & Local Government Review*, 49(4):285–296, 2017. ISSN 0160-323X. URL https://www.jstor.org/stable/48571775. Publisher: [Sage Publications, Inc., Carl Vinson Institute of Government, University of Georgia].

Sandra Håkansson. Do Women Pay a Higher Price for Power? Gender Bias in Political Violence in Sweden. *The Journal of Politics*, 83(2):515–531, April 2021. ISSN 0022-3816. doi: 10.1086/709838. URL https://www.journals.uchicago.edu/doi/full/10.1086/709838. Publisher: The University of Chicago Press.

Sandra Håkansson. The Gendered Representational Costs of Violence against Politicians. *Perspectives on Politics*, 22(1):81–96, March 2024. ISSN 1537-5927, 1541-0986. doi: 10.1017/S1537592723001913. URL https://www.cambridge.org/core/journals/perspectives-on-politics/article/gendered-representational-costs-of-violence-against-politicians/605158B81B587D8AE859ABAA47DFD4C5.

Sandra Håkansson and Nazita Lajevardi. Representation at Risk: Evaluating Levels and Consequences of Violence against Immigrant-Background Politicians. *American Political Science Review*, pages 1–16, December 2024. ISSN 0003-0554, 1537-5943. doi: 10.1017/S000305542400100X. URL https://www.cambridge.org/core/journals/american-political-science-review/article/representation-at-risk-evaluating-levels-and-consequences-of-violence-against-immigrantbackground-politicians/CAF77FD5DAB72B7FF434AD70BB18126A.

IPU. Sexism, harassment and violence against women in parliaments in Europe. Technical report, 2018. URL https://www.ipu.org/resources/publications/issue-briefs/2018-10/sexism-harassment-and-violence-against-women-in-parliaments-in-europe.

Emma A Jane. Systemic misogyny exposed: Translating Rapeglish from the Manosphere with a Random Rape Threat Generator. *International Journal of Cultural Studies*, 21 (6):661–680, November 2018. ISSN 1367-8779. doi: 10.1177/1367877917734042. URL https://doi.org/10.1177/1367877917734042. Publisher: SAGE Publications Ltd.

Nina Jankowicz. How disinformation became a new threat to women, December 2017. URL https://www.codastory.com/disinformation/how-disinformation-became-a-new-threat-to-women/.

Nina Jankowicz. *How to Be a Woman Online: Surviving Abuse and Harassment, and How to Fight Back*. Bloomsbury Publishing, 1 edition, April 2022.

Elisabeth A. Jarman. Clicks and Stones: Women Politicians and Gendered Hostility Online.

Shannon Jenkins. A Woman's Work Is Never Done? Fund-Raising Perception and Effort among Female State Legislative Candidates. *Political Research Quarterly*, 60(2):230–239, June 2007. ISSN 1065-9129. doi: 10.1177/1065912907301682. URL https://doi.org/10.1177/1065912907301682. Publisher: SAGE Publications Inc.

Jackson Katz. Man Enough?:Donald Trump, Hillary Clinton, and the Politics of Presidential Masculinity - Harvard Book Store, April 2016. URL https://www.harvard.com/book/man_enoughdonald_trump_hillary_clinton_and_the_politics_of_presidential_mas/.

Frederik Klaaborg Kjøller and Lene Holm Pedersen. The Gendered Cost of Politics. *British Journal of Political Science*, 55:e59, January 2025. ISSN 0007-1234, 1469-2112. doi: 10.1017/S0007123424000310. URL https://www.cambridge.org/core/journals/british-journal-of-political-science/article/gendered-cost-of-politics/2EADF7209471F07DDFEFB7F4D19A594D.

Karina Kosiara-Pedersen. Single ladies and freedom of speech: gendered explanations for, and effects of, violence in politics | Request PDF. *European Journal of Politics and Gender*, 7(2):221–238, April 2023. doi: https://doi.org/10.1332/251510821X16777808903018. URL https://www.researchgate.net/publication/371716551_Single_ladies_and _freedom_of_speech_gendered_explanations_for_and_effects_of_violence_in_politics.

Mona Krook and Juliana Restrepo. The Cost of Doing Politics? Analyzing Violence and Harassment against Female Politicians. *Perspectives on Politics*, 18:1–16, July 2019. doi: 10.1017/S1537592719001397.

Mona Lena Krook. Violence Against Women in Politics. *Journal of Democracy*, 28(1):74–88, 2017. ISSN 1086-3214. URL https://muse.jhu.edu/pub/1/article/645539. Publisher: Johns Hopkins University Press.

Mona Lena Krook. *Violence against Women in Politics*. Oxford University Press, July 2020. ISBN 978-0-19-008848-4. Google-Books-ID: flXwDwAAQBAJ.

Mona Lena Krook and Juliana Restrepo Sanín. Gender and political violence in Latin America. Concepts, debates and solutions. *Política y gobierno*, 23(1):127–162, June 2016. ISSN 1665-2037. URL http://www.scielo.org.mx/scielo.php?script=sci_abstract&pid=S1665-20372016000100127&lng=es&nrm=iso&tlng=en. Publisher: Centro de Investigación y Docencia Económicas A.C., División de Estudios Políticos.

Jeffrey Lazarus and Amy Steigerwalt. *Gendered Vulnerability: How Women Work Harder to Stay in Office*. University of Michigan Press, 2018. ISBN 978-0-472-13071-9. doi: 10.3998/mpub.9718595. URL https://www.jstor.org/stable/10.3998/mpub.9718595.

Jennifer E. Manning. Membership of the 118th Congress: A Profile. Technical Report R47470, Congressional Research Service, July 2024.

Jane Mansbridge and Shauna L. Shames. Toward a Theory of Backlash: Dynamic Resistance and the Central Role of Power. *Politics & Gender*, 4(4):623–634, December 2008. ISSN 1743-9248, 1743-923X. doi: 10.1017/S1743923X08000500. URL https://www.cambridge.org/core/journals/politics-and-gender/article/abs/toward-a-theory-of-backlash-dynamic-resistance-and-the-central-role-of-power/D368EF191B6ACD06D4ABA028B50D7265. Publisher: Cambridge University Press.

Joshua McCrain and Stephen D. O'Connell. Experience, institutions, and candidate emergence: the political career returns to state legislative service. *Political Science Research and Methods*, 11(4):763–784, October 2023. ISSN 2049-8470, 2049-8489. doi: 10.1017/psrm.2022.16. URL https://www.cambridge.org/core/journals/political-science-research-and-methods/article/experience-institutions-and-candidate-emergence-the-political-career-returns-to-state-legislative-service/CD6333A6699D19E3E71A6C1943FBA55B.

Robert J. McGrath, Josh Ryan, and Jatia Wrighten. High Hurdles: Legislative Professionalism and the Effectiveness of Women State Legislators. *The Journal of Politics*, May 2025. ISSN 0022-3816. doi: 10.1086/736866. URL https://www.journals.uchicago.edu/doi/10.1086/736866. Publisher: The University of Chicago Press.

Amanda Montell. *Wordslut: A Feminist Guide to Taking Back the English Language*. HarperCollins, May 2019. ISBN 978-0-06-286889-3. Google-Books-ID: kKpsDwAAQBAJ.

National Press Foundation. More Women Hold Elective Office, But Not Necessarily More Power, Say Rutgers Researchers, April 2024. URL https://www.youtube.com/watch?v=tCrhl5eszfc.

Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. BERTweet: A pre-trained language model for English Tweets. In Qun Liu and David Schlangen, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.2. URL https://aclanthology.org/2020.emnlp-demos.2.

Sarah Oates, Olya Gurevich, Christopher Walker, and Lucina Di Meco. Running While Female: Using AI to Track how Twitter Commentary Disadvantages Women in the 2020 U.S. Primaries, August 2019. URL https://papers.ssrn.com/abstract=3444200.

Tyler G. Okimoto and Victoria L. Brescoll. The price of power: power seeking and backlash against female politicians. *Personality & Social Psychology Bulletin*, 36(7):923–936, July 2010. ISSN 1552-7433. doi: 10.1177/0146167210371949.

Michele Pathé, Jane Phillips, Elke Perdacher, and Edward Heffernan. The Harassment of Queensland Members of Parliament: A Mental Health Concern. *Psychiatry, Psychology and Law*, 21:577–584, September 2013. doi: 10.1080/13218719.2013.858388.

Kathryn Pearson and Eric McGhee. What It Takes to Win: Questioning "Gender Neutral" Outcomes in U.S. House Elections. *Politics & Gender*, 9(4):439–462, December 2013. ISSN 1743-923X, 1743-9248. doi: 10.1017/S1743923X13000433. URL https://www.cambridge.org/core/journals/politics-and-gender/article/what-it-takes-to-win-questioning-gender-neutral-outcomes-in-us-house-elections/83267D037A804CE682D5A640DA7B27E0. Publisher: Cambridge University Press.

Rasmus T. Pedersen, Niels Bjørn Grund Petersen, and Mads Thau. Online Abuse of Politicians: Experimental Evidence on Politicians' Own Perceptions. *Political Behavior*, June 2024. ISSN 1573-6687. doi: 10.1007/s11109-024-09944-8. URL https://doi.org/10.1007/s11109-024-09944-8.

Rasmus T. Pedersen, Mads Thau, and Niels Bjørn Grund Petersen. Exposure to Online Abuse of Politicians Does Not Scare Citizens Away From Politics. *Journal of Experimental Political Science*, pages 1–14, January 2025. ISSN 2052-2630, 2052-2649. doi: 10.1017/XPS.2024.20. URL https://www.cambridge.org/core/journals/journal-of-experimental-political-science/article/exposure-to-online-abuse-of-politicians-does-not-scare-citizens-away-from-politics/E5D794D11029D46493F93F02B6856295.

Jess Phillips. *Everywoman: One Woman's Truth about Speaking the Truth*. Hutchinson, 2017. ISBN 978-1-78633-077-2. Google-Books-ID: GlVavgAACAAJ.

Nirmal Puwar. *Space Invaders: Race, Gender and Bodies Out of Place*. Berg Publishers, September 2004. ISBN 978-1-85973-654-8. Google-Books-ID: NwK1AAAAIAAJ.

Gowri Ramachandran, Chisun Lee, Maya Kornberg, Kimberly Peeler-Allen, Ruby Edlin, Julia Fishman, Jiyoon Park, and Grady Yuthok Short. Intimidation of State and Local Officeholders: The Threat to Democracy. Technical report, Brennan Center for Justice, January 2024.

Ludovic Rheault, Erica Rayment, and Andreea Musulan. Politicians in the line of fire: Incivility and the treatment of women on social media. *Research & Politics*, 6(1): 2053168018816228, January 2019. ISSN 2053-1680. doi: 10.1177/2053168018816228. URL https://doi.org/10.1177/2053168018816228. Publisher: SAGE Publications Ltd.

Juliana Restrepo Sanín. Violence against women in politics as an unintended consequence of democratization. *International Feminist Journal of Politics*, 24(1):16–39, January 2022. ISSN 1461-6742, 1468-4470. doi: 10.1080/14616742.2021.2014343. URL https://www.tandfonline.com/doi/full/10.1080/14616742.2021.2014343.

Carly Schmitt and Hanna K. Brant. Gender, Ambition, and Legislative Behavior in the United States House. *Journal of Women, Politics & Policy*, 40(2):286–308, April 2019. ISSN 1554-477X. doi: 10.1080/1554477X.2019.1570757. URL https://doi.org/10.1080/1554477X.2019.1570757. Publisher: Routledge _eprint: https://doi.org/10.1080/1554477X.2019.1570757.

Stephanie Schriock and Christina Reynolds. *Run to Win: Lessons in Leadership for Women Changing the World*. Penguin, January 2021. ISBN 978-1-5247-4680-3. Google-Books-ID: JrUlEAAAQBAJ.

Shauna L. Shames. *Out of the Running: Why Millennials Reject Political Careers and Why It Matters*. NYU Press, 2017. ISBN 978-1-4798-2599-8. doi: 10.2307/j.ctt1bj4rv1. URL https://www.jstor.org/stable/j.ctt1bj4rv1.

SheShouldRun. Learn How to Navigate Social Media as a Woman Elected Official, August 2022. URL https://sheshouldrun.org/resources/navigating-social-media-from-your-biggest-fans-to-your-worst-trolls/.

Jure Skubic. Violence Against Women in Politics on Social Media: The Case of Slovenia, May 2025.

Sarah Sobieraj. Bitch, slut, skank, cunt: patterned resistance to women's visibility in digital publics. *Information, Communication & Society*, 21(11):1700–1714, November 2018. ISSN 1369-118X. doi: 10.1080/1369118X.2017.1348535. URL https://doi.org/10.1080/1369118X.2017.1348535. Publisher: Routledge _eprint: https://doi.org/10.1080/1369118X.2017.1348535.

Sarah Sobieraj. *Credible threat: attacks against women online and the future of democracy*. Oxford studies in digital politics. Oxford University Press, New York, NY, 2020. ISBN 978-0-19-008928-3.

Rosalynd Southern and Emily Harmer. Othering Political Women: Online Misogyny, Racism and Ableism Towards Women in Public Life. In *Online Othering: Exploring Digital Violence and Discrimination on the Web*, Palgrave Studies in Cybercrime and Cybersecurity, pages 187–210. Palgrave Macmillan, Cham, April 2019.

Chris Tenove, Tworek , Heidi, Lore , Grace, Buffie , Jordan, , and Trevor Deley. Damage Control: How Campaign Teams Interpret and Respond to Online Incivility. *Political Communication*, 40(3):283–303, May 2023. ISSN 1058-4609. doi: 10.1080/10584609.2022.2137743. URL https://doi.org/10.1080/10584609.2022.2137743. Publisher: Routledge _eprint: https://doi.org/10.1080/10584609.2022.2137743.

Yannis Theocharis, Pablo Barberá, Zoltán Fazekas, and Sebastian Adrian Popa. The Dynamics of Political Incivility on Twitter. *SAGE Open*, 10(2):2158244020919447, April 2020. ISSN 2158-2440. doi: 10.1177/2158244020919447. URL https://doi.org/10.1177/2158244020919447. Publisher: SAGE Publications.

Sue Thomas, Rebekah Herrick, Lori D. Franklin, Marcia L. Godwin, Eveline Gnabasik, and Jean R. Schroedel. Not for the Faint of Heart: Assessing Physical Violence and Psychological Abuse against U.S. Mayors. *State & Local Government Review*, 51(1):57–67, 2019. ISSN 0160-323X. URL https://www.jstor.org/stable/48596710. Publisher: [Sage Publications, Inc., Carl Vinson Institute of Government, University of Georgia].

Danielle M. Thomsen and Bailey K. Sanders. Gender Differences in Legislator Responsiveness. *Perspectives on Politics*, 18(4):1017–1030, December 2020. ISSN 1537-5927, 1541-0986. doi: 10.1017/S1537592719003414. URL https://www.cambridge.org/core/journals/perspectives-on-politics/article/gender-differences-in-legislator-responsiveness/193D2C984B539E564DD776E626F6D294.

Danielle M. Thomsen and Michele L. Swers. Which Women Can Run? Gender, Partisanship, and Candidate Donor Networks. *Political Research Quarterly*, 70(2):449–463, June 2017. ISSN 1065-9129. doi: 10.1177/1065912917698044. URL https://doi.org/10.1177/1065912917698044. Publisher: SAGE Publications Inc.

Angelia Wagner. Tolerating the trolls? Gendered perceptions of online harassment of politicians in Canada. *Feminist Media Studies*, 22(1):32–47, January 2022. ISSN 1468-0777. doi: 10.1080/14680777.2020.1749691. URL https://doi.org/10.1080/14680777.2020.1749691. Publisher: Routledge _eprint: https://doi.org/10.1080/14680777.2020.1749691.

Stephen Ward and Liam McLoughlin. Turds, traitors and tossers: the abuse of UK MPs via Twitter. *The Journal of Legislative Studies*, 26(1):47–73, January 2020. ISSN 1357-2334. doi: 10.1080/13572334.2020.1730502. URL https://doi.org/10.1080/13572334.2020.1730502. Publisher: Routledge _eprint: https://doi.org/10.1080/13572334.2020.1730502.

Claudia Wilhelm and Sven Joeckel. Gendered Morality and Backlash Effects in Online Discussions: An Experimental Study on How Users Respond to Hate Speech Comments Against Women and Sexual Minorities. *Sex Roles*, 80(7):381–392, April 2019. ISSN 1573-2762. doi: 10.1007/s11199-018-0941-5. URL https://doi.org/10.1007/s11199-018-0941-5.

Womens Media Center. Silencing Women in Politics: The Costs to Democracy of Gender-Based Online Harassment, November 2017. URL https://www.youtube.com/watch?v=qwpF-k36CWo.

Natalie Wood. Full- and Part-Time Legislatures, July 2021. URL https://www.ncsl.org/research/about-state-legislatures/full-and-part-time-legislatures.aspx.

Alan N. Yan and Rachel Bernhard. The Silenced Text: Field Experiments on Gendered Experiences of Political Participation. *American Political Science Review*, 118(1):481–487, February 2024. ISSN 0003-0554, 1537-5943. doi: 10.1017/S0003055423000217. URL https://www.cambridge.org/core/journals/american-political-science-review/article/silenced-text-field-experiments-on-gendered-experiences-of-political-participation/E008099CDD9573F7D8ADB128919EFDC4.

Kirsten Zeiter, Sandra Pepera, and Molly Middlehurst. Tweets That Chill: Analyzing Online Violence Against Women In Politics. Technical report, NDI, May 2019.

Supplementary Material for: Toll of Trolls: Gendered Online Hostility and its Impact on Women Politicians Annie JarmanPhD Candidate, Washington University in St. Louis. Email: a.a.jarman@wustl.edu. Web: https://anniejarman.github.io.

Contents

A Regressions in Paper

1 Visibility

Table 5: Visibility Regression Results

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
	(1)	(2)	(3)
Woman Legislator	−2.601	−1.559	2.603*
	(2.106)	(2.583)	(1.308)
Visibility	3.556***	0.949**	−0.012
	(0.359)	(0.380)	(0.135)
Count of Times Tweeted	−0.809***	−0.286	0.130*
	(0.282)	(0.191)	(0.077)
Black or Latino	1.056	0.809	0.359
	(0.816)	(0.756)	(0.228)
Democrat	−6.272***	0.190	1.542***
	(0.667)	(0.714)	(0.254)
Ideology	6.884**	0.584	−0.023
	(3.078)	(2.678)	(0.810)
Positive Sentiment	−10.298***	−4.667***	0.497
	(1.937)	(1.030)	(0.402)
'Masculine' Issues	0.124*	−0.067	−0.001
	(0.070)	(0.052)	(0.037)
'Feminine' Issues	−0.235**	−0.127*	−0.073**
	(0.097)	(0.075)	(0.033)
Opinion	5.326	−10.910**	1.963
	(7.824)	(5.342)	(2.339)
Factual Claim	9.285	−9.067*	−1.556
	(8.981)	(5.307)	(2.571)
Asks for Donation	57.402	−15.325	−1.755
	(42.157)	(54.440)	(7.387)
No Policy Content	4.003	5.923**	0.036
	(4.316)	(2.582)	(1.177)
Woman*Visibility	0.402	0.955**	0.081
	(0.342)	(0.400)	(0.205)
Observations	1,839	1,839	1,839
Adjusted R^2	0.294	0.107	0.225
Note:			*p<0.1; **p<0.05; ***p<0.01 Standard errors in parentheses State fixed effects included in each model

2 Legislator Tone

Table 6: Legislator Tone Regression Results

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
	(1)	(2)	(3)
Woman Legislator	−0.152	2.665	2.731
	(2.649)	(3.232)	(1.813)
Visibility	3.589***	1.007**	−0.011
	(0.367)	(0.383)	(0.139)
Count of Times Tweeted	−6.312***	0.121	1.540***
	(0.661)	(0.709)	(0.258)
Black or Latino	−0.825***	−0.312	0.129
	(0.281)	(0.187)	(0.078)
Democrat	1.104	0.892	0.361
	(0.812)	(0.752)	(0.232)
Ideology	7.064**	0.895	−0.013
	(2.989)	(2.561)	(0.813)
Positive Sentiment	−9.467***	−3.234**	0.540
	(2.086)	(1.270)	(0.469)
'Masculine' Issues	0.126*	−0.063	−0.001
	(0.072)	(0.049)	(0.037)
'Feminine' Issues	−0.236**	−0.128*	−0.073**
	(0.099)	(0.076)	(0.033)
Opinion	5.190	−11.145**	1.956
	(7.660)	(5.030)	(2.334)
Factual Claim	58.754	−12.994	−1.684
	(42.237)	(54.677)	(7.542)
Asks for Donation	4.201	6.264**	0.046
	(4.327)	(2.477)	(1.191)
No Policy Content	9.123	−9.348*	−1.564
	(8.832)	(5.021)	(2.571)
Woman*Visibility	0.304	0.786*	0.076
	(0.348)	(0.416)	(0.221)
Woman*Positive Sentiment	−3.873	−6.681***	−0.202
	(2.506)	(2.238)	(1.091)
Observations	1,839	1,839	1,839
Adjusted R^2	0.295	0.112	0.225

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses
State fixed effects included in each model

3 Growth Rate

Table 7: Growth Rate Regression Results

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
Intercept	3.40	14.41***	2.21
	(6.17)	(5.23)	(2.03)
Woman Legislator	−1.42	−1.22	2.27***
	(2.30)	(1.95)	(0.76)
Visibility	3.35***	0.87***	−0.04
	(0.28)	(0.23)	(0.09)
Democrat	−6.38***	0.08	1.44***
	(0.64)	(0.53)	(0.21)
Count of Times Tweeted	−0.78***	−0.23	0.14**
	(0.19)	(0.16)	(0.06)
Black or Latino	0.93	0.89	0.43
	(0.85)	(0.72)	(0.28)
Ideology	6.80***	1.57	0.28
	(2.38)	(2.01)	(0.78)
Positive Sentiment	−10.98***	−5.03***	0.38
	(1.22)	(1.03)	(0.40)
'Masculine' Issues	0.11**	−0.06	0.00
	(0.05)	(0.05)	(0.02)
'Feminine' Issues	−0.22**	−0.10	−0.07**
	(0.09)	(0.07)	(0.03)
Opinion	6.43	−11.05**	2.38
	(5.91)	(5.03)	(1.95)
Asks for Donation	54.93	−19.85	−1.01
	(58.55)	(50.01)	(19.37)
No Policy Content	4.35	7.21***	0.17
	(2.72)	(2.29)	(0.89)
Factual Claim	9.93	−9.67*	−1.35
	(6.35)	(5.41)	(2.10)
Growth Rate	−0.03	0.03	0.01
	(0.03)	(0.02)	(0.01)
Woman*Visibility	0.34	1.01***	0.11
	(0.38)	(0.32)	(0.12)
Woman*Growth Rate	−0.06*	−0.05*	0.01
	(0.03)	(0.03)	(0.01)
AIC	14,126.00	13,513.84	10,065.67
BIC	14,230.82	13,618.66	10,170.49
Log Likelihood	-7,044.00	-6,737.92	-5,013.83
Observations	1,839	1,839	1,839
Var: state (Intercept)	16.47	2.87	0.67
Var: Residual	120.78	88.38	13.24
Note:			*p<0.1; **p<0.05; ***p<0.01 Standard errors in parentheses

B Labeling Mentions, Part 1

I utilized hand coding, zero-shot classification, and BERT feature representation to label the Twitter corpus for hostile and gendered content. I first classify the mentions with the 'gpt-4' model on the OpenAI API. [27] A team of expert coders given identical instructions to GPT-4 validate the zero-shot classification labels. [28] Next, I used the training corpus to finetune a pre-trained BERT model. [29]Validation of these labels indicated high accuracy and convergence.[30]

1 Instructions for ***Hostile*** Label

Hostility Classification: Label tweet 'hostile' or 'not hostile'. A tweet is 'hostile' if it does ANY of the following:

- Is strongly critical of @legislator
- Uses impolite, rude, or hurtful language
- Uses derogatory comments, swear words, vulgarities, threats, or hate speech
- Uses sarcasm or mocking tones, e.g., thanks a lot
- Has more than one question mark or an ellipsis '...'
- Uses dismissive or condescending language, including diminutives like 'Sweetie' or 'Honey', or suggests a statement by @legislator isn't important
- Includes sexual objectification or harassment
- Manifests racism, misogyny, homophobia, or attacks based on politics, religion, etc
- Uses ad hominem or personal attacks, such as accusing @legislator of having vices, being ineffective, corrupt, bigoted, etc
- Demands @legislator resign, belittles politicians, or advocates for electoral loss
- Cannot be confidently labeled as not hostile
- Includes hostile hashtags or emojis. Ignore URLs

[27]Before labeling the data, I removed all handles except the target legislator's in the text and replaced the legislator's handle with @legislator. The OpenAI annotation was conducted on August 15, 2023, with a model temperature of 0 to facilitate deterministic output. The expert coders annotated between August 6 and 9, 2023.

[28]The ICR between the zero-shot and expert classification for the hostile label is 0.89, and the F1 score is 0.75. The ICR for the gendered label is 0.93, and the F1 score is 0.74. These scores indicate that the zero-shot model has a relatively good balance between precision and recall.

[29]The HateBERT model was pre-trained on banned Reddit posts, making it ideal for detecting hostile content in Twitter mentions CaselliEA2021. A pre-trained BERT model for gendered content does not exist. Instead, I use BERTweet to label the corpus for gendered content nguyen$_{b}ertweet_{2}$020.

[30]A team of expert coders validated the BERT results. Both models are highly accurate, with an F1 score of 0.77 for hostility and 0.81 for gendered content.

2 Instructions for *Gendered* Label

Gendered Classification: Label tweet 'gendered' or 'not gendered'. A tweet is 'gendered' if it does ANY of the following:"

- Contains the specific words 'he', 'his', 'him', 'she', 'her', 'hers'. Pay special attention to these words
- Comments on @legislator's physical appearance, including wardrobe or pictures
- References @legislator's romantic or sexual relationships, marital or parental status, e.g., calling them a mother or father
- Makes judgments about competence based on @legislator's gender
- References stereotyped masculinity or femininity
- Mentions @legislator's gender explicitly, like stating they are a man, woman, congressman, or congresswoman
- Encourages the election of a specific gender
- Compares @legislator to an individual of the opposite gender
- Uses gendered slurs or slang like 'bitch', 'dick', 'asshole', 'dude', 'yes man', or 'girl boss'
- Cannot be confidently labeled as not gendered
- Includes gendered hashtags and emojis. Ignore URLs

3 Partitioning Training Data

To build the training corpus, I randomly sampled five to ten mentions[31] directed toward each legislator. The result is a scoring subset of 41,384 mentions on Twitter. I partitioned the subset into four groups with sizes 37895 (Group A), 388 (Group B), 3101 (Group C), and 388 (Group D) mentions, respectively. The first three sets (n= 41384) were all labeled using gpt-4 zero-shot classification on the OpenAI API, described below. Group B was also hand-labeled by a team of three expert coders to validate the zero-shot classification labels.[32] The 3101 mentions in Group C were labeled but held out from the model building. The 388 mentions in Group D were hand-labeled by the team of expert coders to validate the final measure and excluded from all other stages.

[31]Individuals who received fewer than five mentions are excluded from the analysis. There are 54 legislators who received fewer than five mentions. Because the dependent variables are percentage-based, these individuals frequently represented outliers.

[32]Zero-shot classification is a machine learning technique where a model classifies data into categories without seeing any examples of the categories during training; it uses the instructions it is given and its understanding of language and the world to make predictions.

Table 8: Partitioning training corpus for classification and validation

	Group A (n=37895)	Group B (n=388)	Group C (n=3101)	Group D (n=388)
Zero-shot classified label	✓	✓	✓	
Expert label		✓		✓
Held out from model training			✓	✓

4 Feature Representation and Model Validation

Table 9: BERT Model Prediction Scores

Metric	**Hostility (HateBERT)**	**Gendered (BERTweet)**
Accuracy	0.91	0.96
Recall	0.73	0.84
Precision	0.82	0.79
F1 score	0.77	0.81
ROC_AUC	0.95	0.98

5 Confusion Matrixes

Figure D.1 visualizes the zero-shot and BERT classification's performance. GPT -4 was more prone to Type I errors when classifying hostile content and Type II errors when classifying gendered content. GPT-4 correctly classified 286 tweets as not hostile and 59 tweets as hostile. It falsely classified 7 hostile tweets as neutral and 31 neutral tweets as hostile. GPT-4 correctly classified 332 tweets as not gendered and 30 as gendered. It falsely classified 17 gendered tweets as neutral and 4 neutral tweets as gendered.

BERTweet did not perform as well at identifying gendered content and was nearly as likely to report a false negative as a true positive. As with the zero-shot classification, BERT was more prone to Type I errors when classifying hostile content and Type II errors when classifying gendered content. BERT correctly classified 585 tweets as not hostile and 101 tweets as hostile. It falsely classified 24 hostile tweets as neutral and 67 neutral tweets as hostile. BERT correctly classified 675 tweets as not gendered and 53 as gendered. It falsely classified 43 gendered tweets as neutral and 6 neutral tweets as gendered.

The percentage of hostile and gendered tweets that men receive has been artificially inflated, while women's have been artificially deflated. When I examine confusion matrices by the gender of the legislator recipient, we see a large bias in the algorithms, depicted in Figure 3 below. Across all four models, the false positive rate for men is nearly double that of women. Excepting one model, men also have a lower false negative rate than women. This bias makes it more difficult to find results in line with my hypotheses. The following results, therefore, indicate a conservative test and suggest that this paper understates the true effect of legislator gender on receiving gendered content.[33]

[33]Logistic regression models do not show any evidence that the difference in the likelihood of being labeled as hostile or gendered by BERT vs. human coders is affected by the gender of the legislator the mention is directed at.

Figure 6: Confusion Matrixes

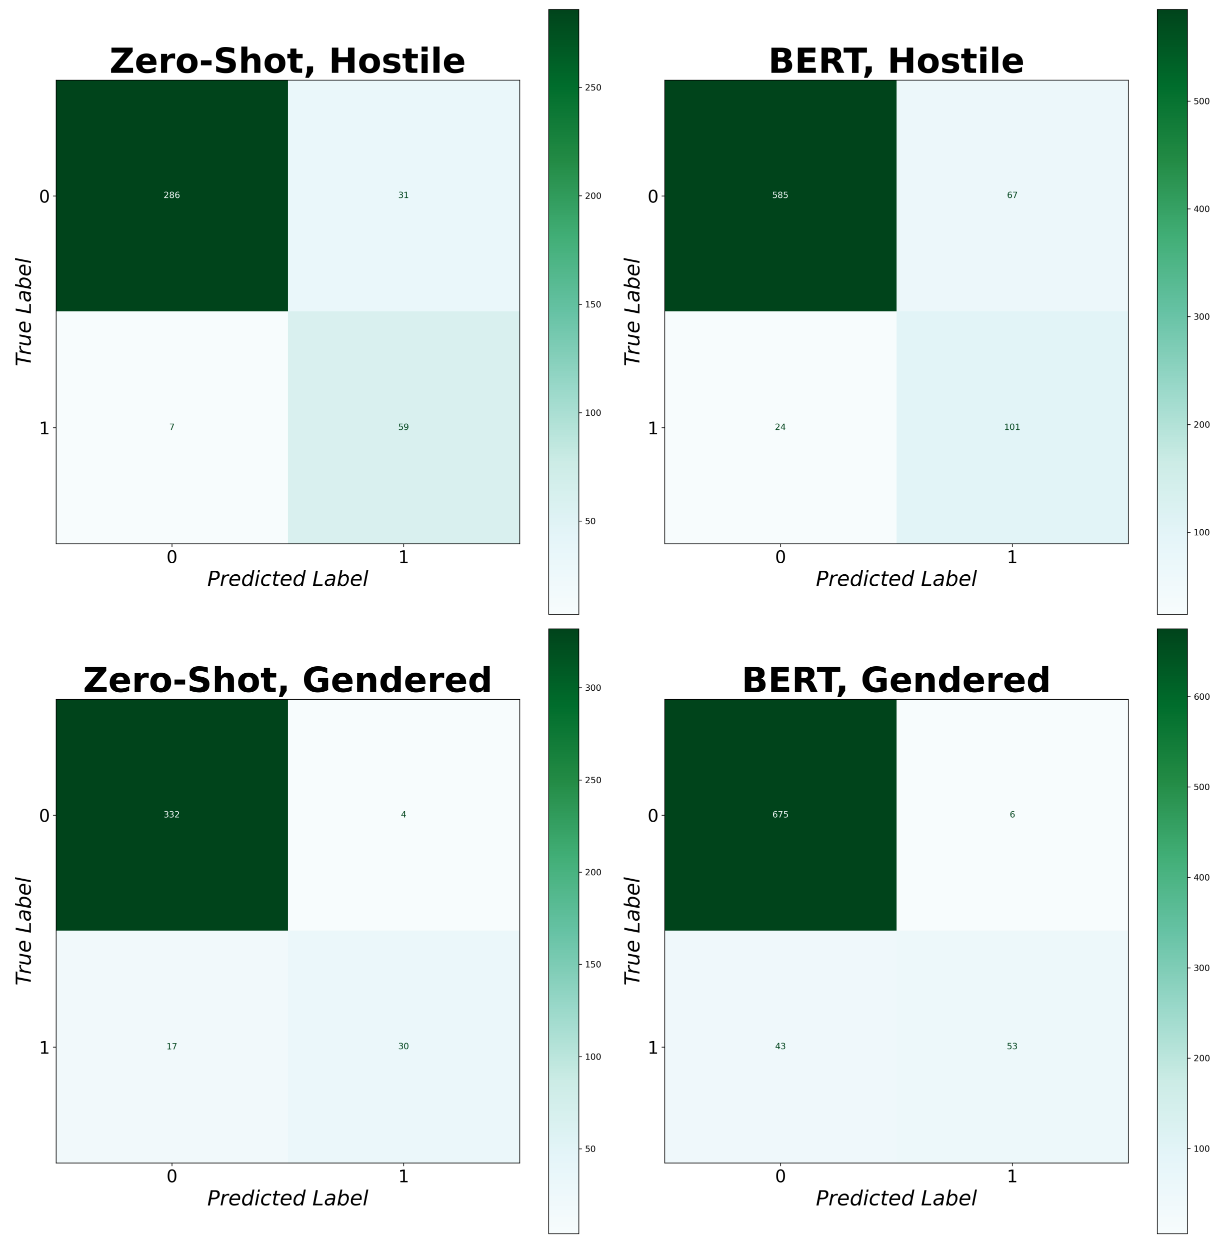

Figure 7: Algorithmic Bias Inflates Men's Percentages and Decreases Women's

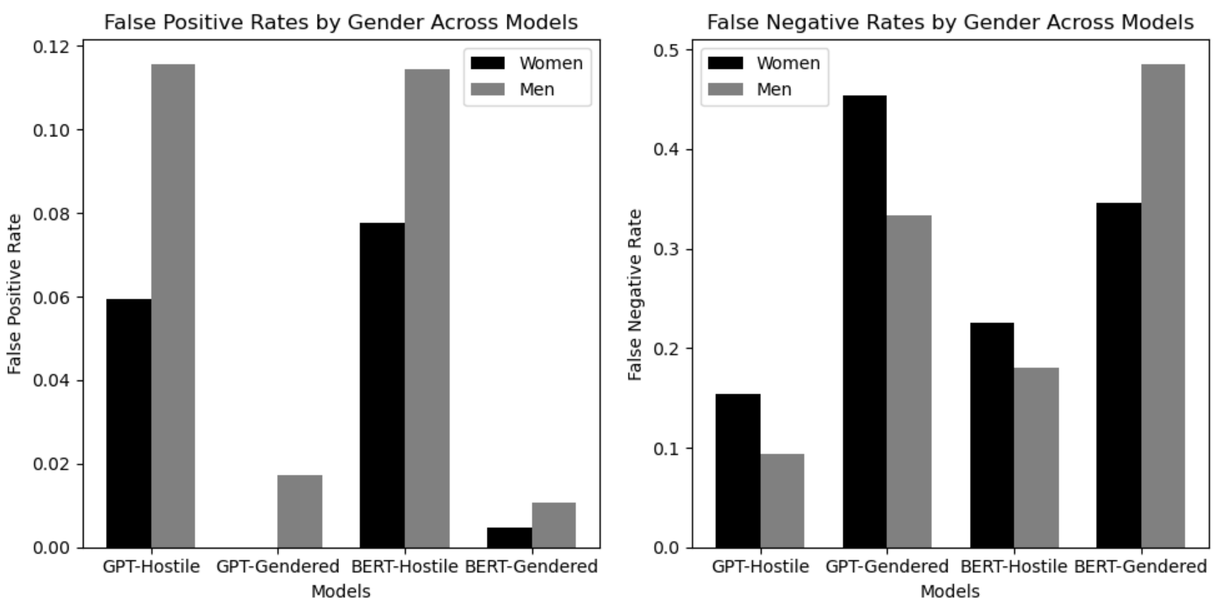

C Labeling Mentions, Part 2

1 Instructions for ***Hostile*** Label

Hostility Classification: Label tweet 'hostile' or 'not hostile'. A tweet is 'hostile' if it does ANY of the following:

- Is strongly critical of @legislator
- Uses impolite, rude, or hurtful language
- Uses derogatory comments, swear words, vulgarities, threats, or hate speech
- Uses sarcasm or mocking tones, e.g., thanks a lot
- Has more than one question mark or an ellipsis '...'
- Uses dismissive or condescending language, including diminutives like 'Sweetie' or 'Honey', or suggests a statement by @legislator isn't important
- Includes sexual objectification or harassment
- Manifests racism, misogyny, homophobia, or attacks based on politics, religion, etc
- Uses ad hominem or personal attacks, such as accusing @legislator of having vices, being ineffective, corrupt, bigoted, etc
- Demands @legislator resign, belittles politicians, or advocates for electoral loss
- Cannot be confidently labeled as not hostile
- Includes hostile hashtags or emojis. Ignore URLs

2 Instructions for ***Gendered*** Label

Gendered Classification: Label tweet 'gendered' or 'not gendered'. A tweet is 'gendered' if it does ANY of the following:"

- Contains the specific words 'he', 'his', 'him', 'she', 'her', 'hers'. Pay special attention to these words
- Comments on @legislator's physical appearance, including wardrobe or pictures
- References @legislator's romantic or sexual relationships, marital or parental status, e.g., calling them a mother or father
- Makes judgments about competence based on @legislator's gender
- References stereotyped masculinity or femininity

- Mentions @legislator's gender explicitly, like stating they are a man, woman, congressman, or congresswoman
- Encourages the election of a specific gender
- Compares @legislator to an individual of the opposite gender
- Uses gendered slurs or slang like 'bitch', 'dick', 'asshole', 'dude', 'yes man', or 'girl boss'
- Cannot be confidently labeled as not gendered
- Includes gendered hashtags and emojis. Ignore URLs

3 Instructions for ***Supportive*** Label

"*Supportive Classification*: Label tweet 'supportive' or 'not supportive'." A tweet is 'supportive' if it does ANY of the following:"

- Any 'thanks' (thanks/TY) in any context.
- Explicit support for the campaign (voting, donating, volunteering).
- Encouraging others to pray for @legislator or highlighting @legislator as a minority.
- Praise for legislative, public, or constituency work.
- Positive comparisons of @legislator to others.
- Compliments (bill passage, milestones, birthdays, sympathy).
- Friendly relationships or concern for @legislator's well-being.
- Positive interactions ('glad to be with')
- Vague niceties ('super', 'nice'), third-party updates, or cooperative actions without praise are not supportive.
- Hashtags/emojis can indicate support. Ignore URLs.

4 Partitioning Training Data

To build the training data, I randomly sampled tweets from each legislator to create a training corpus of 152,522 mentions, equal to 1% of the entire corpus. I partitioned the scoring subset into four groups, described in Table 1. The first three sets were all labeled using gpt-4-turbo zero-shot classification on the OpenAI API, described below. Groups B and D were hand-labeled by expert coders to validate the zero-shot classification labels and final measure, respectively. Groups A and B are the training data, whereas Groups C and D are held out from BERT finetuning to validate the final measure.

Table 10: Partitioning training corpus for classification and validation

	Group A (n=122,615)	**Group B** (n=2,991)	**Group C** (n=23,925)	**Group D** (n=2,991)
Zero-shot classified label	✓	✓	✓	
Expert label		✓		✓
Held out from model training			✓	✓

5 Feature Representation and Model Validation

Table 11: Zero-shot predictions (before DSL) vs. group B

Metric	**Gender**	**Hostile**	**Supportive**
Accuracy	0.9364	0.8929	0.9237
Recall	0.6256	0.8272	0.8610
Precision	0.7651	0.7558	0.7931
F1 score	0.6883	0.7899	0.8256
ROC_AUC	0.8007	0.8706	0.9007

Table 12: BERT predictions (before DSL) vs. groups B & D

Metric	**Gender**	**Hostile**	**Supportive**
Accuracy	0.9718	0.8861	0.9265
Recall	0.9113	0.8295	0.8850
Precision	0.8486	0.7362	0.7900
F1 score	0.8789	0.7801	0.8348
ROC_AUC	0.9454	0.8669	0.9113

D Robustness Checks for Part 2

1 Binomial Logit

Table 13: Logit: Exit vs. Not Exit with Cluster-Robust SEs

	Logit (Exit vs. Not Exit)
(Intercept)	0.520***
	(0.170)
Log Mentions	−0.070**
	(0.030)
Fundraising (Scaled)	0.000
	(0.050)
Democrat	0.010
	(0.130)
Legislative Effectiveness	−0.090
	(0.070)
Nonwhite	−0.280
	(0.170)
Donors (Scaled)	−0.440***
	(0.090)
Year First Elected (Scaled)	−0.270***
	(0.060)
In Office	−2.400***
	(0.150)
Woman	0.180
	(0.180)
General Hostile	1.100***
	(0.420)
Gendered Hostile	0.950
	(0.780)
General Hostile × Woman	−1.910**
	(0.910)
Gendered Hostile × Woman	−0.940
	(2.620)
N	2619
Clusters	1065
LogLik (full)	−936.336
LogLik (null)	−1276.356
AIC	1900.67
BIC	1982.86
McFadden R^2	0.2664
AUC (ROC)	0.7462
Pct. Correct @ 0.5	86.8%

Note: Standard errors clustered by individual legislator. Inactive cases coded as missing and dropped.
*p<0.1; **p<0.05; ***p<0.01.

2 Distinguishing between higher and lower offices

Table 14: Offices by Gender and Power Assessment

Office	Men	Women	Total	Likely to be More Powerful?
Governor	5	4	9	Yes
Lieutenant Governor	7	4	11	No (most states ceremonial)
State Senator	99	51	150	Yes
U.S. Senator	3	2	5	Yes
U.S. Representative	21	16	37	Yes
Secretary of State	5	4	9	Yes
Attorney General	4	1	5	Yes
State Auditor / Comptroller (incl. Missouri State Auditor, Auditor, Comptroller)	3	3	6	Yes
State Treasurer	1	2	3	Yes
Commissioner of Insurance	1	0	1	Yes
Superintendent of Public Instruction	1	0	1	Yes
State Supreme Court Justice (incl. "State Supreme Court")	1	1	2	Yes
Public Service Commission (incl. Montana exception)	2	2	4	No (except MT)*
Court of Appeals (state)	1	0	1	No
Trial Court Judge (civil, superior, county circuit)	2	4	6	No
Mayor	1	1	2	No (except big cities)
County Commission	4	1	5	No
City Council / Alderman	3	1	4	No
County Legislature	1	0	1	No
County Sheriff	1	0	1	No
District Attorney (U.S. Attorney if federal)	1	0	1	Yes
Local School / Education Board (incl. school boards, community college boards)	1	1	2	No
Local Financial / Administrative Officers (incl. county auditors, municipal treasurers, property appraisers)	0	3	3	No
Federal Agency State Director (incl. USDA, Farm Service Agency roles)	0	3	3	Yes
U.S. Ambassador	1	0	1	Yes
Deputy or Assistant to Elected Official (incl. Deputy Secretaries, Borough Presidents, Chiefs of Staff, Policy Directors, Budget Commissioners, Parole Board, Supervisor of Huntington)	11	3	14	No

3 Multinomial Logit, comparing higher vs. lower offices

Table 15: Multinomial Logit with Cluster-Robust SEs (Bootstrapped)

	Model 1 (With Hostility)			**Model 2 (Without Hostility)**		
	Coef.	SE	OR	Coef.	SE	OR
Likelihood of Campaign for Lower Office vs. Exit						
Intercept	-3.65***	(0.84)	0.03	-3.80***	(0.70)	0.02
Log Mentions (6 Months)	0.10	(0.31)	1.11	0.14	(0.32)	1.15
Fundraising (Scaled)	-0.15	(0.20)	0.86	-0.17	(0.23)	0.84
Democrat	0.34***	(0.10)	1.40	0.29***	(0.07)	1.34
SLES17	-0.24***	(0.05)	0.78	-0.24***	(0.04)	0.78
Nonwhite	1.45***	(0.04)	4.28	1.40***	(0.03)	4.06
Donors (Scaled)	0.45	(0.53)	1.57	0.48*	(0.28)	1.61
Year First Elected (Scaled)	0.54***	(0.13)	1.72	0.55***	(0.11)	1.73
In Office	0.65***	(0.06)	1.92	0.53***	(0.06)	1.70
Woman	-1.07***	(0.21)	0.34	-0.07	(0.48)	0.93
General Hostility	-1.26**	(0.61)	0.28	—	—	—
Gendered Hostility	1.31***	(0.14)	3.71	—	—	—
General Hostility × Woman	4.44***	(0.32)	84.54	—	—	—
Gendered Hostility × Woman	4.19***	(0.10)	66.09	—	—	—
Likelihood of Campaign for Higher Office vs. Exit						
Intercept	-1.96***	(0.07)	0.14	-2.04***	(0.22)	0.13
Log Mentions (6 Months)	0.18	(0.52)	1.20	0.16	(0.14)	1.17
Fundraising (Scaled)	-0.03	(0.27)	0.97	-0.04	(0.28)	0.97
Democrat	-0.30*	(0.18)	0.74	-0.27***	(0.10)	0.76
SLES17	0.07	(0.21)	1.07	0.08	(0.07)	1.08
Nonwhite	0.04	(0.13)	1.04	0.04	(0.43)	1.04
Donors (Scaled)	0.47***	(0.10)	1.59	0.49*	(0.28)	1.63
Year First Elected (Scaled)	0.42	(0.39)	1.52	0.42***	(0.16)	1.52
In Office	0.91***	(0.10)	2.49	0.91***	(0.15)	2.48
Woman	0.25	(0.66)	1.29	0.43***	(0.11)	1.54
General Hostility	-0.54***	(0.06)	0.58	—	—	—
Gendered Hostility	-1.24***	(0.25)	0.29	—	—	—
General Hostility × Woman	0.74***	(0.18)	2.09	—	—	—
Gendered Hostility × Woman	2.52*	(1.34)	12.41	—	—	—
Likelihood of Campaign for Reelection vs. Exit						
Intercept	-0.98	(0.69)	0.37	-1.13***	(0.09)	0.32
Log Mentions (6 Months)	0.05	(0.45)	1.05	0.02	(0.26)	1.02
Fundraising (Scaled)	0.03	(1.00)	1.03	0.03	(0.11)	1.03
Democrat	0.04	(0.27)	1.04	0.09	(0.06)	1.10
SLES17	0.10	(0.19)	1.10	0.10	(0.49)	1.11
Nonwhite	0.27	(2.20)	1.31	0.26	(0.24)	1.29
Donors (Scaled)	0.44	(1.45)	1.55	0.47***	(0.17)	1.59
Year First Elected (Scaled)	0.25	(1.05)	1.29	0.24	(0.34)	1.28
In Office	2.83	(2.92)	16.94	2.83***	(0.20)	17.01
Woman	-0.22	(0.99)	0.80	0.10	(0.14)	1.11
General Hostility	-1.16	(1.35)	0.32	—	—	—
Gendered Hostility	-1.02	(5.58)	0.36	—	—	—
General Hostility × Woman	2.16	(3.93)	8.68	—	—	—
Gendered Hostility × Woman	-0.41	(3.50)	0.66	—	—	—
N		2,925			2,925	
Log-Likelihood		-1732.83			-1741.70	
AIC		3549.70			3543.40	
BIC		3796.20			3719.50	
McFadden R^2		0.12			0.12	
Cox–Snell R^2		0.16			0.16	
Nagelkerke R^2		0.21			0.20	

Note: Bootstrapped (n=10,000) standard errors are clustered at the individual level. Odds ratios (OR) are exponentiated coefficients. Dashes (—) indicate terms not included in that model. The comparison group is those who exited politics. Significance: *p<0.1; **p<0.05; ***p<0.01.

4 Multinomial Logit, Odds Ratios instead of Coefficients

Table 16: Multinomial Logit with Cluster-Robust SEs (Bootstrapped)

	Different Office vs. Exit			**Reelection vs. Exit**		
	Coef.	SE	OR	Coef.	SE	OR
Constant	-1.760^{***}	(0.290)	0.17	-0.980^{***}	(0.100)	0.37
Log Mentions	0.170	(0.200)	1.18	0.050	(0.060)	1.05
Fundraising (Scaled)	-0.040	(0.050)	0.96	0.030	(0.230)	1.03
Democrat	-0.220^{***}	(0.040)	0.80	0.040	(0.180)	1.04
Legislative Effectiveness	0.050	(0.130)	1.05	0.100	(0.640)	1.10
Non-white	0.320^{***}	(0.060)	1.38	0.270	(0.450)	1.31
Count Donors (Scaled)	0.460^{**}	(0.210)	1.59	0.440^{*}	(0.260)	1.55
Year First Elected (Scaled)	0.440^{***}	(0.140)	1.55	0.250	(0.190)	1.29
In Office	0.880^{***}	(0.100)	2.41	2.830^{**}	(1.180)	16.96
Woman	0.070	(0.250)	1.07	-0.220	(1.300)	0.80
General Hostility	-0.670^{***}	(0.070)	0.51	-1.160	(1.010)	0.31
Gendered Hostility	-0.600^{***}	(0.180)	0.55	-1.010	(1.000)	0.36
General Hostility × Woman	1.310^{***}	(0.130)	3.71	2.160	(3.600)	8.63
Gendered Hostility × Woman	2.590^{***}	(0.100)	13.37	-0.440	(3.480)	0.64
N			2619			
LogLik			-1650.49			
AIC			3357.00			
BIC			3521.40			
McFadden R^2			0.12			
Cox–Snell R^2			0.15			
Nagelkerke R^2			0.20			

Note: Bootstrapped cluster-robust standard errors in parentheses. Odds ratios (OR) are exponentiated coefficients. The comparison group is *Exit*. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The odds ratios reported in Table **??** provide a more intuitive interpretation of the logit coefficients presented in the main text. Odds ratios above one indicate an increase in the likelihood of the outcome relative to exit, while odds ratios below one indicate a decrease. Among key predictors, general hostility and gendered hostility both reduce the odds of a man campaigning for a different office by about half. These results suggest that both forms of hostility act as deterrents to men's continued candidacy. However, the interaction terms show that women respond differently: the odds of women running for a different office increase by nearly four times under general hostility and by more than thirteen times under gendered hostility . This pattern indicates that women adapt to hostile environments by shifting offices rather than exiting politics altogether. Other predictors operate as expected. Higher fundraising, more donors, and earlier years of first election all increase the likelihood of continued candidacy, while nonwhite legislators are about 1.4 times more likely to campaign for a different office than white legislators. In contrast, being a Democrat decreases the odds of pursuing a different office. Taken together, these odds ratios complement the coefficient estimates by showing not only the direction and significance of effects but also their substantive magnitude in shaping political career trajectories.

E Measuring Visibility, California Example

Is the count of mentions a politician receives on Twitter a credible indicator of visibility? I check the accuracy of this measurement by examining a single state in depth. I do not have access to alternate measures of Twitter visibility – such as follower count or retweets. Instead, I examine six additional measures of offline visibility.

The first three measures relate to public notability. The first measure is the number of times the legislator was searched on Google[34]. I limited the search results to the state of California and within the topic of 'Laws and Government'. The second and third measurements are the number of times the legislator was referenced in a news article or report, obtained from Nexis Uni[35] and Factiva[36] respectively. For each of these sources, I limited the time frame to the same period as the tweets are collected from.[37] I limited the results to the state of California[38] and searched for the legislator's name with proper capitalization and inside of quotation marks.[39] I did not group duplicates (publications with similar language).

The remaining measures relate to situational notability. These measures are leadership position within the state assembly[40], whether the legislator campaigned for another

[34]This data was collected on Google Trends. I limited the search results to the state of California and within the topic of 'Laws and Government'. I used the proper capitalization of their name. Google Trends sometimes includes the option of specifying 'Legislator Name – State Assembly Member'. As this option is not available for all individuals, I did not use this specification. I checked the associated topics to ensure that they were related to government.

[35]Nexis Uni describes itself thus: "Contains more than 15,000 news, legal, and business sources, including journals, television and radio broadcasts, newswires and blogs; local, regional, national and international newspapers with deep archives; extensive legal sources for federal and state cases and statutes, including U.S. Supreme Court decisions since 1970; business information on more than 80 million U.S. and international companies and more than 75 million executives."

[36]Factiva describes itself thus: "Combines more than 35,000 sources for access to premium content from 200 countries in 26 languages. Worldwide full text coverage of local and regional newspapers, web and social media, tweets, digital video and audio clips, analyst reports, trade publications, business newswires, press release wires, media transcripts, news photos, market research reports, country and regional profiles, historical market data."

[37]October 1, 2015 – July 31, 2018

[38]By limiting my results to the state of California, I am not capturing national visibility. However, this precaution helps ensure that my results relate to the actual individual of interest. The speaker of the assembly at the time —Anthony Rendon— was likely mentioned in news articles beyond those sourced from Californian publications. His Factiva results within the state of California number at 860 references. This increases to 1245 results when the region is expanded to the United States. Without individually parsing through these 1245 articles, however, it is difficult to know whether they reference Anthony Rendon (speaker of the California state assembly) or Anthony Rendon (third baseman for the Washington Nationals MLB team). To avoid artificially inflating the numbers, I limit my results to California and assume that national visibility is closely correlated to state visibility.

[39]This ensures that the results contain the full name of the individual, ruling out articles that may use both the first and last name in separate cases.

[40]There were eleven leadership positions held by members of the state assembly in the sample. I grouped these positions based on their relative power to form an ordinal categorical variable. Higher numbers communicate greater power. A value of '0' signifies no leadership position. A '1' includes Republican floor manager, assistant majority whip, and deputy republican leader. A '2' includes assistant speaker pro tempore, majority whip, assistant majority leader, Republican caucus chair, and assistant Republican leader. A '3' includes speaker, minority leader, and speaker pro tempore.

office during the legislative term[41], and whether the legislator experienced any high-profile scandals[42] during the legislative term. This data was collected from individual pages on Ballotpedia and cross-referenced against individual pages on Wikipedia. Because only three individuals ran for another office[43] and only two experienced high-profile scandal, I do not include these measures in my final analysis.[44]

I run a linear regression to establish the correlation between these alternate measures of visibility and the indicator used in the paper. My dependent variable is the number of mentions a legislator received. The independent variables are the alternate measures of visibility detailed above, as well as the number of times the legislator tweeted in the time frame of interest. Factiva and Nexis Uni references show a high correlation of 0.78, so I run separate regressions to see the effect of each measure on the dependent variable without the influence of the other.

The results indicate that the count of Twitter mentions correlates strongly with alternate measures of visibility. A one-unit increase in Google searches, references in the news, the number of times a legislator tweeted, and leadership position are all statistically significant and positively correlated with the number of mentions a legislator received. Factiva's database is more than double the size of Nexis Uni, leading Model 2 to outperform Model 1 in adjusted R^2. The high adjusted R^2 value of 0.736 further lends confidence to the use of Twitter mentions as the indicator of visibility.

[41] Three individuals campaigned for another office during the legislative term.

[42] A scandal is considered high-profile if it merited mention on either the Ballotpedia or Wikipedia page. Two legislators fit this criterion – both for accusations of sexual harassment. One of the legislators resigned because of the accusations.

[43] Travis Allen ran for governor of California in 2018. He gained internet notoriety as "California's Trump" and used slogans like "Make California Great Again". He represented a clear outlier in the data and was removed from the analysis.

[44] This was problematic because it led the results to imply that experiencing a scandal or campaigning for another office decreased a legislator's visibility, when intuition and data firmly suggest the opposite.

Table 17: Regression Results

	Dependent variable:	
	Twitter Mentions	
	(1)	(2)
Google Trends	1.666	1.515^{*}
	(1.101)	(0.852)
Nexis Uni	38.153^{***}	
	(9.934)	
Factiva		21.247^{***}
		(2.984)
Leadership Position	$1,687.549^{***}$	$1,442.634^{***}$
	(533.351)	(421.145)
Count of Tweets	1.788^{***}	1.515^{***}
	(0.465)	(0.368)
Constant	$-4,168.194^{**}$	$-4,656.064^{***}$
	(1,722.087)	(1,342.153)
Observations	48	48
Adjusted R^2	0.572	0.736
Note:	*p<0.1; **p<0.05; ***p<0.01 Standard errors in parentheses	

F Race and Ethnicity

Previous research has found that politicians who are women of color are especially targeted on social media TumultyEA2020, marvelousai$_gender_2019, AmnestyInternational2019.Thisappendixexaminestheeffectofvisibilityonhosti$

First, I tally the percent of hostile mentions which are sent to white men, nonwhite men, white women, and nonwhite women. I repeat this for hostile and non-hostile gendered mentions. I further break down this relationship and examine it by party. The results suggest there may be a racial effect, however, these are not grouped by individual and do not have any tests of statistical significance.

To examine further, I run the same regressions from the paper, this time replacing the interaction between gender*visibility with race*visibility for the three hypotheses. The results are shown below. There is no indication of statistically significant relationships.

Next, I subset the data to just women and again create the plots with the interaction between race*visibility. The results are shown below. Again, there is no indication of a statistically significant relationship.

Finally, I subset the data to non-white politicians and replicate the models from the paper with the interaction between gender*visibility. The results are shown below and largely replicate the findings from the paper. For non-white politicians, visibility does not have a gendered impact on hostility, but it does have a gendered impact on gendered language. Interestingly, H3 about hostile mentions with gendered content is not replicated. This may be an artifact of the smaller sample size, however, as the slopes move in the expected directions.

The lack of racial effect with visibility and hostility may be due to the small number of non-white individuals in the dataset. The lack of effect between race and gendered language overall is not surprising. It does appear that nonwhite women receive similar percentages of gendered language as white women. The comparably smaller percentage of hostile mentions which are gendered for nonwhite women may be also due to smaller sample size, or it could be that a larger percentage of hostile mentions with racial content is driving down this number.

Table 18: Percent of Mentions Directed at Race & Gender Group

Demographic Group	**Hostile**	**Gendered**	**Hostile & Gendered**
Man, white or other	57.8%	45.5%	47.5%
Man, Black or Latino	9.8%	7.2%	6.5%
Woman, white or other	23.2%	34.6%	31.1%
Woman, Black or Latino	9.2%	12.8%	15.0%

1 Visibility * Black or Latino Legislator

Figure 8: Comparing Mentions Content by Race/Ethnicity.
Note: The x-axis shows *Visibility*, measured as the logged number of mentions a legislator received (e.g., a score of x equals 10^{x} mentions). From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile.

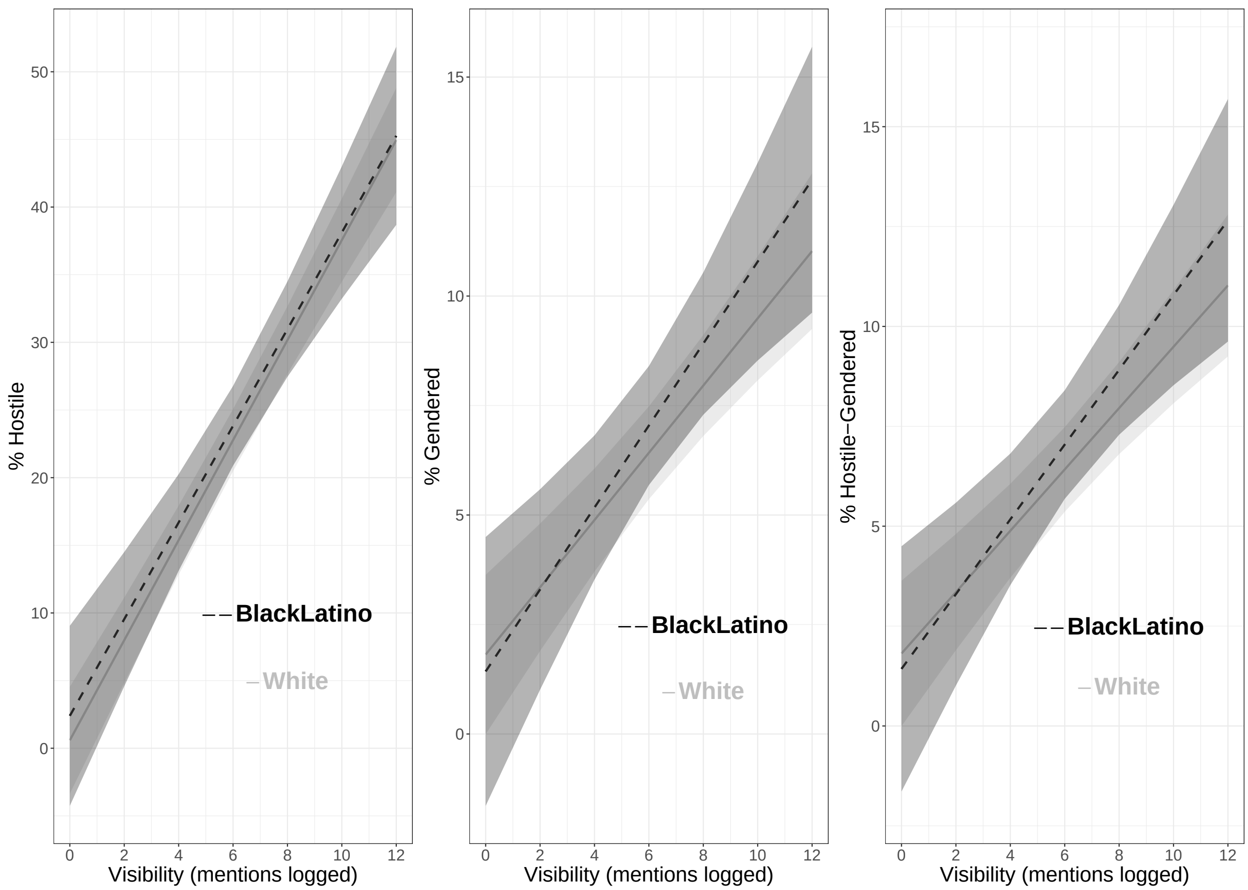

Table 19: Visibility Regression Results, Centered on Race

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
Black or Latino	1.80	1.52	−0.35
	(3.41)	(5.90)	(1.17)
Visibility	3.70***	1.26***	−0.00
	(0.35)	(0.31)	(0.13)
Democrat	−6.28***	0.20	1.55***
	(0.67)	(0.70)	(0.25)
Count of Times Tweeted	−0.82***	−0.31	0.13
	(0.28)	(0.19)	(0.08)
Woman Legislator	−0.26	4.00***	3.08***
	(0.60)	(0.57)	(0.29)
Ideology	6.81**	0.42	−0.03
	(3.11)	(2.63)	(0.81)
Positive Sentiment	−10.28***	−4.65***	0.48
	(1.96)	(1.02)	(0.41)
'Masculine' Issues	0.12*	−0.08	−0.00
	(0.07)	(0.05)	(0.04)
'Feminine' Issues	−0.24**	−0.13*	−0.07**
	(0.10)	(0.08)	(0.03)
Opinion	5.68	−10.15*	1.96
	(7.96)	(5.52)	(2.31)
Asks for Donation	58.80	−12.41	−1.80
	(42.26)	(55.85)	(7.34)
No Policy Content	4.01	5.96**	0.06
	(4.33)	(2.53)	(1.17)
Factual Claim	9.58	−8.48	−1.59
	(9.10)	(5.47)	(2.55)
Black or Latino × Visibility	−0.12	−0.12	0.12
	(0.56)	(0.89)	(0.19)
Observations	1,839	1,839	1,839
R^2 (Full Model)	0.32	0.13	0.25
R^2 (Projected)	0.23	0.08	0.20
Adj. R^2 (Full Model)	0.29	0.10	0.23
Adj. R^2 (Projected)	0.21	0.04	0.18
Num. Groups (state)	48	48	48

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses
State fixed effects included in each model

2 Visibility * Black or Latino Legislator, for Women Only

Figure 9: Comparing Mentions Content by Race/Ethnicity for Women Only.
Note: The x-axis shows *Visibility*, measured as the logged number of mentions a legislator received (e.g., a score of x equals 10^x mentions). From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile.

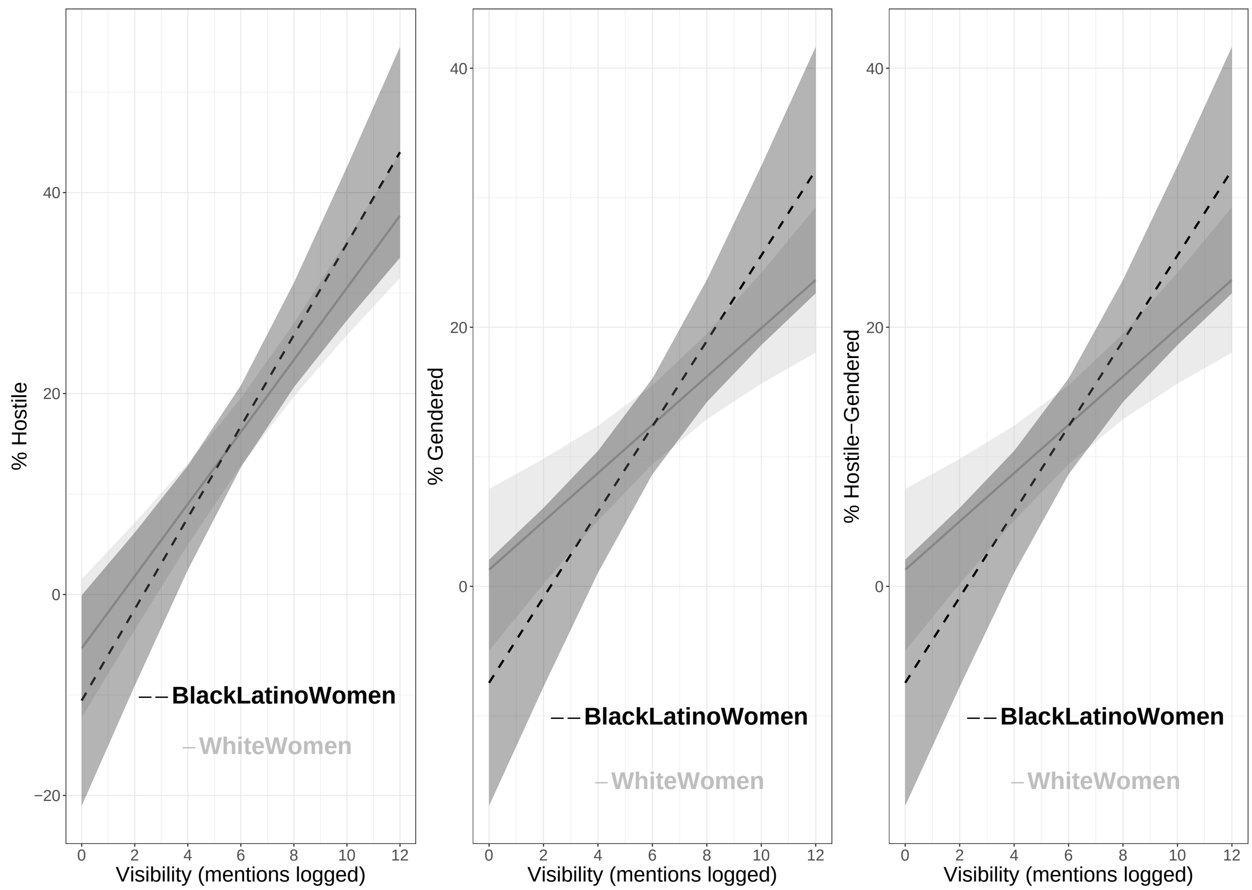

Table 20: Visibility Regression Results, Centered on Race (Subset: Non-White Legislators)

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
Black or Latino	−5.20	−8.73	−1.81
	(5.61)	(5.68)	(3.17)
Visibility	3.59***	1.87***	−0.11
	(0.52)	(0.52)	(0.26)
Democrat	−6.51***	−1.35	2.04***
	(1.35)	(1.27)	(0.51)
Count of Times Tweeted	−1.18***	−0.62*	0.28
	(0.38)	(0.37)	(0.18)
Ideology	7.09	−3.17	−4.72*
	(6.71)	(5.01)	(2.35)
Positive Sentiment	−15.82***	−11.06***	1.01
	(3.12)	(2.32)	(0.97)
'Masculine' Issues	0.09	−0.02	0.09**
	(0.08)	(0.07)	(0.04)
'Feminine' Issues	−0.02	−0.05	−0.18***
	(0.19)	(0.15)	(0.06)
Opinion	14.95	−13.93	0.99
	(15.99)	(12.69)	(5.29)
Asks for Donation	−661.48	437.21	439.89**
	(429.72)	(589.43)	(179.10)
No Policy Content	10.18	11.44*	0.87
	(8.37)	(6.21)	(2.71)
Factual Claim	14.75	−15.62	−7.43
	(17.43)	(12.38)	(5.38)
Black or Latino × Visibility	0.96	1.43*	0.35
	(0.92)	(0.84)	(0.50)
Observations	550	550	550
R^2 (Full Model)	0.40	0.24	0.27
R^2 (Projected)	0.27	0.12	0.16
Adj. R^2 (Full Model)	0.33	0.15	0.18
Adj. R^2 (Projected)	0.18	0.02	0.06
Num. Groups (state)	47	47	47
Note:			*p<0.1; **p<0.05; ***p<0.01 Standard errors in parentheses State fixed effects included in each model

3 Visibility * Women, for Black or Latino Only

Figure 10: Comparing Mentions Content by Gender for Black and Latinx Legislators.
Note: The x-axis shows *Visibility*, measured as the logged number of mentions a legislator received (e.g., a score of x equals 10^x mentions). From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile.

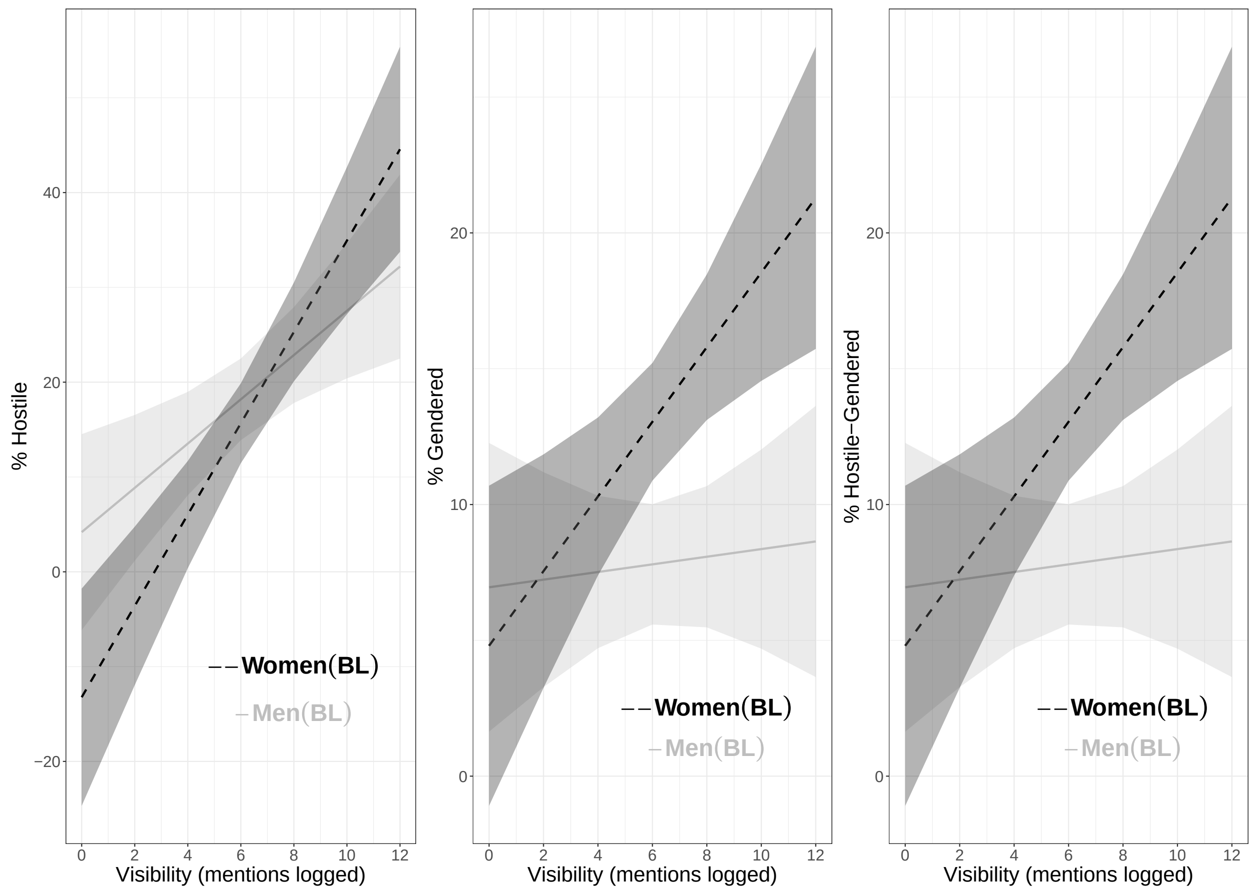

Table 21: Regression Results - Gender*Visibility for Non-white Only

	Dependent variable:		
	% Hostile	% Gendered	% Hostile Gendered
Woman Legislator	-17.415^{***}	-2.153	-12.417
	(5.321)	(4.304)	(9.309)
Visibility	2.334^{***}	0.141	0.059
	(0.650)	(0.355)	(1.200)
Count of Times Tweeted	-2.515	1.795^{**}	2.719
	(2.004)	(0.822)	(2.241)
Democrat	-0.943	0.465^{**}	-0.245
	(0.694)	(0.201)	(0.489)
Ideology	11.793^{**}	1.234	-8.961
	(5.589)	(2.878)	(7.955)
Positive Sentiment	-7.358	-2.031	-3.027
	(5.265)	(1.890)	(7.267)
'Masculine' Issues	0.126	-0.075	-0.260
	(0.260)	(0.064)	(0.196)
'Feminine' Issues	-0.409	-0.103	-0.100
	(0.253)	(0.087)	(0.117)
Opinion	36.269^{*}	1.565	-2.347
	(17.982)	(8.113)	(15.476)
Asks for Donation	-61.731	-114.407	-239.549
	(393.678)	(208.707)	(221.921)
No Policy Content	5.896	1.493	8.279
	(8.544)	(4.068)	(9.528)
Factual Claim	44.627^{**}	2.553	2.876
	(20.100)	(9.936)	(16.347)
Women Legislator*Visibility	2.482^{***}	1.234^{*}	2.458^{*}
	(0.879)	(0.662)	(1.328)
Observations	240	240	240
Adjusted R^2	0.350	0.265	0.157
Note:			*p<0.1; **p<0.05; ***p<0.01 Standard errors in parentheses State fixed effects included in each model

G Gendered Policy Areas

Some citizens believe men and women are differentially equipped to address specific issues Bauer2020, HerrnsonEA2003, Lawless2004, sanbonmatsu$_p oised_2 009.Butleretal.classifyeachrepresentative'stweetbytheissueareaitaddresses, follow$ $hostilecontentintheirmentionsthanwomenwhotweetabout\backslash masculine" topics.Noothersignificanteffects$

Figure 11: Comparing Mentions Content by Gender for Policy Areas.
Note: The x-axis on the top row shows *% Tweets on 'Feminine' Topics*, measured as the percentage of their tweets related to education or healthcare. The x-axis on the bottom row shows *% Tweets on 'Masculine' Topics*, measured as the percentage of their tweets related to macroeconomics and national security. From left to right, the y-axes represent the percentage of (1) mentions that were hostile, (2) hostile mentions that were gendered, and (3) mentions that were gendered but not hostile.

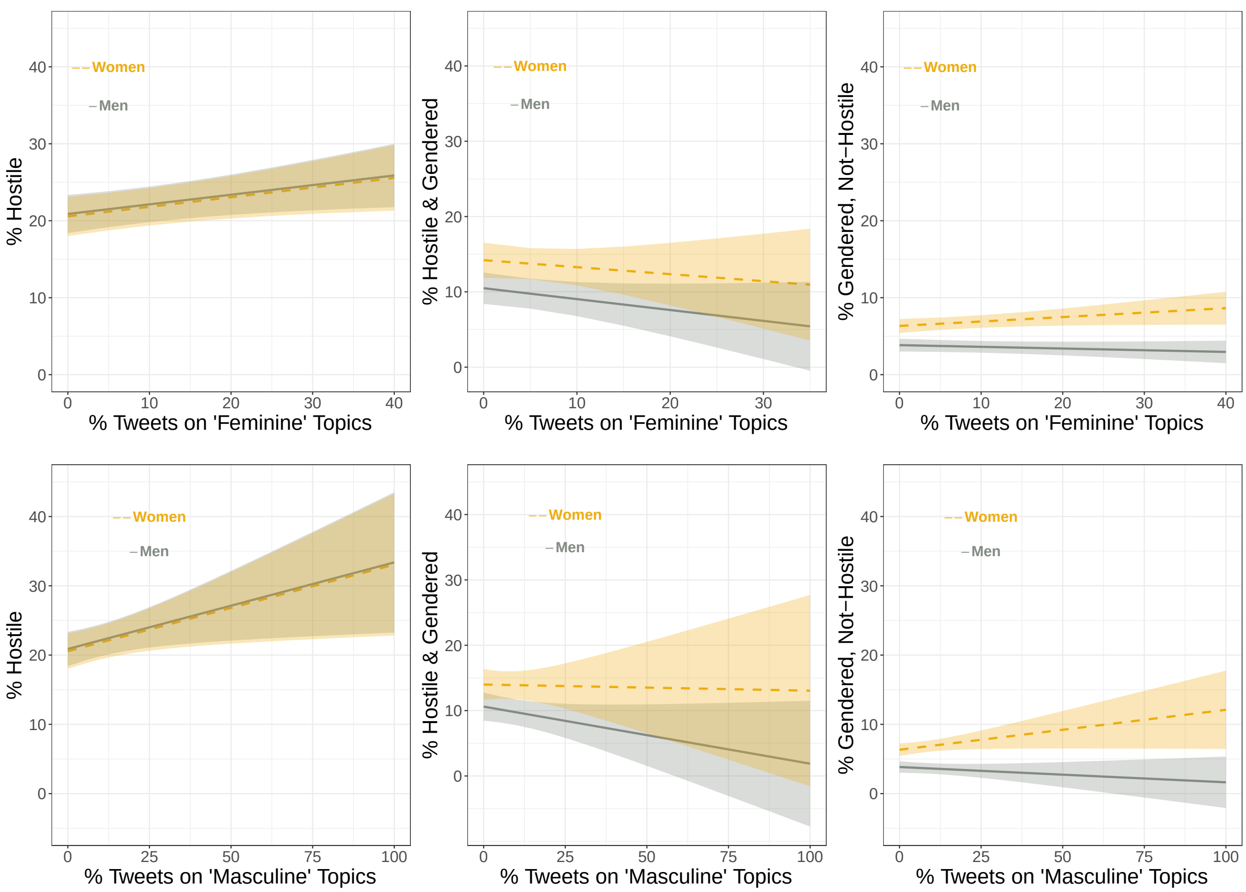

Table 22: 'Feminine' Issues Regression Results

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
	(1)	(2)	(3)
Woman Legislator	−3.014	−1.750	3.267**
	(2.104)	(2.630)	(1.357)
Visibility	3.560***	0.950**	−0.020
	(0.360)	(0.380)	(0.140)
Democrat	−6.270***	0.190	1.530***
	(0.667)	(0.710)	(0.250)
Count of Times Tweeted	−0.810***	−0.280	0.130
	(0.280)	(0.190)	(0.080)
Black or Latino	1.080	0.820	0.310
	(0.810)	(0.760)	(0.230)
Ideology	6.850**	0.570	0.040
	(3.070)	(2.670)	(0.800)
Positive Sentiment	−10.280***	−4.660***	0.460
	(1.930)	(1.030)	(0.410)
'Masculine' Issues	0.130*	−0.070	−0.000
	(0.070)	(0.050)	(0.040)
'Feminine' Issues	−0.270***	−0.140*	−0.010
	(0.100)	(0.080)	(0.050)
Opinion	5.410	−10.870**	1.820
	(7.800)	(5.330)	(2.330)
Asks for Donation	56.890	−15.560	−0.920
	(42.500)	(54.480)	(7.360)
No Policy Content	4.020	5.930**	0.010
	(4.330)	(2.590)	(1.140)
Factual Claim	9.410	−9.010*	−1.760
	(8.990)	(5.300)	(2.610)
Woman*Visibility	0.390	0.950**	0.100
	(0.340)	(0.400)	(0.200)
Woman*'Feminine' Issues	0.110	0.050	−0.180**
	(0.160)	(0.090)	(0.070)
Observations	1,839	1,839	1,839
R^2 (full model)	0.32	0.14	0.26
R^2 (proj model)	0.23	0.08	0.21
Adjusted R^2 (full model)	0.29	0.11	0.23
Adjusted R^2 (proj model)	0.21	0.05	0.18
Num. groups: state	48	48	48
Note:			*p<0.1; **p<0.05; ***p<0.01
			Standard errors in parentheses
			State fixed effects included in each model

Table 23: 'Masculine' Issues Regression Results

	Dependent variable:		
	% Hostile	% Hostile Gendered	% Gendered, Not-Hostile
	(1)	(2)	(3)
Woman Legislator	−2.270	−2.110	2.040*
	(2.020)	(2.770)	(1.080)
Visibility	3.550***	0.960**	−0.000
	(0.360)	(0.380)	(0.140)
Democrat	−6.270***	0.180	1.530***
	(0.670)	(0.720)	(0.260)
Count of Times Tweeted	−0.810***	−0.290	0.130*
	(0.280)	(0.190)	(0.070)
Black or Latino	1.050	0.810	0.360
	(0.820)	(0.760)	(0.230)
Ideology	6.930**	0.510	−0.100
	(3.090)	(2.700)	(0.820)
Positive Sentiment	−10.290***	−4.680***	0.480
	(1.940)	(1.030)	(0.390)
'Masculine' Issues	0.140	−0.090	−0.020
	(0.080)	(0.060)	(0.020)
'Feminine' Issues	−0.240**	−0.130	−0.070**
	(0.100)	(0.080)	(0.030)
Opinion	5.260	−10.800**	2.080
	(7.840)	(5.350)	(2.350)
Asks for Donation	57.970	−16.270	−2.720
	(42.310)	(55.030)	(7.440)
No Policy Content	3.990	5.950**	0.060
	(4.300)	(2.560)	(1.140)
Factual Claim	9.200	−8.920	−1.410
	(9.000)	(5.320)	(2.570)
Woman*Visibility	0.400	0.950**	0.080
	(0.340)	(0.400)	(0.200)
Woman*'Masculine' Issues	−0.050	0.080	0.080
	(0.090)	(0.090)	(0.070)
Observations	1,839	1,839	1,839
R^2 (full model)	0.32	0.14	0.25
R^2 (proj model)	0.23	0.08	0.21
Adjusted R^2 (full model)	0.29	0.11	0.23
Adjusted R^2 (proj model)	0.21	0.05	0.18
Num. groups: state	48	48	48
Note:			*p<0.1; **p<0.05; ***p<0.01 Standard errors in parentheses State fixed effects included in each model