

RESEARCH

Open Access



Gender differences in online abuse: the case of Dutch politicians

Isabelle van der Vegt^{1*}

Abstract

Online abuse and threats towards politicians have become a significant concern in the Netherlands, like in many other countries across the world. This paper analyses gender differences in abuse received by Dutch politicians on X (formerly Twitter), while taking into account the possible additional impact of ethnic minority status. All tweets directed at party leaders throughout the entire year of 2022 were collected. The effect of gender and ethnic minority status were estimated for six different linguistic measures of abuse, namely, toxicity, severe toxicity, identity attacks, profanity, insults, and threats. Contrary to expectations, tweets directed at male politicians scored higher on all forms of abuse. Significant interaction effects between gender and ethnic minority status were found for a number of abuse measures. Tweets directed at ethnic minority female politicians scored higher on severe toxicity, identity attacks, and profanity, compared to those directed at ethnic majority female politicians. Importantly, female ethnic minority politicians received the highest levels of threats compared to all groups. Given that online abuse and threats are reported to have a negative effect on political participation and retention, these results are particularly worrying.

Keywords Online abuse, Threats against politicians, Twitter data, Abuse detection

Introduction

In 2022, 1,125 reports of threats made to politicians were filed to the Dutch police (Openbaar Ministerie, 2023). This figure stands in stark contrast with the 588 reports filed in 2021, and 200 in 2015 (Jonker, 2022). A survey with members of the House of Representatives in the Netherlands showed that 12% reported to have kept an opinion to themselves due to fear of threats or intimidation (Jonker & Parre, 2022). While a large number of threats are made online, some threateners also physically approach politicians. In January 2022, the then-Dutch minister of Finance was approached at her home by a man shouting conspiratorial slogans while carrying a burning torch (NOS Nieuws, 2022). For women in

politics, the issue appears to be particularly salient. Such threats not only put individual women at risk, but they may also have a negative effect on women's participation in politics as a whole. The same Dutch Finance Minister announced her withdrawal from politics in 2023 due to the continued threats she endured (NOS nieuws, 2023). In 2019, eighteen female UK MPs stood down from their seat, with several of them announcing that the abuse they received was a factor in their decision (Scott, 2019). A study drawing on interviews with 101 Canadian politicians showed that online harassment did not necessarily stifle political ambition, but the hostile work environment reportedly affected the ability of these politicians to do their job and their willingness to stay in the job (Wagner, 2022).

Considering the increasing trend of online abuse and threats directed at female politicians and its far-reaching consequences, the current study empirically analyses gender differences in online abuse received by Dutch politicians on X (formerly Twitter), while taking into account the growing body of evidence that ethnic

*Correspondence:

Isabelle van der Vegt
i.w.j.vandervegt@uu.nl

¹ Department of Sociology, Utrecht University, Padualaan 14, 3584 CH Utrecht, The Netherlands



© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

minority women are particularly affected. This study offers a novel perspective from the Netherlands of online abuse of politicians, which has been predominantly studied in the UK (e.g., Esposito & Breeze, 2022; Gorrell et al., 2020; Ward & McLoughlin, 2020). In addition, it uniquely examines a broad timeline of tweets directed at Dutch politicians throughout the entire year of 2022, while considering several different measures of abuse.

Background

The following sections cover empirical evidence on the prevalence of online abuse and gender differences herein, followed by theoretical explanations of online abuse, gendered abuse of politicians, and the abuse of ethnic minority politicians.

Prevalence of (gendered) online abuse

Several survey studies have examined the levels of online harassment, abuse and threats experienced by politicians. Surveys of MPs from the early 2000s already showed that 10% of UK MPs experienced unwanted contact via social media in 2010 (James et al., 2016), to 60% of New Zealand MPs in 2014 (Every-Palmer et al., 2015). More recent surveys report a further rise in abusive online messages in recent years, with increases from 40% in 2013 to 70% in 2021 in Norway (Bjørge et al., 2022). Others have shown that 100% of MPs in the UK (Akhtar & Morrison, 2019) and Victorian (Phillips et al., 2023) parliament reported experiencing abuse via social media, in 2018 and 2021, respectively. Similar patterns have been reported when examining social media data. Gorrell et al. (2020) studied X data surrounding the 2019 parliamentary elections in the UK and observed that 4.46% of replies to MPs could be considered abusive, compared to 3.27% in the same period in 2017.

Case studies leveraging social media data have clearly demonstrated the severity of online abuse and threats experienced by female politicians. Examples include a study of abusive language and hate speech directed at Japanese female politicians (Fuchs & Schäfer, 2021), and sexual, physical, and psychological threats on X directed at pro-choice female politicians in Chile (Pérez-Arredondo & Graells-Garrido, 2021). However, empirical evidence on gender differences in online abuse and threats received by politicians is mixed. The majority of studies that have made direct comparisons between male and female politicians focus on the UK. Ward and McLoughlin (2020) report that male UK MPs received significantly more abusive tweets than female MPs in the UK between November 2016 and January 2017. However, female MPs received more hate speech than their male counterparts, partially

explained by the inclusion of gendered slurs as hate speech. Gorrell et al. (2020) similarly found that male politicians received more general and political abuse, while women received more sexist abuse. Southern and Harmer (2021) collected 117,802 tweets directed at UK MPs from 2 weeks in 2018 to assess gender differences. They found that female MPs received significantly more incivility, and were more likely to receive tweets that stereotyped them or questioned their position as representatives. Esposito and Breeze (2022) examined tweets for 3 weeks preceding the UK general elections in December 2019, and found no marked semantic differences (i.e., linguistic measures of emotion, references to appearance and intelligence) in tweets directed at male and female MPs. However, they did find large inter-individual differences, in that certain female MPs received a disproportionate amount of tweets containing words referring to appearance, sexual relations, emotions, and violence (Esposito & Breeze, 2022).

Two noteworthy studies have empirically examined the relationship between the gender of politicians and online messages within the Dutch context. Tromble and Koole (2020) studied tweets directed at Dutch politicians in October 2013, and found a marginal effect for gender in predicting tone (i.e., negative/positive language in tweets). Dutch, female, non-populist politicians received somewhat 'friendlier' tweets than their male counterparts. The authors also qualitatively examined negative tweets, and found very few tweets that could be considered sexist (e.g., mocking a female politician's appearance or using gendered slurs) and even fewer that could be considered a physical threat. A second investigation published in weekly opinion magazine *De Groene Amsterdammer* examined tweets directed at female Dutch politicians from October 2020 to February 2021 (Saris & Ven, 2021; Veerbeek, 2021). A subset of tweets were manually labelled as being hateful, threatening, or non-problematic. The hateful and threatening tweets were secondarily labelled for containing comments about (negative) female stereotypes, age, appearance, ethnicity, and religion. Thereafter, machine learning models were trained on this data and used to classify the remaining tweets in the dataset. The authors found that 10% of tweets directed at female politicians could be classified as hateful or threatening. Furthermore, it was found that female politicians were more frequently addressed with their first name and terms such as 'lady' and 'girl' ('vrouwje' and 'meid' in Dutch) than their male counterparts with male equivalent terms (e.g., 'boy', 'little boy'). It was also found that male politicians were addressed by their surname in 57% of the tweets, versus 27% for female politicians.

Explaining online abuse

Defining online abuse is seen as one of the key challenges by researchers in the field (Vidgen et al., 2019), with no single agreed-upon definition of the concept. One influential study separated online abuse into abuse directed against a group and that directed at an individual, referred to as hate speech and harassment, respectively (Waseem et al., 2017). Definitions of these concepts tend to differ between individual studies (Waseem et al., 2017). In some cases, researchers have relied on official definitions of hate speech (Vidgen et al., 2019), such as that put forward by the European Commission against Racism and Intolerance: *'the advocacy, promotion or incitement of the denigration, hatred or vilification of a person or group of persons, as well as any harassment, insult, negative stereotyping, stigmatization or threat of such person or persons and any justification of all these forms of expression on the ground of race colour, descent, national or ethnic origin, age, disability, language, religion or belief, sex, gender, gender identity, sexual orientation and other personal characteristics or status'* (ECRI, 2015).

Several frameworks have been put forward to explain abuse directed at politicians, as well as online abuse more generally. The online disinhibition effect, as described by Suler (2004), refers to the tendency of individuals to behave in ways online that they would not in person, due to the anonymity and minimization of authority on social media platforms, among other factors. This can lead to a sense of detachment from social norms, which may result in "acting out" online, for example in the form of harassment and threats directed at politicians (Trifiro et al., 2021; Tromble & Koole, 2020). Gorrell et al. (2020) proposed a heuristic framework specifically aimed at understanding the abuse politicians receive on X. The framework consists of four factors which may explain the distribution of abuse among politicians, specifically, (1) prominence: abuse focuses on individuals most in the public eye, (2) event surge: specific political or media events may lead to a surge in attention and/or hostility towards a politician, (3) engagement: abuse may be the result of a specific tweet by the politician themselves, and (4) identity: personal characteristics of a politician may affect the abuse they receive. Gorrell et al. (2020) found support for their model of abuse in their dataset of tweets from 2019, with the majority of abuse focusing on high profile politicians. Specific events and engagement, such as television appearances and opinionated tweets, were associated with 'spikes' in levels of abusive. Finally, Gorrell et al. (2020) found that certain personal characteristics, such as political stance and gender, led to differences in the level and forms of abuse received.

Explaining gendered online abuse

The distinction between (group-based) hate speech and individual harassment is less clear-cut when considering online abuse directed at female politicians (see also Vidgen et al., 2019). Even when female politicians are attacked based on their individual characteristics, the abuse may nevertheless be motivated by their identity as a woman. This phenomenon can be understood by leveraging theoretical frameworks from political science aimed at explaining violence directed at women in politics. Krook and Sanín (2020) suggest that violence against women in politics stems from misogyny, defined as "a system that polices and enforces patriarchal norms and expectations" (p. 742). The authors propose a three-factor model explaining violence against women in politics. First, the origins of this violence are structural, resulting from longstanding political theories associating men with the public and women with the private sphere. This, in turn, inspires and rationalizes hostility against female leaders, because they violate aforementioned female gender roles by operating in politics, a traditionally male-dominated domain. Second, Krook and Sanín (2020) put forward cultural violence as the means by which violence against women in politics is perpetrated. Specifically, cultural violence is the phenomenon of tolerating violence when it is directed at a particular group. Examples of this in the context of violence towards women in politics include sexist jokes and sexual objectification. Third, symbolic violence is seen as the (intended) outcome of violence against women in politics. This takes the form of male domination, aimed at "putting women who deviate from prescribed norms back in their place" (p. 743). Bardall et al. (2020) put forward a related theoretical framework for gendered political violence. They separate gendered motives, forms, and impacts in order to distinguish between political violence and *gendered* political violence. Gendered motives are present when political violence is committed in order to preserve politics as a male domain. Gendered forms constitute the way in which violence is perpetrated using gendered roles, such as sexualized language. Finally, gendered impacts refer to the possible gender differences in the meaning and consequences of political violence. Bardall et al. (2020) mention the way in which the media and society narrate an incident of political violence and possible differences in the extent to which political violence causes women versus men to retreat from politics, as examples of gendered impacts.

Online abuse of (ethnic) minorities

In their theoretical framework of gendered political violence, Bardall et al. (2020) further specified that this form of violence can serve to preserve political power of the

hegemonic male group. That is, politicians not belonging to the dominant cultural, ethnic, or religious group (e.g., gay men or ethnic minority women) may be similarly or even more strongly impacted than female politicians with majority characteristics. This idea lends support from previous research exploring the intersection between white and male hegemony, for example to explain the ideology of white supremacist movements (Berbrier, 2000). Belief in white heterosexual male supremacy has also been put forward to explain other related phenomena, such as anti-feminist movements defending perpetrators of domestic abuse against women (Blais, 2020).

Various investigations into the gender differences associated with online abuse directed at politicians have found that ethnic minority women were targeted disproportionately. As a consequence, Esposito and Breeze (2022) proposed that ‘gender bias is activated by some women more than others’ (p. 320). Analysis conducted by Amnesty International on tweets directed at female MPs in the UK in 2017 showed that 41% of all abusive tweets directed at women MPs were sent to Black, Asian and Ethnic minority (BAME) women MPs, even though almost 89% of the sample of MPs were white (Amnesty International UK, 2017). A disproportionate amount (32%) of abusive tweets were directed at Diane Abbott, the first black woman MP. The authors note that abuse sent to her often included threats of sexual violence and frequently mentioned her gender and race. In the Dutch investigation of Saris & Ven (2021) it was also found that female politicians of color or those with minority religious beliefs were impacted more strongly. For example, up to 30% of tweets directed at a Muslim second chamber member were classified as abusive, versus the 10% average for female politicians in the sample.

The current study

In explaining online abuse towards female politicians, both Krook and Sanín (2020) and Bardall et al. (2020) propose that this phenomenon, as with other forms of gendered political violence, may be rooted in the (conscious or unconscious) wish of abusers to ‘punish’ women who operate in a traditionally hegemonic male domain. Bardall et al. (2020) extend their theoretical framework to all politicians who do not fall within the dominant cultural, ethnic, or religious group. In the current paper, the focus is on differences in online abuse based on gender and ethnic minority status. Firstly, it is expected that Dutch female politicians receive higher levels of abuse on X than male politicians. Secondly, it is hypothesized that Dutch female politicians from a minority ethnic background receive higher levels of abuse than their ethnic majority counterparts. This study distinguishes between different types of abuse on X, namely, toxic tweets,

severely toxic tweets, profanity, insults, identity attacks, and threats. Taking into account the framework for online abuse directed at politicians (Gorrell et al., 2020), control variables include the political orientation of politicians, their prominence, and the extent to which they engage on social media.

Method

This study received approval from the ethics review board of the Faculty of Social and Behavioural Sciences at Utrecht University. Data and code for analysis are available via the Open Science Framework: <https://osf.io/vf6xt/>. In line with X terms of service, we can only make tweet IDs available publicly. The full dataset is available upon request.

Collecting X data

All tweets mentioning at least one of the Dutch political party leaders posted between 10 January 2022 (inauguration of the House of Representatives after the 2021 general elections) to 31 December 2022 were collected using the *academictwitter* R package (Barrie & Ho, 2022) and the now-defunct Twitter Academic API. In addition to party leaders ($n = 18$) independent second chamber members who separated from their party at some point during 2022 or 2021 ($n = 4$) were also included. Some party leaders have two X accounts because they also hold a ministerial post, such as the finance minister (i.e., @SigridKaag and @Minister_Fin). In those cases ($n = 3$) both accounts were included for data collection. In the API call, it was further specified that only tweets in Dutch should be collected and promoted Tweets should be excluded. In total, this resulted in a sample of 1,909,844 tweets directed at 22 politicians (see Appendix A), of which 9 (40.91%) are female.

Linguistic outcomes

The Google Perspective API was used to measure the linguistic outcomes (i.e., dependent variables). API calls were executed using the *peRerspective* R package (Votta, 2021). The tool uses machine learning models to predict ‘the perceived impact a comment may have on a conversation by evaluating that comment across a range of emotional concepts.’¹ At time of writing, measures of toxicity, severe toxicity, identity attacks, insults, profanity, and threats were available in Dutch, see Table 1 for a definition of each measure provided by Google. Training data for the machine learning models consists of ‘millions of comments from a variety of sources, including comments

¹ See https://developers.perspectiveapi.com/s/about-the-api-key-concepts?language=en_US.

Table 1 Definitions of linguistic measures from Google Perspective

Measure	Definition
1. Toxicity	A rude, disrespectful, or unreasonable comment that is likely to make people leave a discussion
2. Severe toxicity	A very hateful, aggressive, disrespectful comment or otherwise very likely to make a user leave a discussion or give up on sharing their perspective. This attribute is much less sensitive to more mild forms of toxicity, such as comments that include positive uses of curse words
3. Identity attack	Negative or hateful comments targeting someone because of their identity, including but not limited to race or ethnicity, religion, gender, nationality or citizenship, disability, age, or sexual orientation
4. Insult	Insulting, inflammatory, or negative comment towards a person or a group of people (not identity specific)
5. Profanity	Swear words, curse words, or other obscene or profane language
6. Threat	Describes an intention to inflict pain, injury, or violence against an individual or group

from online forums such as Wikipedia and The New York Times, across a range of languages.' Precise training data for the Dutch Perspective models are not reported. Each comment in the training data was scored on aforementioned measures by 3–10 annotators. Probability scores range between 0 and 1, where 1 represents a tweet in which the model predicts that all annotators would agree a tweet is toxic (or severely toxic, a threat, etc.). Model performance measured by the AUC on holdout test sets in Dutch range between 0.95 for insults to 1.0 for threats, where a score of 1 represents a model which predictions are 100% correct (see also Huang & Ling, 2005). Google recommends that a score of 0.7 or higher (i.e., 7 out of 10 annotators would agree that a tweet is toxic) can be used as a threshold in social science research for considering a tweet as toxic (or severely toxic, a threat, etc.). Before obtaining the linguistic measures via the Google Perspective API, no text preprocessing (e.g., removing mentions, URLs) was conducted because the Perspective API is developed for use on social media data. Table 2 shows examples of tweets with all six measures.

Independent, moderator and control variables

Gender of politicians (male or female) is entered into the regression models as independent variable. Ethnic minority status (0 = ethnic majority, 1 = ethnic minority) is included as a control variable in Model 1 and additionally as moderator for gender in Model 2. In both models, further control variables include prominence, political position, and engagement. Prominence is measured by the X follower count of each politician on 31 December 2022. Online engagement of politicians is measured through the number of tweets sent between January 10 and 31 December 2022. Both variables are log-transformed to account for skewness in the data.

Political position is based on the visualization of the Dutch political landscape along two axes developed by 'Kieskompas'. The tool is developed by independent political researchers who make use of the official viewpoints

of political parties in their manifestos (for the Dutch general elections of 2021) in order to characterize them.² Originally, political position was included in the form of two variables. First, the economic stance of the party a politician belongs to, including the categories left-wing, center, and right-wing. The second refers to a party's cultural stance, being either conservative or progressive. However, when conducting diagnostic checks for the regression analysis below, we identified high multicollinearity (based on the Variance Inflation Factor) between both variables. Therefore, analyses were continued with economic stance alone (left-wing, center or right-wing) as measure of political position. For independent second chamber members who left their party, the political position of the party they were previously a member of was used. See Appendix A for an overview and further justification of the political position variables.

Analytical approach

Two general linear models are constructed for each of the six linguistic outcomes. Model 1 tests for a relationship between gender and the linguistic outcome, while controlling for ethnic minority status, follower count, political position, and number of tweets. In Model 2 for each outcome, the interaction between gender and ethnic minority status is entered. Importantly, analyses are performed at the tweet-level. That is, the dependent variable consists of the linguistic scores per tweet, and these are not aggregated per politician. To account for nesting in the data (i.e., tweets are nested within politicians), cluster-robust standard errors are computed using the *miceadds* R package (Robitzsch & Grund, 2023). Confidence intervals reported are based on the clustered standard errors. Due to skewness in the data, all linguistic outcomes are log-transformed. In order to calculate the

² See <https://www.kieskompas.nl/nl/de-kieskompas-methode/> for a description of the method (in Dutch).

Table 2 Examples of tweets and linguistic outcome measures

Dutch (original)	English	Gender	Toxicity	Severe toxicity	Identity attack	Insult	Profanity	Threat
hang jezelf aub op, vieze kanker jood. Je verneukt heel nederland, hoop dat je word doorzeefd met kogels	Please hang yourself, dirty "cancer" jew. You're ruining the Netherlands, hope you get riddled with bullets	M	0.96	0.82	0.84	0.93	0.95	0.86
Fuck you smerige kankerhoer, jij staat in je uppie. Helemaal alleen. Je bent een vieze landverrader. Een nsber. Een harteloos kankerwijf!!!	Fuck you dirty "cancer" whore, you're all alone. Completely alone. You are a dirty traitor to this country. An NSBer**.	F	0.95	0.82	0.40	0.93	0.96	0.28
Je bent een gore kanker nazi, een loser die nooit iets zal bereiken, de letterlijke reetkanker van de samenleving	You are a dirty "cancer" nazi, a loser who will never achieve anything, you are the "anal cancer" of society	M	0.95	0.82	0.59	0.94	0.94	0.19
Smerige Kuthoer,, kogel in je kop moet je krijgen	Dirty "cunt" "cancer" whore,, you should get a bullet in the head	F	0.93	0.82	0.11	0.87	0.89	0.85
ook zoiets,,kut kanker makro,,kom maar op,,9 mm ligt klaar schiet je zo overhoop,,en daarna vermoord ik je hele fam tot in de 7e gen.,	Again,, moroccan,, come on,, 9 mm is ready to shoot you to pieces,, and after that I kill you and your whole family to the 7th gen.,	M	0.93	0.82	0.32	0.80	0.90	0.90
Rot jij maar lekker op naar je vaderland met je klote islam	Piss off to your homecountry with your damn islam	M	0.92	0.82	0.85	0.80	0.76	0.35
Domme muts ben je. Serieus wat een achtelijk schijtwijs. Geen andere woorsen voor deze domme domme domme opmerking van je	Stupid bimbo you are. Seriously what a retarded "crap" woman. No other words for this stupid stupid stupid remark of yours	F	0.89	0.66	0.06	0.88	0.84	0.02
Ik doe straks een strik om je nek heen	I'm going to put a noose around your neck later	M	0.83	0.71	0.03	0.63	0.42	0.79

English translations by the author

* Cancer (and other diseases) is a common swear word in Dutch

** NSB was the Dutch national socialist movement in the 1930-1940s

Table 3 Descriptive statistics by gender

	Female			Male		
	M	SD	%	M	SD	%
Follower count	129,987	66,840		565,585*	496,059	
Number of tweets	4,522*	3,665		2,001	1,514	
Linguistic outcomes						
Toxicity	0.16	0.18	1.31	0.19*	0.19	1.81*
Severe toxicity	0.05	0.12	0.16	0.07*	0.14	0.24*
Identity attack	0.04	0.09	0.05	0.05*	0.11	0.11*
Insult	0.15	0.20	1.66	0.19*	0.22	2.45*
Profanity	0.10	0.15	1.47	0.12*	0.16	1.59*
Threat	0.02	0.05	0.01	0.02*	0.06	0.02*

* $p < 0.001$

percent change in linguistic outcomes, regression coefficients b can be interpreted as follows: $100 * (e^b - 1)$.

Statistically significant interactions in Model 2 are further examined through pairwise comparisons of estimates for all possible combinations of gender and ethnic minority status. The *marginaleffects* R package (Arel-Bundock, 2023) is used to obtain the estimates for tweets directed at ethnic majority female, ethnic minority female, ethnic majority male, and ethnic minority male politicians. Control variables are set to their mean for continuous variables (i.e., number of followers and number of tweets) or to the reference levels used in the original models for categorical variables (i.e., political position). For each comparison, confidence intervals based on cluster-robust standard errors are computed.

Finally, a robustness check is performed. To verify whether regression results are not driven by the scores of an individual politician, the procedure above is repeated for each linguistic outcome, with the tweets of individual politicians who scored the highest on the respective measure removed from the data.

Results

Descriptive results

Table 3 shows the descriptive statistics separated by gender. Univariate t-tests show that male politicians have significantly more followers, whereas female politicians on average have a higher level of engagement, measured by the number of tweets. Comparing mean scores for all linguistic outcomes, male politicians score significantly higher than female politicians on all measures. Nevertheless, the average level of abuse can be considered relatively low, given that the maximum score is 1. The table also shows the percentage of tweets by gender that can be classified as toxic, severely toxic etc., (i.e., the tweet received a score ≥ 0.7 for the measure). Again, a small

percentage of tweets can be considered abusive, ranging between 0.01% of tweets directed at female politicians classified as a threat, to 2.45% of tweets directed at male politicians classified as an insult. Chi-square tests indicated that there were significant associations with gender for each measure, with male politicians receiving more abusive tweets across all measures. In Appendix B, linguistic outcomes per individual politician are shown.

Regression models

Results for the regression analysis are shown in Table 4. Contrary to Hypothesis 1, tweets directed at female politicians score lower (as indicated by the negative regression coefficient) on (1) *toxicity*, (2) *severe toxicity*, (3) *identity attacks*, (4) *insults*, (5) *profanity*, and (6) *threats* compared to their male counterparts when controlling for number of followers, ethnic minority status, political position, and number of tweets. Recall that the dependent variables were log-transformed. This means, for example, that female politicians score $100 * (e^{-0.25} - 1) = 22.1\%$ lower on toxicity than male politicians, and $100 * (e^{-0.28} - 1) = 24.4\%$ lower on severe toxicity, and so on. Significant effects for ethnic minority status can be observed in Model 1 for the measures of severe toxicity and profanity, with ethnic minority politicians receiving 29.7% and 20.92% higher scores on these measures, respectively.

When introducing the interaction between gender and ethnic minority status in Model 2 for each measure, a significant positive interaction for (2) *severe toxicity*, (3) *identity attacks*, (5) *profanity*, and (6) *threats* can be observed. No significant interaction effect between gender and ethnic minority status was found for (1) *toxicity* nor for (4) *insults*. The models explain only a small amount of variance for each linguistic outcome, ranging

Table 4 Regression results for linguistic outcomes per model

Model	Linguistic outcome											
	1. Toxicity		2. Severe toxicity		3. Identity attack		4. Insult		5. Profanity		6. Threat	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Gender (ref=M)	-0.25** [-0.41, -0.09]	-0.27*** [-0.42, -0.12]	-0.28*** [-0.44, -0.13]	-0.33*** [-0.45, -0.20]	-0.33*** [-0.47, -0.19]	-0.38*** [-0.49, -0.27]	-0.22* [-0.39, -0.05]	-0.24** [-0.40, -0.08]	-0.17** [-0.29, -0.06]	-0.20*** [-0.30, -0.10]	-0.04* [-0.08, 0.00]	-0.06*** [-0.09, -0.02]
Ethnic minority (ref=maj)	0.16 [-0.06, 0.37]	0.08 [-0.21, 0.36]	0.26* [0.03, 0.50]	0.08 [-0.11, 0.28]	0.14 [-0.13, 0.41]	-0.06 [-0.23, 0.10]	0.18 [-0.05, 0.40]	0.10 [-0.21, 0.40]	0.19* [0.04, 0.35]	0.10 [-0.07, 0.27]	0.07 [0.00, 0.15]	0.01 [-0.04, 0.05]
Political position (ref=left)	-0.15 [-0.37, 0.07]	-0.18 [-0.42, 0.06]	-0.22* [-0.43, -0.01]	-0.30** [-0.52, -0.08]	-0.30* [-0.54, -0.07]	-0.39** [-0.63, -0.15]	-0.14 [-0.37, 0.09]	-0.17 [-0.42, 0.07]	-0.10 [-0.26, 0.05]	-0.14 [-0.31, 0.02]	-0.05 [-0.12, 0.01]	-0.08* [-0.15, -0.02]
<i>Center</i>												
<i>Right-wing</i>	0.28* [0.03, 0.52]	0.25 [0.00, 0.50]	0.34** [0.09, 0.59]	0.28* [0.03, 0.52]	0.20 [-0.06, 0.45]	0.13 [-0.13, 0.38]	0.27* [0.01, 0.52]	0.24 [-0.02, 0.50]	0.31*** [0.13, 0.49]	0.28** [0.10, 0.45]	0.14*** [0.07, 0.21]	0.11*** [0.05, 0.18]
Tweets (log)	-0.24*** [-0.30, -0.18]	-0.24*** [-0.31, -0.18]	-0.31*** [-0.39, -0.24]	-0.32*** [-0.39, -0.26]	-0.30*** [-0.38, -0.23]	-0.31*** [-0.38, -0.24]	-0.24*** [-0.30, -0.18]	-0.24*** [-0.31, -0.18]	-0.20*** [-0.25, -0.15]	-0.21*** [-0.25, -0.16]	-0.11*** [-0.13, -0.09]	-0.11*** [-0.13, -0.09]
Followers (log)	-0.06 [-0.14, 0.02]	-0.05 [-0.14, 0.03]	-0.05 [-0.14, 0.03]	-0.03 [-0.12, 0.05]	-0.03 [-0.11, 0.05]	-0.01 [-0.09, 0.06]	-0.05 [-0.14, 0.04]	-0.04 [-0.13, 0.05]	-0.06* [-0.12, 0.00]	-0.05 [-0.11, 0.00]	0.01 [-0.02, 0.04]	0.02 [-0.01, 0.04]
Gender* ethnicity	0.21	0.21	0.47***	0.47***	0.51***	0.51***	0.20	0.20	0.25**	0.25**	0.17***	0.17***
Pseudo R ²	0.032	0.032	0.045	0.047	0.038	0.040	0.031	0.031	0.031	0.032	0.025	0.026

Unstandardized regression coefficients *b* and confidence intervals based on cluster-robust standard errors. Linguistic outcomes were log-transformed

* *p* < 0.05; ** *p* < 0.01; *** *p* < 0.001

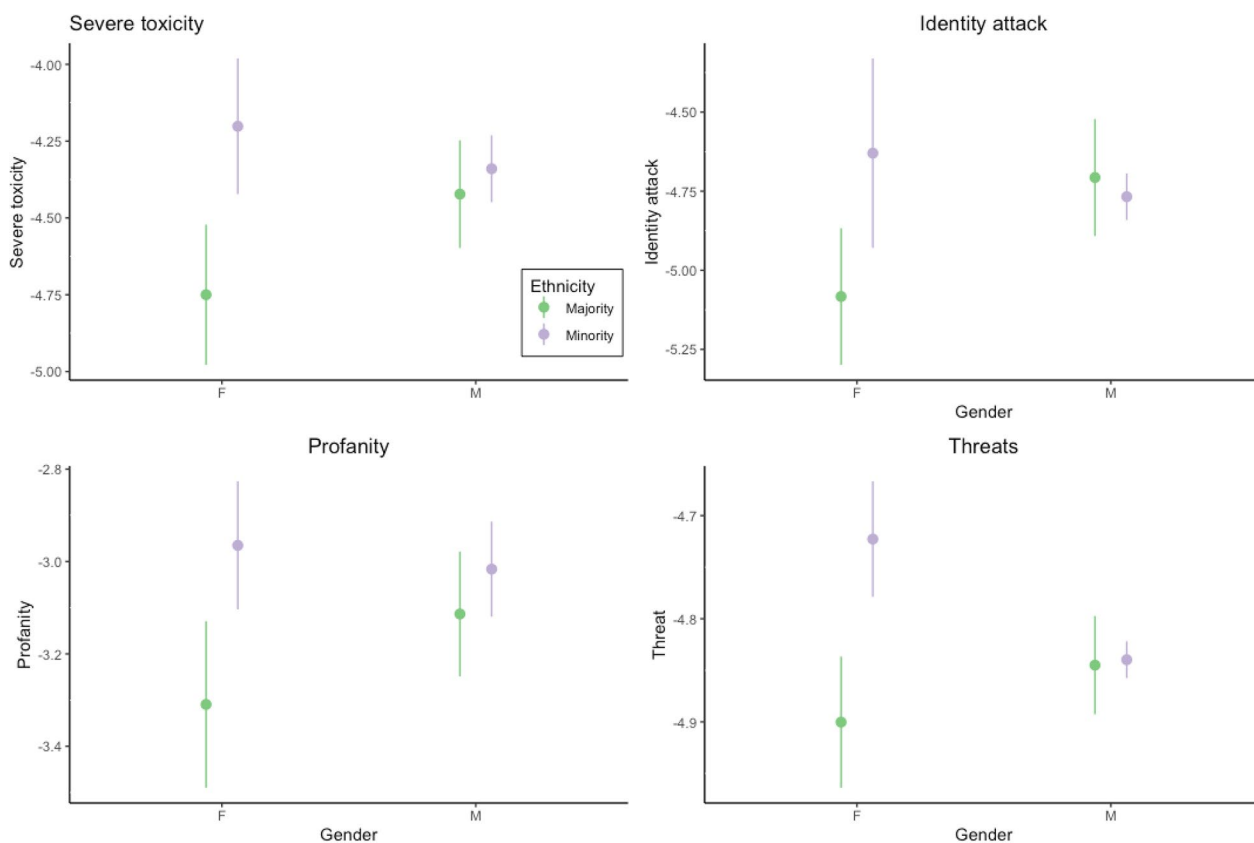


Fig. 1 Interaction plots. Estimates are based on a model where number of followers and number of tweets are set to their mean, and where the original reference levels are used for political position (i.e., left-wing). Lines around point estimates represent cluster-robust standard errors

from 2.6% of variance for threats to 4.7% for severe toxicity (for Model 2 in both cases).

Interaction plots for (2) *severe toxicity*, (3) *identity attacks*, (5) *profanity*, and (6) *threats* are shown in Fig. 1, with pairwise comparisons between each group reported in Appendix C. Examining the plot for (2) *severe toxicity*, tweets directed at ethnic minority female politicians score higher than tweets directed at majority female politicians ($p < 0.001$, see Appendix C) and those directed at majority male politicians ($p < 0.05$) but not significantly different from tweets directed at ethnic minority male politicians. For the measure of (3) *identity attacks*, a significant difference between tweets directed at ethnic majority and minority female politicians is found ($p < 0.01$), but tweets for ethnic minority female politicians do not score significantly different from tweets mentioning ethnic majority and minority male politicians. For (5) *profanity*, a similar pattern to (2) *severe toxicity* emerges. For the measure of (6) *threats*, tweets directed at minority ethnic female politicians score the highest. The estimate for this group is significantly higher than that for tweets directed at female ethnic majority, male ethnic majority, and male ethnic minority politicians (all $p < 0.001$).

The results for the robustness check are shown in Appendix D. For each linguistic outcome, the tweets directed at the politician that scored the highest on that measure (as shown in Appendix B) were removed from the data and both models were constructed. Results show similar effects in terms of directionality and statistical significance as those shown in Table 4.

Discussion

The results of this study paint a mixed picture of the abuse Dutch politicians receive online via X. When tweets are directed at female politicians the levels of all forms of abuse are lower compared to tweets directed at their male counterparts. Ethnic minority status significantly increases the level of severe toxicity, identity attacks, profanity, and threats female politicians receive. For the measures of severe toxicity, identity attacks, and profanity, scores for tweets directed at ethnic minority female politicians are heightened to the level of that of tweets directed at (ethnic minority) male politicians. Importantly, for the measure of threats only, tweets directed at ethnic minority female politicians score the highest when compared to tweets directed at all other

groups. In short, no support for the hypothesis that Dutch female politicians receive higher levels of abuse than male politicians on X is found. The results offer partial support for the hypothesis stating that ethnic minority status increases the level of abuse female politicians receive. This holds for the measures of severe toxicity, identity attacks, profanity, and threats only.

Perhaps the most striking result of this study is that tweets directed at ethnic minority women in Dutch politics contain the highest levels of threats, arguably the most severe form of abuse. While tweets directed at ethnic minority female politicians do not score higher than those directed at male politicians on the other forms of abuse measured in this study, it is clear that these tweets are generally worse than those directed at ethnic majority female politicians. Looking back at the theoretical frameworks for gendered political violence (Bardall et al., 2020; Krook & Sanín, 2020), a few remarks can be made. Expectations were derived from the idea that online abuse of female politicians can be explained by the motivation to preserve politics as hegemonic male domain. Some support for this explanation was found, albeit not in the way that was expected. Tweets directed at female politicians did not score significantly higher on all of measures of abuse. However, the expected pattern of results was obtained for tweets directed at female politicians with an ethnic minority background for some of the outcome measures. Therefore, the possibility still exists that abuse (and threats in particular) are used to 'punish' ethnic minority women for their participation in the political arena, in line with preserving politics as a hegemonic male domain. Although significant differences were found based on gender in combination with ethnic minority status, it cannot be concluded that the abuse expressed in this sample of tweets is *explicitly* motivated by these characteristics. Based on the probability scores for abuse, we cannot establish whether gender and ethnicity attributes of politicians are, for example, mentioned in a derogatory way. The measure of 'identity attacks' does concern abuse based on someone's identity, but the probability score does not show whether the attack is based on gender, ethnicity, religion, sexual orientation, or any other characteristic. In future research, more in-depth analysis of the content of tweets will be necessary.

In short, results appear to be more in line with previous work suggesting that male politicians receive higher levels of general abuse online (Gorrell et al., 2020; Ward & McLoughlin, 2020). At the same time, some of the results align with earlier investigations showing that female ethnic minority politicians are disproportionately affected by online abuse (Amnesty International UK, 2017; Esposito & Breeze, 2022; Saris & Ven, 2021). However,

direct comparisons of current results with previous studies are difficult since linguistic measures and analytical approaches were different. For instance, average levels of abuse directed at female politicians in this study vary between 0.01% for threatening and 1.66% for insulting tweets, which can be considered sufficiently lower than the 10% of tweets classified as abusive or threatening by Saris & Ven (2021) in their Dutch sample. However, due to differences in definitions of abuse and data collection procedures, these findings cannot be directly compared, and thus it would not be appropriate to say that levels of abuse directed at female politicians have decreased. Finally, it should be noted that this study only examined tweets directed at party leaders and independent second chamber members. That is, the data does not cover the full landscape of Dutch politicians, and thus the findings may not generalize to this wider population or to the new second chamber members installed after the 2023 elections. Nevertheless, we argue that the current study provides the most comprehensive picture of online abuse directed at Dutch politicians thus far, given that we have examined a full year of tweets and several different measures of abuse.

The theoretical frameworks leveraged in this study to explain abuse directed at female politicians cannot be utilized to understand the high(er) levels of abuse directed at male politicians in the data. Alternative explanations can perhaps be found in the Gorrell et al., (2020) four factor framework, which considers prominence, event surge, engagement, and identity as main drivers of online abuse of politicians. The possibility exists that male politicians receive more abuse than their female counterparts due to their prominence or engagement with the public, and that these factors were not adequately controlled for with the measures of follower count and number of tweets, respectively. Furthermore, the current study did not take into account so-called 'event surge', that is, specific political or media events that may have led to high volumes of abuse. Another possible explanation may lie in the linguistic measures utilized in this study. For instance, the measure of 'identity attacks' is defined by Google Perspective as any attack based on someone's gender, ethnicity, race, or other personal characteristics. As a consequence, one might expect ethnic minority politicians and women in particular to receive a high level of such messages. Surprisingly, tweets directed at minority ethnic female politicians do not score significantly different on this measure from tweets directed at majority and minority ethnic male politicians. One explanation for this could be that in some cases tweets that mentioned a politician included (racist) abuse or threats directed at another individual or group and not at the politician him/herself. This seemed to be particularly the case for

tweets directed at the leader of a political party with an anti-Islam agenda (Witteveen, 2017). As mentioned above, the tools used for linguistic measurement in this study do not provide specific information on the target of abuse or the personal characteristics of a target motivating an identity attack. In addition to future qualitative examination of tweets to determine explicit motivations and targets, the application of more sophisticated target detection algorithms could offer some further insight on this matter.

Some additional imitations to this study need to be considered. First, the data collection procedure only retrieved tweets that directly mentioned a politician by their username on X. Therefore, tweets in which a politician was simply named or implicitly referred to were not included in the analysis. The possibility exists that such tweets are more or less abusive than what was captured in this dataset. Second, on October 27, 2022, Elon Musk became the new owner and CEO of Twitter (later rebranded to X). Various reports have stated that content moderation has suffered or at least changed as a result of this (Paul & Dang, 2022), as Musk worked to instate his “free speech” agenda on the platform (Zakrzewski et al., 2022). These events may have had an effect on the results observed in this paper. At the same time, the possibility also exists that the levels of abuse posted on X were underestimated due to content moderation, with us failing to capture tweets that were removed from the platform as a result of this. All in all, these factors urge us to interpret (the magnitude and direction of) the results presented with some caution. Future research may also analyze the data from a timeseries perspective, in order to assess the potential effect of the Musk acquisition or specific political events. Additionally, free access to the X/Twitter Academic API has been curtailed since early 2023 (Calma, 2023), complicating future replication efforts or additional data collection. Finally, a third-party tool such as the Google Perspective API does not allow for fully transparent linguistic measurements. That is, the documentation does not specify the precise training data used for the Dutch models and possible bias as a result of this. Furthermore, since the tool merely produces a probability score for each measure of abuse, it was not possible to scrutinize the precise features the models use to predict the levels of toxicity, identity attacks, and other measures.

Conclusion

Politicians around the world, including in the Netherlands, have reported a significant and rising trend of online abuse. This paper set out to add to growing field of empirical inquiry into gender differences in abuse received by politicians online, while taking into account

the additional impact of ethnic minority status. Six different forms of abuse were measured in a full year of tweets on X. Contrary to the expectations set out at the beginning of this study, tweets directed at male politicians contain higher levels of most forms of abuse. At the same time, tweets directed at female ethnic minority politicians did score higher than those directed at female ethnic majority politicians on some forms of abuse. Notably, tweets directed at female ethnic minority politicians scored the highest in terms of threats, the most severe form of abuse measured in this study. While it cannot be said how many of such threats may lead to actual physical violence, the mere prevalence and nature of abuse in the online domain are reportedly already enough to have women leave or refrain from participating in the political arena. Therefore, it is important that we continue to invest in understanding and combatting this phenomenon.

Appendix A: Political position measures

In order to determine political position of politicians, a visualization of the Dutch political landscape for the 2021 general elections developed by the independent organization ‘Kieskompas’ was used (see <https://tweedekamer2021.kieskompas.nl/nl/>). In the visualization, each political party is placed along the horizontal axis of socio-economic viewpoints (left to right-wing) and the vertical axis of cultural (progressive to conservative) viewpoints. All political parties left of the mid-point of the horizontal axis were coded as left-wing in terms of socio-economic viewpoints, all parties to the right of the mid-point as right-wing, and those who intersected with the mid-point of the horizontal axis were coded as ‘center’.

The vertical axis of the Kieskompas represents the cultural stance of political parties. All political parties above the mid-point of the vertical axis were considered as progressive, and those below as conservative. No political parties crossed the mid-point of the vertical axis for cultural viewpoints. However, model diagnostic checks indicated this variable was highly correlated with the economic stance. Therefore, it was not included in any of the final regression models. The below table shows the coding for each politician.

Party leader	Political party	Economic	Cultural
Mark Rutte	VVD	Right-wing	Conservative
Sigrid Kaag	D66	Center	Progressive
Jesse Klaver	GroenLinks	Left-wing	Progressive
Nilüfer Gundogan	Independent (former VOLT)	Center	Progressive

Party leader	Political party	Economic	Cultural	Party leader	Political party	Economic	Cultural
Wopke Hoekstra	CDA	Center	Conservative	Attje Kuiken	PvdA (from 22 April 2022)	Left-wing	Progressive
Sylvana Simons	BIJ1	Left-wing	Progressive	Esther Ouwehand	Partij voor de Dieren	Left-wing	Progressive
Laurens Dassen	VOLT	Center	Progressive	Wybren van Haga	Independent (former FvD)	Right-wing	Conservative
Liane den Haan	Independent (former 50Plus)	Left-wing	Progressive	Gertjan Segers	ChristenUnie	Left-wing	Progressive
Farid Azarkan	DENK	Left-wing	Progressive	Caroline van der Plas	BBB	Right-wing	Conservative
Thierry Baudet	Forum voor Democratie	Right-wing	Conservative	Lilian Marijnissen	SP	Left-wing	Progressive
Geert Wilders	PVV	Left-wing	Conservative	Kees van der Staaij	SGP	Right-wing	Conservative
Joost Eerdmans	JA21	Right-wing	Conservative	Pieter Omtzigt	Independent (former CDA)	Center	Conservative
Lilianne Ploumen	PvdA (until 21 April 2022)	Left-wing	Progressive	Martin van Rooijen	50Plus	Left-wing	Conservative

Appendix B: Linguistic outcomes per politician

Politician	Gender	Minority ethnic	Toxic	Severely toxic	Identity attack	Insult	Profanity	Threat
Mark Rutte	M	N	3.58	0.49	0.06	4.64	2.90	0.03
Sigrid Kaag	F	N	3.41	0.53	0.12	3.91	3.28	0.02
Jesse Klaver	M	Y	3.21	0.38	0.17	4.32	2.38	0.01
Nilüfer Gundogan	F	Y	2.86	0.49	0.22	3.75	2.05	0.01
Wopke Hoekstra	M	N	2.25	0.22	0.05	3.15	1.84	0.01
Sylvana Simons	F	Y	2.17	0.33	0.14	2.63	1.88	0.00
Laurens Dassen	M	N	2.14	0.21	0.06	2.56	2.05	0.01
Attje Kuiken	F	N	1.95	0.11	0.01	2.84	1.75	0.00
Liane den Haan	F	N	1.90	0.12	0.01	2.00	4.30	0.00
Farid Azarkan	M	Y	1.61	0.50	0.62	1.45	1.40	0.01
Thierry Baudet	M	N	1.58	0.12	0.02	2.56	1.57	0.01
Geert Wilders	M	N	1.57	0.31	0.27	1.79	1.43	0.04
Joost Eerdmans	M	N	1.27	0.02	0.02	2.04	1.31	0.00
Lilianne Ploumen	F	N	1.22	0.06	0.02	1.56	1.18	0.00
Esther Ouwehand	F	N	1.22	0.11	0.03	1.62	1.54	0.01
Wybren van Haga	M	N	0.86	0.06	0.02	1.54	0.99	0.00
Gertjan Segers	M	N	0.79	0.09	0.07	1.11	0.67	0.00
Caroline van der Plas	F	N	0.32	0.03	0.01	0.49	0.64	0.00
Lilian Marijnissen	F	N	0.31	0.02	0.02	0.46	0.68	0.00
Kees van der Staaij	M	N	0.27	0.06	0.02	0.56	0.56	0.00
Pieter Omtzigt	M	N	0.27	0.02	0.01	0.43	0.50	0.00
Martin van Rooijen	M	N	0.24	0.02	0.01	0.36	0.32	0.00

Number of tweets out of 100 tweets directed at a politician that can be classified as each form of abuse are shown. Highest score per measure in bold. In the table above, linguistic outcomes are separated by individual politicians, showing the number of tweets out of 100 directed at each politician that is classified as toxic, severely toxic, an identity attack, an insult, profanity, or a threat (i.e., ≥ 0.7). Mark Rutte, the male prime minister, receives the highest proportion of toxic tweets, with 3.58 out of 100 tweets directed at him being classified as such. Sigrid Kaag, the female vice prime minister and finance minister, receives the second highest proportion of toxic tweets (3.41/100), but the highest proportion of *severely* toxic tweets (0.53/100). Identity attacks are less common, but here Farid Azarkan (0.62/100), Nilüfer Gundogan (0.22/100), and Sylvana Simons (0.14/100), all with an ethnic minority background, receive the most. Profanity is most commonly found in tweets directed at Liane Den Haan (4.30/100), Sigrid Kaag (3.28/100) and Mark Rutte (2.90/100). The high score for Liane Den Haan is somewhat misleading, since on 13/7/2022 she posted a tweet condemning the use of #kutland ('shit country'). Many tweets in reply to this included the expletive 'kut', hence the high score on profanity for this politician. Finally, although threats are the least common, Geert Wilders receives the most (0.04/100), followed by Mark Rutte (0.03/100) and Sigrid Kaag (0.02/100). Here, it must be noted that Geert Wilders is the leader of a political party that is regarded as xenophobic and anti-Islam (Witteveen, 2017). Manual inspection of tweets directed at Geert Wilders show that they frequently contain threats directed at others (mainly Muslims) rather than at the politician himself.

Appendix C: Pairwise comparisons (contrasts)

A positive estimated difference suggests that the first group in the comparison scores higher than the second group, and vice versa. Confidence intervals are computed based on cluster-robust standard errors, with: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Contrast	Est. diff	CI _{low}	CI _{high}
Severe toxicity			
Majority male–majority female	0.327***	0.201	0.454
Majority male–minority male	-0.083	-0.277	0.112
Majority male–minority female	-0.221*	-0.436	-0.006
Majority female–minority male	-0.410***	-0.659	-0.161
Majority female–minority female	-0.548***	-0.756	-0.341
Minority male–minority female	-0.138	-0.389	0.113
Identity attack			
Majority male–majority female	0.376***	0.267	0.486
Majority male–minority male	0.061	-0.104	0.225
Majority male–minority female	-0.077	-0.398	0.244
Majority female–minority male	-0.316**	-0.520	-0.111
Majority female–minority female	-0.453**	-0.763	-0.143
Minority male–minority female	-0.138	-0.439	0.164
Profanity			
Majority male–majority female	0.196***	0.096	0.295
Majority male–minority male	-0.097	-0.268	0.074
Majority male–minority female	-0.148*	-0.281	-0.016
Majority female–minority male	-0.293**	-0.507	-0.079
Majority female–minority female	-0.344***	-0.479	-0.209
Minority male–minority female	-0.051	-0.235	0.132
Threat			
Majority male–majority female	0.055***	0.025	0.086
Majority male–minority male	-0.005	-0.052	0.042
Majority male–minority female	-0.122***	-0.167	-0.077
Majority female–minority male	-0.061*	-0.121	0.000
Majority female–minority female	-0.177***	-0.220	-0.135
Minority male–minority female	-0.117***	-0.175	-0.059

Appendix D: Robustness check

Linguistic outcome												
Model	1. Toxicity		2. Severe toxicity		3. Identity attack		4. Insult		5. Profanity		6. Threat	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
Gender (ref=M)	-0.25**	-0.26**	-0.30**	-0.37***	-0.34***	-0.38***	-0.22*	-0.23*	-0.17**	-0.20***	-0.04*	-0.06***
	[-0.42, -0.08]	[-0.43, -0.09]	[-0.49, -0.12]	[-0.53, -0.21]	[-0.48, -0.21]	[-0.49, -0.27]	[-0.41, -0.04]	[-0.41, -0.05]	[-0.28, -0.06]	[-0.29, -0.10]	[-0.08, 0.00]	[-0.09, -0.03]
Ethnic minority (ref=maj.)	0.19	0.14	0.26*	0.07	0.17	-0.04	0.20	0.16	0.20*	0.10	0.07	0.00
	-0.03, 0.41]	[-0.18, 0.46]	[0.03, 0.50]	[-0.12, 0.26]	[-0.12, 0.47]	[-0.22, 0.14]	[-0.03, 0.44]	[-0.19, 0.51]	[0.04, 0.36]	[-0.07, 0.27]	[-0.01, 0.15]	[-0.05, 0.05]
Political position (ref=left) <i>Center</i>	-0.10	-0.12	-0.25*	-0.36**	-0.30*	-0.39**	-0.09	-0.11	-0.09	-0.14	-0.05	-0.09**
	[-0.33, 0.12]	[-0.38, 0.13]	[-0.48, -0.01]	[-0.62, -0.10]	[-0.53, -0.07]	[-0.63, -0.14]	[-0.32, 0.14]	[-0.38, 0.15]	[-0.26, 0.07]	[-0.31, 0.03]	[-0.12, 0.02]	[-0.15, -0.02]
<i>Right-wing</i>	0.27*	0.26	0.32**	0.24*	0.22	0.14	0.26	0.25	0.31***	0.27**	0.14***	0.11**
	[0.02, 0.53]	[0.00, 0.52]	[0.08, 0.56]	[0.00, 0.47]	[-0.04, 0.47]	[-0.13, 0.40]	[0.00, 0.53]	[-0.02, 0.52]	[0.13, 0.49]	[0.09, 0.46]	[0.07, 0.21]	[0.04, 0.18]
Tweets (log)	-0.22***	-0.22***	-0.30***	-0.30***	-0.31***	-0.31***	-0.22***	-0.22***	-0.20***	-0.20***	-0.11***	-0.11***
	[-0.29, -0.15]	[-0.30, -0.15]	[-0.41, -0.19]	[-0.40, -0.20]	[-0.38, -0.23]	[-0.38, -0.24]	[-0.29, -0.15]	[-0.30, -0.15]	[-0.25, -0.15]	[-0.25, -0.16]	[-0.13, -0.09]	[-0.13, -0.10]
Followers (log)	-0.08*	-0.08	-0.05	-0.03	-0.04	-0.02	-0.08	-0.07	-0.06	-0.05	0.01	0.02
	[-0.17, 0.00]	[-0.17, 0.02]	[-0.13, 0.04]	[-0.12, 0.06]	[-0.12, 0.04]	[-0.10, 0.07]	[-0.16, 0.01]	[-0.17, 0.03]	[-0.12, 0.00]	[-0.11, 0.01]	[-0.02, 0.04]	[-0.01, 0.05]
Gender*ethnicity		0.11		0.50***		0.49***		0.11		0.26**		0.18***
		[-0.23, 0.46]		[0.30, 0.71]		[0.20, 0.78]		[-0.28, 0.49]		[0.09, 0.42]		[0.13, 0.23]
Pseudo R ²	0.026	0.026	0.045	0.046	0.039	0.040	0.025	0.025	0.031	0.032	0.028	0.029
N	1,720,403		1,807,688		1,877,889		1,720,403		1,890,893		1,674,094	
Politician removed	Mark Rutte		Sigrid Kaag		Farid Azarkan		Mark Rutte		Liane den Haan		Geert Wilders	

For each linguistic outcome, the politician with the highest score on that outcome (see Table 4) is removed from the data. Unstandardized regression coefficients *b* and cluster-robust standard errors are reported. Linguistic outcomes were log-transformed. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Author contributions

IV conducted all aspects of this manuscript.

Funding

This work was not supported by any funding.

Availability of data and materials

Data (tweets IDs) and code are available via: <https://osf.io/vf6xt/>. The full dataset is available upon request.

Declarations

Competing interests

The author declares no competing interests.

Received: 28 September 2023 Accepted: 3 February 2024
Published online: 26 February 2024

References

Akhtar, S., & Morrison, C. M. (2019). The prevalence and impact of online trolling of UK members of parliament. *Computers in Human Behavior, 99*, 322–327. <https://doi.org/10.1016/j.chb.2019.05.015>

Amnesty International UK. (2017). *Black and Asian women MPs abused more online*. <https://www.amnesty.org.uk/online-violence-women-mps>. Accessed 22 Dec 2023.

Arel-Bundock, V. (2023). *marginaleffects: Predictions, Comparisons, Slopes, Marginal Means, and Hypothesis Tests*. (R package version 0.12.0) [Computer software]

Bardall, G., Bjarnegård, E., & Piscopo, J. M. (2020). How is political violence gendered? Disentangling motives, forms, and impacts. *Political Studies, 68*(4), 916–935. <https://doi.org/10.1177/0032321719881812>

Barrie, C., & Ho, J. (2022). *Using the twitter academic API with R for social science research*. SAGE Publications Ltd. <https://doi.org/10.4135/9781529609233>

Berbrier, M. (2000). The victim ideology of white supremacists and white separatists in the United States. *Sociological Focus, 33*(2), 175–191. <https://doi.org/10.1080/00380237.2000.10571164>

Bjørgero, T., Jupskås, A. R., Thomassen, G., & Strype, J. (2022). Patterns and consequences of threats towards politicians: Results from surveys of National and Local Politicians in Norway. *Perspectives on Terrorism, 16*(6), 100.

- Blais, M. (2020). *Masculinist discourses on intimate partner violence: Antifeminist men defending white heterosexual male supremacy*. Routledge.
- Calma, J. (2023, May 31). *Scientists say they can't rely on Twitter anymore*. The Verge. <https://www.theverge.com/2023/5/31/23739084/twitter-elon-musk-api-policy-chilling-academic-research>. Accessed 22 Dec 2023.
- ECRI (2015). General Policy Recommendation No. 15 on Combating Hate Speech
- Esposito, E., & Breeze, R. (2022). Gender and politics in a digitalised world: Investigating online hostility against UK female MPs. *Discourse & Society*, 33(3), 303–323. <https://doi.org/10.1177/09579265221076608>
- Every-Palmer, S., Barry-Walsh, J., & Pathé, M. (2015). Harassment, stalking, threats and attacks targeting New Zealand politicians: A mental health issue. *Australian & New Zealand Journal of Psychiatry*, 49(7), 634–641. <https://doi.org/10.1177/0004867415583700>
- Fuchs, T., & Schäfer, F. (2021). Normalizing misogyny: Hate speech and verbal abuse of female politicians on Japanese Twitter. *Japan Forum*, 33(4), 553–579. <https://doi.org/10.1080/09555803.2019.1687564>
- Gorrell, G., Bakir, M. E., Roberts, I., Greenwood, M. A., & Bontcheva, K. (2020). Which politicians receive abuse? Four factors illuminated in the UK general election 2019. *EPJ Data Science*, 9(1), 18. <https://doi.org/10.1140/epjds/s13688-020-00236-9>
- Huang, J., & Ling, C. X. (2005). Using AUC and accuracy in evaluating learning algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 17(3), 299–310. <https://doi.org/10.1109/TKDE.2005.50>
- James, D. V., Farnham, F. R., Sukhwai, S., Jones, K., Carlisle, J., & Henley, S. (2016). Aggressive/intrusive behaviours, harassment and stalking of members of the United Kingdom parliament: A prevalence study and cross-national comparison. *The Journal of Forensic Psychiatry & Psychology*, 27(2), 177–197. <https://doi.org/10.1080/14789949.2015.1124908>
- Jonker, J., & van der Parre, H. (2022, November 20). *Onderzoek NOS: Kamerleden voelen zich geïntimideerd, houden mening soms voor zich*. <https://nos.nl/artikel/2453170-onderzoek-nos-kamerleden-voelen-zich-geintimideerd-houden-mening-soms-voor-zich>. Accessed 22 Dec 2023.
- Jonker, J. (2022, December 15). *Recordaantal meldingen bedreigde politici*. <https://nos.nl/artikel/2456538-recordaantal-meldingen-bedreigde-politici>. Accessed 22 Dec 2023.
- Krook, M. L., & Sanín, J. R. (2020). The cost of doing politics? Analyzing violence and harassment against female politicians. *Perspectives on Politics*, 18(3), 740–755. <https://doi.org/10.1017/S1537592719001397>
- NOS nieuws. (2022, January 5). *Politie pakt man met brandende fakkel op bij huis D66-leider Sigrid Kaag* (Accessed: December 22, 2023). <https://nos.nl/artikel/2412074-politie-pakt-man-met-brandende-fakkel-op-bij-huis-d66-leider-sigrid-kaag>
- NOS nieuws. (2023, July 13). *Kaag vertrekt als D66-leider: "Mijn werk trekt een wissel op mijn gezin"*. <https://nos.nl/artikel/2482592-kaag-vertrekt-als-d66-leider-mijn-werk-trekt-een-wissel-op-mijn-gezin>. Accessed 22 Dec 2023.
- Openbaar Ministerie. (2023, May 22). *Niet eerder zoveel meldingen van bedreiging politici—Nieuwsbericht—Openbaar Ministerie* [Nieuwsbericht]. Ministerie van Justitie en Veiligheid. <https://www.om.nl/actueel/nieuws/2023/05/22/niet-eerder-zoveel-meldingen-van-bedreiging-politici>. Accessed 22 Dec 2023.
- Paul, K., & Dang, S. (2022, December 5). Exclusive: Twitter leans on automation to moderate content as harmful speech surges. *Reuters*. <https://www.reuters.com/technology/twitter-exec-says-moving-fast-moderation-harmful-content-surges-2022-12-03/>. Accessed 22 Dec 2023.
- Pérez-Arredondo, C., & Graells-Garrido, E. (2021). Twitter and abortion: Online hate against pro-choice female politicians in Chile. *Journal of Language Aggression and Conflict*, 9(1), 127–154. <https://doi.org/10.1075/jlac.00056.per>
- Phillips, L. J., Pathé, M., & McEwan, T. (2023). Gender differences in stalking, threats and online abuse reported by Victorian politicians. *Psychiatry, Psychology and Law*. <https://doi.org/10.1080/13218719.2022.2142975>
- Robitzsch, A., & Grund, S. (2023). *miceadds: Some Additional Multiple Imputation Functions, Especially for "mice"*. (R package version 3.16–18) [Computer software]. <https://CRAN.R-project.org/package=miceadds>
- Saris, K., & van de Ven, C. (2021, March 3). *De online haat dreigt vrouwen uit de politieke arena te verdrijven*. De Groene Amsterdammer. <https://www.groene.nl/artikel/misogynie-als-politiek-wapen>
- Scott, J. (2019, October 31). Women MPs say abuse forcing them from politics. *BBC News*. <https://www.bbc.com/news/election-2019-50246969>. Accessed 22 Dec 2023.
- Southern, R., & Harmer, E. (2021). Twitter, incivility and "everyday" gendered othering: An analysis of tweets sent to UK Members of Parliament. *Social Science Computer Review*, 39(2), 259–275. <https://doi.org/10.1177/0894439319865519>
- Trifiro, B. M., Paik, S., Fang, Z., & Zhang, L. (2021). Politics and politeness: Analysis of incivility on twitter during the 2020 democratic presidential primary. *Social Media + Society*, 7(3), 20563051211036940. <https://doi.org/10.1177/20563051211036939>
- Tromble, R., & Koole, K. (2020). She belongs in the kitchen, not in Congress? Political engagement and sexism on twitter. *Journal of Applied Journalism & Media Studies*, 9(2), 191–214. https://doi.org/10.1386/ajms_00022_1
- Veerbeek, J. (2021, March 3). *Verantwoording bij het onderzoek naar seksisme in de politieke arena*. De Groene Amsterdammer. <https://www.groene.nl/artikel/verantwoording-bij-het-onderzoek-naar-seksisme-in-de-politieke-arena>
- Vidgen, B., Margetts, H., & Harris, A. (2019). *How much online abuse is there? A systematic review of evidence for the UK*. https://www.turing.ac.uk/sites/default/files/2019-11/online_abuse_prevalence_full_24.11.2019_-_format_0.pdf
- Votta, F. (2021). *peRpective: Interface to the "Perspective" API* [Computer software]. <https://cran.r-project.org/web/packages/peRpective/index.html>
- Wagner, A. (2022). Tolerating the trolls? Gendered perceptions of online harassment of politicians in Canada. *Feminist Media Studies*, 22(1), 32–47. <https://doi.org/10.1080/14680777.2020.1749691>
- Ward, S., & McLoughlin, L. (2020). Turds, traitors and tossers: The abuse of UK MPs via Twitter. *The Journal of Legislative Studies*, 26(1), 47–73. <https://doi.org/10.1080/13572334.2020.1730502>
- Waseem, Z., Davidson, T., Warmesley, D., & Weber, I. (2017). *Understanding Abuse: A Typology of Abusive Language Detection Subtasks* (arXiv:1705.09899). arXiv. <https://doi.org/10.48550/arXiv.1705.09899>
- Witteveen, D. (2017). The Rise of Mainstream Nationalism and Xenophobia in Dutch Politics. *Journal of Labor and Society*, 20(3), 373–378. <https://doi.org/10.1163/24714607-02003006>
- Zakrzewski, C., Siddiqui, F., & Menn, J. (2022, November 23). Musk's 'free speech' agenda dismantles safety work at Twitter, insiders say. *Washington Post*. <https://www.washingtonpost.com/technology/2022/11/22/elon-musk-twitter-content-moderations/>. Accessed 22 Dec 2023.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.