

Gendered Perceptions and the Costs of Political Toxicity: Experimental Evidence from Politicians and Citizens in Four Democracies

GREGORY EADY *University of Copenhagen, Denmark*

ANNE RASMUSSEN *King's College London, United Kingdom, and University of Copenhagen, Denmark*


Politicians frequently face toxic behaviors. We argue that these behaviors impose a double burden on women, who may not only face higher exposure to toxicity, but experience attacks that they and others understand to be motivated by prejudice and designed to push them out of office. Using large-scale image-based conjoint experiments in the United States, Denmark, Belgium, and Chile, we demonstrate that both politicians themselves and citizens regard messages targeting women politicians as more toxic than otherwise equivalent messages targeting men. This perception intensifies when messages mention gender or come from perpetrators who are men. A second experiment to investigate the mechanisms shows that hostile behaviors toward women are more frequently understood as driven by prejudice and attempts to remove women from politics. These findings highlight the importance of understanding how perceptions of perpetrators' motives affect the severity of political toxicity, and provide insights into the gendered effects of political hostility.


INTRODUCTION

Online harassment, abuse, and intimidation of politicians are on the rise (Collignon, Campbell, and Rüdig 2022).¹ Women politicians are widely suggested to bear the highest burden of these behaviors (Astor 2018; Dhrodia 2017; Specia 2019). A cross-country survey of women parliamentarians shows, for example, that a majority report facing contemptuous comments online and having sexually explicit material concerning them shared on social media (Inter-Parliamentary Union 2016). There is also evidence that women politicians—especially high-profile women—encounter these behaviors more frequently than their counterparts who are men (Collignon and Rüdig 2021; Daniele, Dipoppa, and Pulejo 2023; Håkansson 2021; Herrick and Thomas 2021; Rheault, Rayment, and Musulan 2019). For some women politicians, toxic messages and harassment are

now considered “just part of the job” (Terris 2016). Yet, differences in both exposure to and interpretations of toxic behaviors directed at women and men politicians could have substantial implications for political representation. Toxic behavior can affect both the willingness of women to remain in politics, and whether ordinary women who observe this behavior aspire to and run for public office (Every-Palmer, Barry-Walsh, and Pathé 2015; Maisel 2012). It could also affect women's willingness to engage in democratic dialogue in the public sphere altogether (Every-Palmer, Barry-Walsh, and Pathé 2015; Maisel 2012; Theocharis et al. 2020; Tromble 2018). In this article, we argue that understanding these downstream consequences requires that researchers document not only differences in the frequency and content of toxic behaviors faced by women and men politicians but also *how* these behaviors are perceived by politicians themselves, who may be the targets of toxic behaviors, and by ordinary citizens who witness them.

Theoretically, we argue that women politicians face what we refer to as a double burden with respect to toxic behavior: women politicians may not only be exposed to higher rates of toxic behaviors—the focus of much of the literature to date—but also experience these behaviors as more severe due to beliefs about the gendered prejudices that drive them. Our expectation is that politicians and citizens witnessing hostile behaviors toward women politicians will perceive them as more severe than otherwise equivalent behaviors toward men. The mechanism, we argue, is that the severity of hostile behavior is a function not only of a behavior's content but also of the perceived motives and underlying prejudices that drive it. Hostile messages sent toward women, we expect, will be more

Corresponding author: Gregory Eady , Associate Professor, Department of Political Science and the Center for Social Data Science (SODAS), University of Copenhagen, Denmark, gregory.eady@gmail.com.

Anne Rasmussen , Professor, Department of Political Economy, King's College London, United Kingdom; Professor, Department of Political Science, University of Copenhagen, Denmark, anne.rasmussen@kcl.ac.uk.

Received: August 16, 2023; revised: April 20, 2024; accepted: September 20, 2024.

¹ We refer to political “toxic behavior” and “hostility” throughout as a general term that encompasses forms of aggressive behaviors or expressions directed at others that include, for example, incivility, intolerance, cyber-bullying, harassment, hate speech, and “trolling” (Sandberg and Segesten 2022).

likely to be perceived as motivated by prejudices and by a perpetrator's goal to push women out of office (as compared to equivalent messages sent to men). If true, this would suggest that the severity of toxicity directed at women and men politicians (i.e., its potential cost to the recipient) depends on the underlying meaning given to it by the targets of toxic behavior and those who observe it.

This article thus investigates an under-examined and pernicious side of toxic behavior, whereby the perceived prejudices that drive negative behaviors in the political workplace affect understandings of the severity of their harm. As the literature on discrimination shows, being subjected to behavior that is perceived as motivated by prejudice magnifies its harm through, for example, decreasing mental health, increasing physical illness, and lowering job satisfaction (Dardenne, Dumont, and Bollier 2007; Dover, Hunger, and Major 2021; Pascoe and Smart Richman 2009).² Thus, understanding how and why toxic behaviors toward women are considered more severe by the recipients of that behavior (women politicians) and by those who may enter politics in the future (citizens) speaks to how the burden of toxicity can vary by gender even in the face of what can appear superficially as equivalent behaviors.

To investigate this, we use a series of image-based conjoint experiments with elected politicians and citizens. We examine the effects of the gender of a politician who is targeted with a politically toxic behavior by presenting politician and citizen respondents with pairs of images of social media conversations that are designed to visually mimic hostile exchanges on social media. The experiments, which consist of millions of possible images, allow us to isolate the causal effect of the attributes of the perpetrator who sends a hostile message and those of the politician receiving it, holding a variety of conversation attributes constant. This enables us to examine how various characteristics that are commonly present in online toxic exchanges with politicians influence the way that politicians and citizens interpret toxic behavior directed at politicians. For example, this allows us to determine whether contextual characteristics of an exchange (e.g., the gender of the perpetrator, or whether a message explicitly highlights a politician's gender) amplify or diminish the impact of a politician's gender on perceptions of its severity. This also makes it possible for us to document how sensitive assessments of the link between a politician's gender and toxic behavior are to characteristics of the respondents who perceive them (e.g., political affiliation, co-partisanship, whether a respondent is a politician or citizen, gender, etc.).

In a second experiment, we then investigate the mechanisms that drive differences in understandings of the severity of hostility toward women and men politicians. This enables us to examine whether the gender of a politician who is the target of politically hostile behavior affects politicians' and citizens' interpretations of the

motivations of the perpetrator (e.g., whether driven by policy differences, prejudice, or a desire to push a politician out of office). To ensure that our findings are robust to the idiosyncrasies of individual countries, we collect data from currently elected politicians and ordinary citizens in the United States, Denmark, Belgium, and Chile. Due to potential differences in online exposure and tolerance of toxic behavior between political elites and ordinary citizens—especially considering that the former are the targets of the toxic behavior we examine—we conduct experiments among both types of actors in each of these four political systems.

We find strong evidence across countries and among both politicians and citizens that toxic behaviors toward women politicians are assessed as more severe than otherwise equivalent behaviors toward men politicians. This result is consistent across respondents, regardless of their gender, partisanship, political ideology, and relevant subgroups. Further, in line with expectations, the effect of targeting women politicians with hostile messages is stronger when the perpetrator is a man, and when the message highlights the gender of the politician.

We also document how the inferred motivations of the person sending a politically toxic message differ based on the gender of the targeted politician. In our second experiment, which we design to evaluate the underlying mechanisms, we show that respondents interpret the motives of those who send hostile messages toward women politicians differently than equivalent messages sent toward men: messages sent to women politicians are less likely to be viewed as driven by policy and more likely to be considered motivated by prejudice or a desire to push a politician out of politics.

Overall, our study makes significant contributions to the literature on gender bias in politics and political toxicity by demonstrating the divergence in the perceived severity of toxic behaviors when directed at women and men politicians, even when the behavior's content is comparable. Further, we show how interpretations of gendered prejudices and the use of toxicity to deter women from participating in politics are mechanisms driving this result. The implication is that political toxicity may have larger downstream effects on women than men through channels beyond the frequency or type of toxic behaviors themselves. Dealing with toxic behavior interpreted to be motivated by prejudice and a desire to expel them from office could, for example, further decrease women's willingness to continue in politics, to run for national office, to enter politics more generally, or to engage in public dialogue altogether. It speaks, furthermore, to the reasons why women politicians express more concern about political toxicity generally (see Section D of the Supplementary Material), and anecdotally, why a toxic political atmosphere is often cited by women, relative to men, as a reason for leaving office. This article thus highlights the need to investigate how politicians themselves, and citizens, understand the toxic behaviors that they may face or observe in their daily political lives, and the importance of investigating understandings of the motives that drive these behaviors.

² The literature on prejudice-based violent crime unsurprisingly shows similar findings (McDevitt et al. 2001).

THEORETICAL FRAMEWORK

When women run for office, research finds that they tend to win elections at similar rates as men (Darcy, Welch, and Clark 1994; Fox 2005; King and Leigh 2010; Thomsen 2020). According to candidate choice experiments, voters slightly prefer women candidates to men candidates (Schwarz and Coppock 2022). Yet, the proportion of women politicians who aspire to public office remains stubbornly below parity with men in most countries (Thomsen and King 2020). Research in political science has thus investigated the barriers to representation for women in politics, and highlighted differences in the costs of being a politician for women and men. These barriers and costs are many, and are shown to include, among others, being subjected to gender-based stereotypes, facing weaker opportunities for development, being expected by voters to hold the double burden of family and politician, and having to meet higher standards to qualify for leadership positions (e.g., Bauer 2019; Bernhard, Shames, and Teele 2021; Bjarnegård 2013; Heath, Schwindt-Bayer, and Taylor-Robinson 2005; Teele, Kalla, and Rosenbluth 2018).³

A recent important avenue of research examines the darker costs of public office by investigating the prevalence of hostile, threatening, and violent interactions with members of the general public. Studies show that majorities of elected politicians in many countries are subjected to psychologically and physically violent interactions with the public at some point during their political careers (e.g., Ballington 2018; Bjarnegård 2023; Inter-Parliamentary Union 2016).

Research also shows that such behavior is not equally distributed across gender. Campaign volunteers who are women receive more hostile responses than men volunteers, and women politicians are more likely to be the targets of psychological or physical violence compared to men (Daniele, Dipoppa, and Pulejo 2023; Herrick et al. 2019; Yan and Bernhard 2024). This is especially the case for high-profile women (Håkansson 2021; Herrick and Thomas 2021; Rheault, Rayment, and Musulan 2019). As recent research suggests, the disproportionate number of attacks on women are unlikely to be due to differences in policy positions between women and men politicians (Daniele, Dipoppa, and Pulejo 2023).

The consequences for politicians and political representation are decidedly negative. Herrick and Franklin (2019) suggest, for example, that exposure to violence is associated with psychological harm and decreases in ambitions to remain in office. As one woman candidate who ran in a recent Canadian election remarked: “I would never ever, ever subject myself to [the abuse I was subjected to as a candidate] again ... [I]t has damaged my mental health. It has made me fear for the safety of my family. It has made

me fear for my safety” (Lamensch 2021). In the United Kingdom, women candidates were shown to modify their campaign efforts as a result (Collignon, Campbell, and Rüdiger 2022).

Situations in which politicians encounter toxic political behavior have greatly increased with the emergence of social media. As a technology, social media is double-edged (Tucker et al. 2017). On the one hand, it is considered by many politicians and staffers as an essential tool for political communication and representation (McGregor 2020). On the other hand, by lowering barriers to entry for the general public to interact with their representatives, social media has increased opportunities for politicians to be exposed to uncivil and threatening messages, hate speech, and harassment, among other toxic behaviors (Collignon, Campbell, and Rüdiger 2022).

Research into toxic behavior directed at politicians online (and offline) has increasingly looked at differences in the frequency and content of these behaviors, using, for instance, automated text analysis and machine learning with social media data. Theocharis et al. (2016; 2020) find, for example, that roughly 18% of posts on Twitter that mention U.S. members of Congress are uncivil; that members adopting extreme positions are more likely to be the targets; and that levels of engagement on social media increase politicians’ exposure to toxic messages and harassment. Rheault, Rayment, and Musulan (2019) find that toxic messages are more frequently directed at high-profile women politicians than high-profile men (with mixed results among low-profile politicians).

Yet, women politicians may be faced with a double burden regarding toxic behaviors. For one, they may be more likely to be exposed to higher rates of toxic behaviors (the focus of the existing literature). They may also, however, bear a second burden that exacerbates the severity of these behaviors: political toxicity for women can indicate efforts by perpetrators that are motivated by gendered prejudices and by a desire to use toxicity to increase the cost for women to run for or remain in politics—motivations that will not be recognized similarly for comparable attacks on their counterparts who are men. Thus, even if there were minimal evidence of differences in the frequency of toxic behaviors, the severity of those behaviors can still vary depending on how the motivations behind them are understood. This matters because perceptions of the motives and prejudices behind toxic behavior will shape how politicians targeted by this behavior assess their severity, and how the public understands the toxicity of the political environment that women may encounter if they enter politics. As evidence suggests in other domains, behaviors driven by prejudice can have especially harmful effects on the mental and physical health of those targeted (e.g., Dardenne, Dumont, and Bollier 2007; Dover, Hunger, and Major 2021).

To date, however, we have relatively limited knowledge regarding how politicians and citizens understand the severity of these interactions (Bardall, Bjarnegård, and Piscopo 2020) and how perceptions of the motives

³ See also, ACE Electoral Knowledge Network (2023).

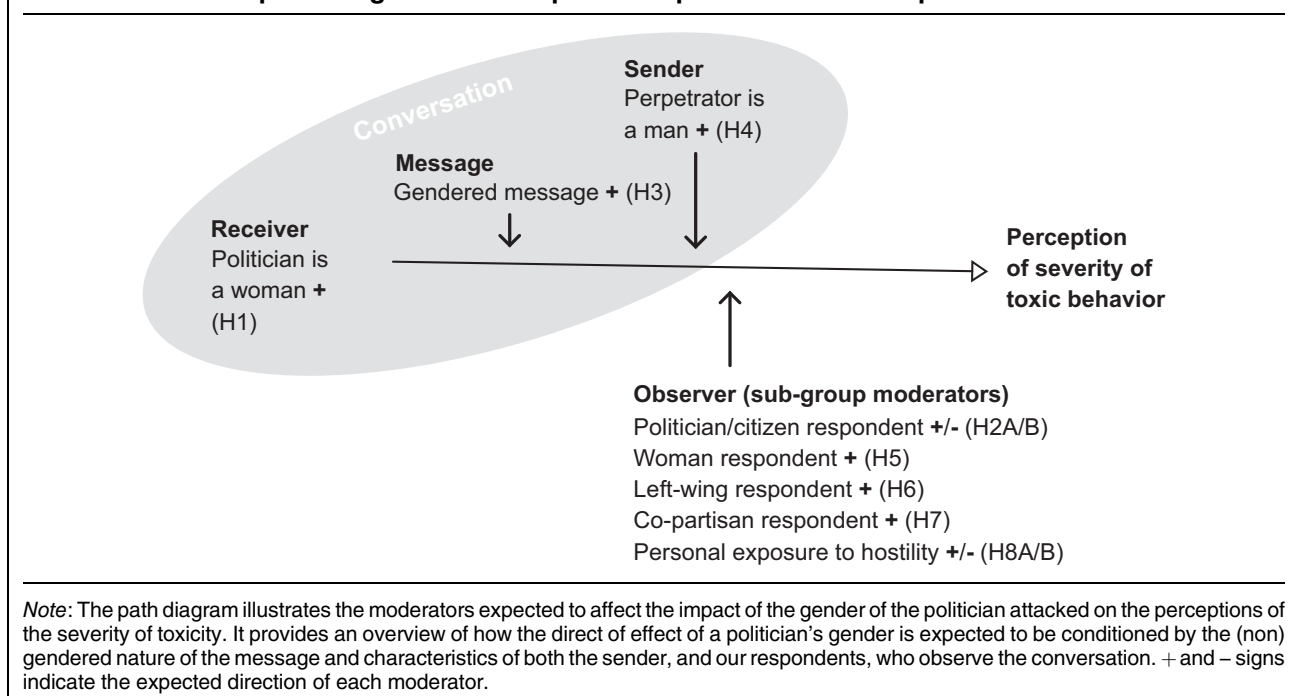
that drive toxic behavior toward politicians speak to understandings of prejudices in everyday political interactions. To investigate this, we conduct experimental studies among elected politicians and citizens in four democracies to determine how key contextual factors regarding gender affect perceptions of the severity of toxic messages directed toward politicians. In doing so, we contribute to a small body of work on perceptions of online toxicity. Examples include work on how the public perceives the incivility of news comments (Kenski, Coe, and Rains 2017); its impact on news credibility, political efficacy, and political trust (Borah 2013); and on preferences for regulation of hate speech (Rasmussen 2024). These studies have predominantly examined citizens' perceptions (e.g., Kenski, Coe, and Rains 2017; Rasmussen 2024; Stryker, Conway, and Danielson 2016). By conducting experiments among both politicians and citizens, we uniquely investigate how the gender of politicians attacked online influences perceptions of toxic behavior among those who experience it (i.e., the politicians themselves) and also among those who witness these interactions (i.e., citizens). Furthermore, by varying additional characteristics of the conversation and considering heterogeneous treatment effects among respondents, we can examine how the effect of a politician's gender on assessments of its severity is influenced by the attributes of the message, its sender, and the individuals observing the conversation. In Figure 1, we provide a conceptual illustration of the hypotheses that we discuss in the subsequent sections, and which are rooted in existing research on gender stereotypes, social identity, and the role of women in politics.

Hypotheses

Gender of the Elected Representative

Women politicians are widely portrayed (e.g., Astor 2018; Camut 2023; Hunt, Evershed, and Liu 2016; Lamensch 2021; Mekouar 2019; Morgan 2020; Specia 2019) and shown empirically to experience substantial amounts of toxic behavior (Collignon and Rüdiger 2021; Daniele, Dipoppa, and Pulejo 2023; Håkansson 2021; Krook 2020; Rheault, Rayment, and Musulan 2019). These behaviors have been linked to prejudices against women in politics. Recent evidence suggests that attacks on women politicians often result from a perpetrator's aim to silence or push them out of office as a form of misogynistic backlash, rather than from policy disagreements (Daniele, Dipoppa, and Pulejo 2023, see also Krook 2020; Krook and Sanín 2019). Prejudices that drive these behaviors could include, for example, beliefs that women make worse or less legitimate political leaders than men (Mo 2015; Vial, Napier, and Brescoll 2016); that women are less emotionally suitable than men for public office (Carnevale, Smith, and Campbell 2019); that women do not have the necessary personal qualities for political leadership (Banwart 2010); or that women in politics are failing to uphold traditional gender norms, triggering hostility and backlash (Karpowitz and Mendelberg 2014; Lawless and Fox 2010; Okimoto and Brescoll 2010; Smith, Paul, and Paul 2007; Teele, Kalla, and Rosenbluth 2018). Yan and Bernhard (2024), for example, suggest that the potential mechanisms driving a higher frequency of attacks on female-named campaign volunteers could

FIGURE 1. Conceptual Diagram of the Expected Experimental and Respondent-Level Moderators



be related to the internalized sexist attitudes of perpetrators.

Such prejudices we expect are, in general, recognized by politicians and by many citizens as driving and magnifying the harm of toxic behaviors toward women. This would align with the empirical literature showing that prejudiced behaviors can affect the mental and physical health outcomes of those who experience or witness it (Dardenne, Dumont, and Bollier 2007; Dover, Hunger, and Major 2021; Pascoe and Smart Richman 2009). This also aligns with existing evidence indicating that the general public, particularly those who uphold egalitarian ideals, exhibit a heightened awareness of behaviors stemming from prejudice (Schmader et al. 2012). In other words, hostile behaviors motivated by prejudice are regarded as more insidious than those driven by, for instance, intense policy disagreements. The heightened sensitivity and acknowledgment of the negative impacts of prejudice toward women should logically also have intensified alongside the overall reduction in gender-based biases.⁴ While this internalization of egalitarian values pertaining to women's suitability and competence in politics may not be universal, we can reasonably anticipate that it has progressively raised awareness among politicians and the public regarding the magnitude and consequences of gender-related biases. In sum, we expect that politicians—who know firsthand the harm of toxic behavior—and citizens, will perceive toxic messages sent to women politicians as more severe than equivalent messages sent to men politicians in our first conjoint experimental study (H1). We test the proposed mechanism concerning prejudice-driven behavior in a second experimental study, which we describe further below.⁵

This hypothesized difference in understandings of the severity of toxic behaviors toward women and men politicians could manifest differently among politician and citizen respondents. On the one hand, politicians could exhibit heightened sensitivity to prejudices directed at women members of their own social and professional in-group (Smith 1993; Tajfel 1981). Alternatively, this sensitivity might stem from personal exposure to instances where women politicians have encountered noxious behaviors and the subsequent repercussions, either through direct observation or firsthand accounts (H2A). On the other hand, there are reasons why politicians may be less sensitive than citizens to toxic behavior toward women politicians (H2B). Evidence suggests that women preemptively adopt male-dominant behavior patterns to avoid gender-biased judgments (Dittmar, Sanbonmatsu, and

Carroll 2018; Lazarus and Steigerwalt 2018). Politicians may therefore be cautious to categorize hostile behavior toward women politicians as more severe to prevent reinforcing negative gender stereotypes associated with women in politics.

Gender of the Perpetrator and Gendered Message Content

The tendency to perceive messages to women politicians as more toxic than those to men politicians should be stronger if they explicitly indicate that a politician is a woman (H3) and are from a perpetrator who is a man rather than a woman (H4). We expect that messages indicating the gender of a politician, even if stated seemingly neutrally, will serve as a heuristic for whether the sender of a toxic message is motivated by gender-based considerations. Similarly, messages that are sent by men can be expected to be perceived as driven by different motives than those sent by women, even if the message text is equivalent. When men engage in political hostility toward women, we expect that they will be more likely to be perceived as driven by gendered-based prejudices than if similar hostility were perpetrated by women. Finally, differences in assessments of the severity of toxic behavior sent by men and women could also signify a reduced tolerance for toxic behaviors directed at women politicians who might be seen as recurrent targets of such behavior from men (Amnesty International 2018).

Respondent Characteristics and Experiences

Finally, we consider how a number of characteristics of politician and citizen respondents moderate the extent to which toxic messages toward women politicians are perceived as more toxic than those toward men. We expect, first, that tolerance toward hostile attacks on women relative to those on men will be lower for respondents who are women than are men (H5). Gender might act as a social identity (Tajfel 1978): those of the same gender could feel stronger affinity to each other, for example, due to feelings of group solidarity with those whose fortunes are linked to one's own (Dolan 2008). Studies of voting behavior have found support for the idea that gender affinity matters, demonstrating that women prefer candidates who are women, and men prefer candidates who are men (e.g., Dolan 1998; Plutzer and Zipp 1996; Sanbonmatsu 2002, but see Dolan 2008; Teele, Kalla, and Rosenbluth 2018). Women have also been shown to be more sensitive to toxic behavior in general (Kenski, Coe, and Rains 2017), potentially due to different norms in expression for men and women, or from experiencing or witnessing higher frequencies of toxic behavior.

Second, we expect respondents' political ideology to affect how they judge toxic behavior toward women politicians, with left-wing respondents being more sensitive to attacks on women than right-wing respondents (H6). Right-wing parties are frequently associated with support for traditional gender roles and regard sexism as less of an issue than left-wing politicians (Craig and Cossette 2022; Pratto, Stallworth, and Sidanius 1997).

⁴ For example, in the US in 1975, approximately half of American citizens expressed a belief that men were more emotionally suited for politics than women. In 2018, that number had dropped nearly fourfold, to 13% (Carnevale, Smith, and Campbell 2019).

⁵ The ordering of the hypotheses in this article differs from that in the preregistration. Section B of the Supplementary Material shows the link between the hypothesis numbers presented here and those in the preregistration. All hypotheses documented in the preregistration are tested and presented in this article.

In the U.S. context, for example, the populist turn of the Republican party has also used the purported defense of traditional culture (e.g., restrictions on affirmative action, quotas, abortion, and gender equality) as a means to gain support (Norris and Inglehart 2019).

Third, judgments of any political action, and certainly speech by politicians and reactions to it by others, is frequently subject to partisan-motivated reasoning, with judgments interpreted in ways that confirm existing beliefs (Bolsen, Druckman, and Cook 2014). Politicians and citizens will likely perceive the motivations that drive out-partisan perpetrators as more likely driven by prejudice. We expect, therefore, that the effect of toxic messages to women politicians compared to men politicians will be stronger when a respondent is a co-partisan of the politician being attacked (H7).

Finally, we expect respondents' own personal exposure to online toxic behavior to play a role, but that the direction of the conditional effect could go both ways. On the one hand, those who have experienced online toxic behavior may be better equipped to empathize with or be sensitive to women politicians who are underrepresented in politics (UN Women 2023) and more likely to be perceived as targeted for prejudicial and other gender-based reasons. Understanding the hardships faced by others, either directly or, for example, through perspective-taking, is shown to foster support for others in disadvantaged positions (Adida, Lo, and Platas 2018; Bor and Simonovits 2021; Kubn et al. 2021). On the other hand, exposure to toxicity could desensitize respondents to these toxic behaviors in general and make them less sensitive to potential differences in the nature and motivations of attacks on men and women politicians (Collignon, Campbell, and Rüdiger 2022). This desensitization may reduce respondents' ability to recognize that these behaviors meaningfully affect others, limiting their capacity to empathize more with women politicians who may be especially targeted due to their gender. In sum, the tendency to regard attacks on women politicians as more severe could thus either be stronger (H8A) or weaker (H8B) among politicians (and citizens) who report being subjected to online attacks themselves.

EXPERIMENTAL DESIGN

We test our hypotheses with two sets of preregistered image-based conjoint experiments (Vecchiato and Munger 2022) that we included in surveys of elected politicians ($N = 2,821$) and citizens ($N = 5,376$) in the US, Denmark, Belgium, and Chile.⁶ We ran the experiment in four relatively disparate democracies. These cases vary, for example, by whether their system is parliamentary or presidential (DK, CL, BE/US) and federal or

unitary (DK, CL/BE, US); the percentage of national politicians who are women (DK 45%, BE 42%, CL 36%, US 29%; Inter-Parliamentary Union 2024); their gender inequality rankings (DK 1st, BE 11th, US 44th, CL 49th; UNDP 2024); ongoing events or social movement actions concerning women's rights (e.g., leak of/overturning of *Roe v. Wade* in the US; the *Ni Una Menos* movement against gender-based violence in South America); and beliefs about the acceptability of intimate-partner violence against women (DK 0%, BE 3%, US 14%, CL 31%; OECD 2024). These, among other cross-national differences, enable us to investigate whether the effects presented in the "Results" section are specific to one particular political, institutional, or cultural context. For this reason, as specified in the experiment's preregistration, effect estimates are presented both by country and in aggregate.

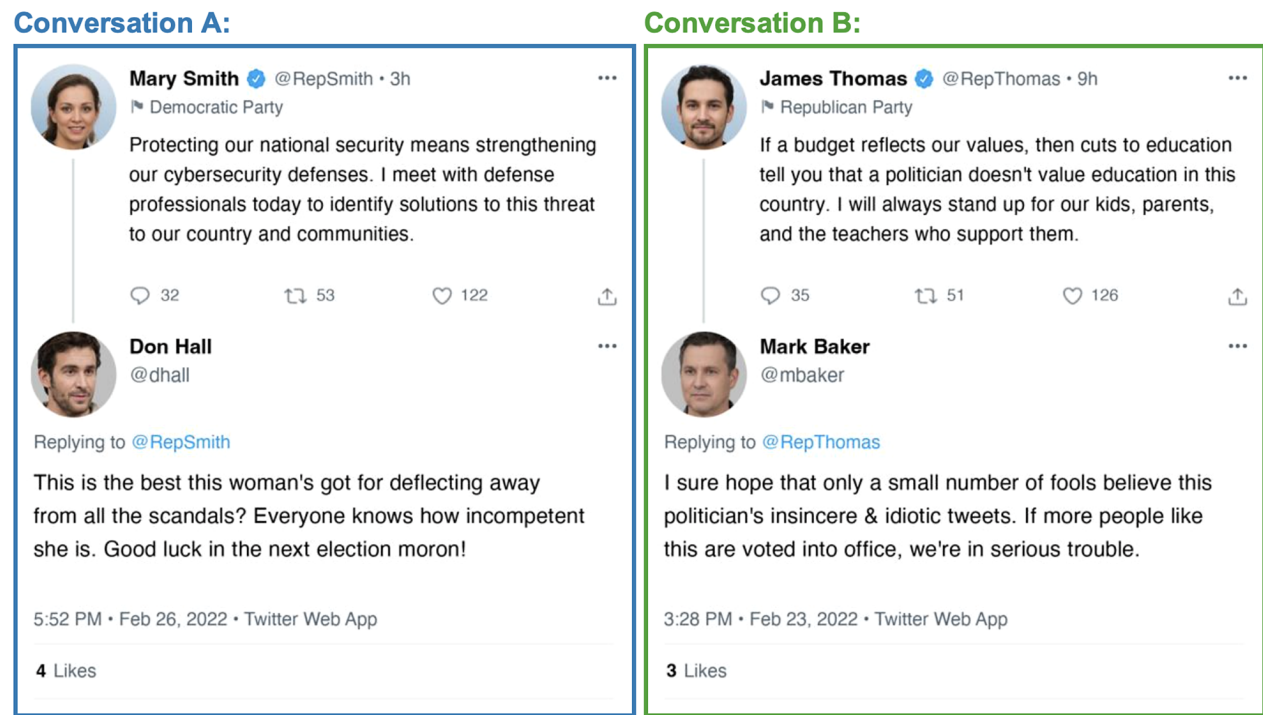
The surveys of politicians, which consist of a mix of national, regional and local-level representatives, were conducted by the authors in Denmark, Belgium (Flanders), and Chile between March and July 2022, and by CivicPulse—a research firm dedicated to surveying politicians—in the US, between September and October 2022 (for details of the politician sample, see Section C of the Supplementary Material). Surveys of politicians in the US and Belgium do not include national politicians due to pragmatic challenges of, for example, obtaining access to politicians at the highest level (e.g., sitting U.S. senators and congresspersons). We note that an important benefit of examining local and national politicians is that, while national politicians may face more online toxic behavior, attacks on local politicians are not rare,⁷ and local politicians constitute an important pool of potential candidates for electoral office. Women also are more likely than men to begin their career in politics at the local level (Berevoescu and Ballington 2021).⁸ As we note further below, we examine (the absence of) differences in our results for local and national politicians in Section K of the Supplementary Material. Finally, surveys of citizens were conducted by the survey firm Dynata (formerly SSI Research) between May and October 2022, and are representative of the voting-age population for each country (for details, see Section C of the Supplementary Material). Citizens can be regarded as third party observers to the actions under study, but also embody potential politicians or participants in other political arenas. This is particularly noteworthy due to the underrepresentation of women in politics generally. The potential indirect costs of witnessing toxic behavior toward women politicians may not only affect those currently in office, but also prospective ones, potentially dampening their ambitions to enter politics.

We conduct two experimental studies to assess how the characteristics of conversations between ordinary

⁶ The research design was approved by the institutional review board at the corresponding author's university. The study was conducted online and involved informed and voluntary consent, and respondents were informed that they could end participation in the study at any time. Citizen respondents were also provided monetary compensation consistent with the survey firm's market rate.

⁷ In our surveys of national and local politicians, 60% of national politicians and 42% of local politicians report exposure to online toxic behavior.

⁸ This latter fact has been linked to differences in perceptions between men and women about their suitability for national office (Lawless and Fox 2005).

FIGURE 2. Example of the Paired Visual Conjoint Experiment

Note: This figure displays an example of the paired conjoint design, in which the attributes varied are each politician's gender (name, photo); ordinary user's gender (name, photo); politicians' partisanship; text sent by the politician; toxic message sent by the ordinary user; and whether the gender of the politician is highlighted in the toxic message from the ordinary user.

citizens and politicians on social media affect how politician and citizen respondents perceive political toxicity directed toward women and men. In the first study, we create images of social media conversations at scale (millions of images)⁹ to represent all permutations of a large set of treatment conditions, which we detail below. These images, which visually mimic toxic exchanges on the micro-blogging platform Twitter are presented to politician and citizen respondents in pairs. An example is presented in Figure 2. In each conversation, an ordinary social media user was shown responding to a politician. Respondents were told that they would be shown a series of conversations between fictitious citizens and politicians, and asked which of the conversations in each conjoint task they deemed more disrespectful. We chose to use this term because it is a relatively neutral term compared to a term like "toxicity" (which has stronger normative connotations and is often associated in media with behavior toward women), or a term like "uncivil" (which is less accessible/common).¹⁰

In a second study, we then investigate the mechanisms underlying politicians' and citizens' perceptions

of toxic messages sent to politicians. To do so, we present respondents with single images of toxic exchanges, and ask them a set of seven (randomly ordered) questions that probe how they perceive the motives of the user who is shown sending the toxic message. Respondents are asked about the extent to which they believe that the user is driven by prejudice or discrimination; by aims to discourage the politician from staying in politics; by opinion differences with the politician; by dislike of the politician's party; by the user's dissatisfaction with his/her own life; by a desire to get a reaction from the politician; and by a desire to get a reaction from other users. These questions allow us to examine whether citizens who send toxic messages to women or men politicians are perceived to be driven by different motivations even if the content of the message is held constant.

For treatment conditions in the two sets of experiments, we vary six attributes in the social media images. First, the gender of the politician is varied, as indicated both by the photo and name of the politician. In total, eighteen fictitious politicians were used, comprising nine women and nine men. Given that the experiment is conducted in four culturally distinct countries, the names of politicians were chosen to be culturally appropriate for each country, drawn from lists of the most popular first and last names in each country. The photos of politicians remained consistent across all countries. Second, and relatedly, the gender of the user

⁹ See Section A of the Supplementary Material for details.

¹⁰ The term "disrespect" (unlike, say, "uncivil") can also be more straightforwardly translated for the three non-English language surveys, with the meaning of the term similar across contexts.

was varied, with eighteen sets of names and photos that were distinct from those of the politicians.¹¹

Third, we vary the partisanship of the politician by displaying a party label under each politician's name. In the US, these parties are the Republican Party and Democratic Party; in the other three multiparty contexts, we select major parties from each bloc and across the ideological spectrum. Respondents whose self-described party ID matches that of the politicians being displayed in a social media conversation is considered an in-partisan, and those indicating a different (or no) party ID are coded as out-partisans.

Fourth, we create sixteen toxic messages that users are shown sending to politicians to ensure that respondents are not repeatedly shown the same text. We develop these messages to allow us to vary whether the text of the message explicitly highlights the gender of the politician. As Rasmussen (2024) shows, the intensity of toxic messages—from mild insults to threats of violence—heavily influences what ordinary citizens perceive as toxic. Language that threatens violence, for instance, is much more likely to be considered toxic than personal insults. However, because our aim is to examine the toxicity that elected representatives face in their day-to-day work environment as politicians, we create a set of messages that are similar in language and tone and mimic the majority of insulting language used daily on social media toward politicians. To develop these texts, we collected replies from ordinary users sent to politicians on Twitter, and classified them by their level of toxicity using Google's Perspective API (Lees et al. 2022; Wulczyn, Thain, and Dixon 2017). We then qualitatively assessed messages classified above the 80th percentile in toxicity, and used these messages to create a set of 16 messages that are similar in tone to these messages that would (1) be plausible responses to ordinary posts from politicians (detailed below), (2) fit cross-culturally across the four countries in which the study is run, and (3) could be edited to highlight the gender of the politician being targeted by the message. Political scientists with domain expertise in each country then translated and commented on whether each message would be plausible in each country's political context to enable the use of the same toxic comments across countries. For the text from each of the 16 messages, we then created neutral and gendered variants. Gendered versions of each post have each message explicitly refer to the politician as a man or a woman. For example, in "Conversation B" of Figure 2, the user refers to the politician as "this man"/"a man," rather than simply as a politician (in "Conversation A").

Fifth, we created texts for ordinary posts from politicians for each social media post. We developed these by collecting tweets from politicians, and using them as

a template for posts on a number of topics that would be plausible cross-culturally. We developed posts from politicians concerning valence issues (rather than directional ones), to ensure that they could be used equivalently across the four country contexts. In total, we created 20 separate posts from politicians that were evenly split on topics concerning the economy, health care, education, crime, and national security. These texts were translated and commented on by political scientists with domain expertise in each country (see Section A.2 of the Supplementary Material).

In total, the number of permutations of all attributes per country ranges between 414,720 and 829,440, (e.g., in the two-party U.S. system, there are 414,720 possible combinations) (details for all treatment conditions and attributes are presented in Section A of the Supplementary Material). As a result, we generate roughly 5 million images of social media conversations, which include all permutations of attributes across the four countries. We then program survey software to randomly select from the set of attributes and display the relevant image. In the survey, each respondent completed five paired conjoint tasks to assess differences in perceptions of toxicity toward politicians, and completed two vignette (i.e., single image) tasks to assess the potential mechanisms. Because each image contains a relatively large amount of information, respondents were first asked to complete the single vignette tasks—which are less complex—after which they completed the paired conjoint tasks.¹²

We note that one of the benefits of using conjoint experiments to investigate gender biases is that they aid in minimizing social desirability bias—a potential problem in studies of sensitive issues, for example, concerning gender. As Horiuchi, Markovich, and Yamamoto (2022) show, social desirability bias is minimized when conjoint designs contain fully randomized attributes (as we do here), that is, when respondents have the possibility of seeing, for example, attacks on two women politicians, two men politicians, or one of each. In the study and experiment, we also note that no cues related to gender were given to respondents, who were told only that the study concerned disrespectful social media behavior. The purpose of the study was given only at the end of the survey when respondents were debriefed on its broader goals. Finally, although recent research suggests that demand effects are, in general, null or minimal in experimental research (Mummolo and Peterson 2019), in the "Results" section, we test whether respondents show evidence of having

¹¹ The set of politicians and users that were to be shown in images to respondents were checked with subject-matter experts for each country in which the experiment was run. See Section A.2 of the Supplementary Material.

¹² After each single vignette conjoint task, respondents were asked a randomly ordered set of questions about the potential motivations of the perpetrator. In Section J of the Supplementary Material, we test whether respondents may be affected by answering these or other questions. We find no evidence that respondents, either politicians or citizens, provide different answers to the first single vignette conjoint task (before they have answered any other question) compared to any subsequent conjoint task.

“learned” the purpose of the experiment and exhibited demand effects by changing their conjoint choices in later tasks. We find no such evidence.

To analyze the paired conjoint experiment, we calculate the Average Marginal Component Effects (AMCE) for each attribute of interest, and Average Marginal Component Interaction Effects (AMCIE) to test our conditional hypotheses (Hainmueller, Hopkins, and Yamamoto 2014). For the paired conjoint design, the outcome variable indicates whether a respondent chose a given image as the more toxic between a pair of images. Because respondents complete multiple paired conjoint tasks, we cluster the standard errors at the level of the respondent. For the second study, which analyzes a single vignette experiment, we use linear regression models with standard errors clustered at the respondent level. Each outcome in the vignette design is a six category Likert scale (“strongly disagree” to “strongly agree”) for each of the seven questions concerning the perceived motivations of the user who is shown sending a toxic message to the politician.

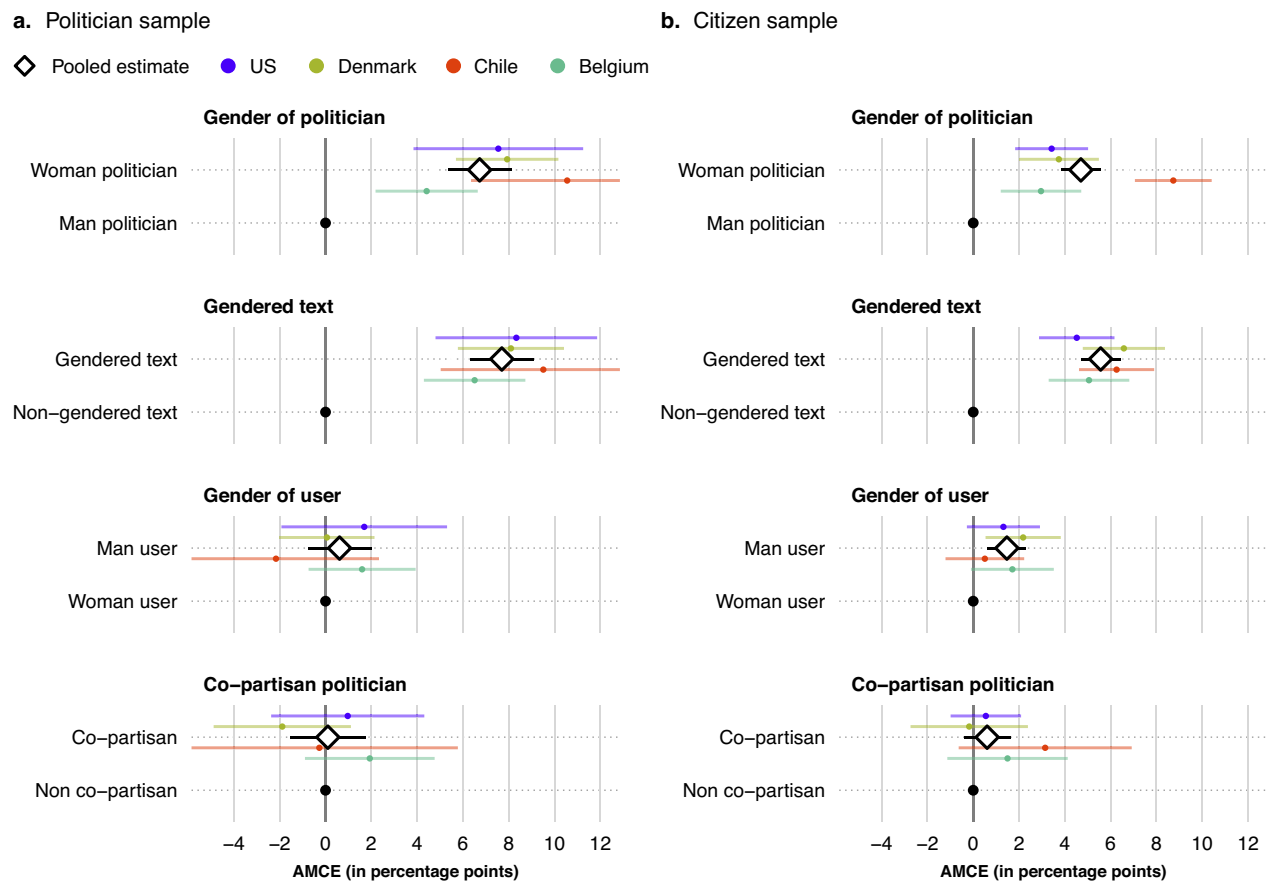
Code and data to reproduce the analysis in this article are available at the American Political Science Review Dataverse (Eady and Rasmussen 2024).

RESULTS

Paired Conjoint Experiment

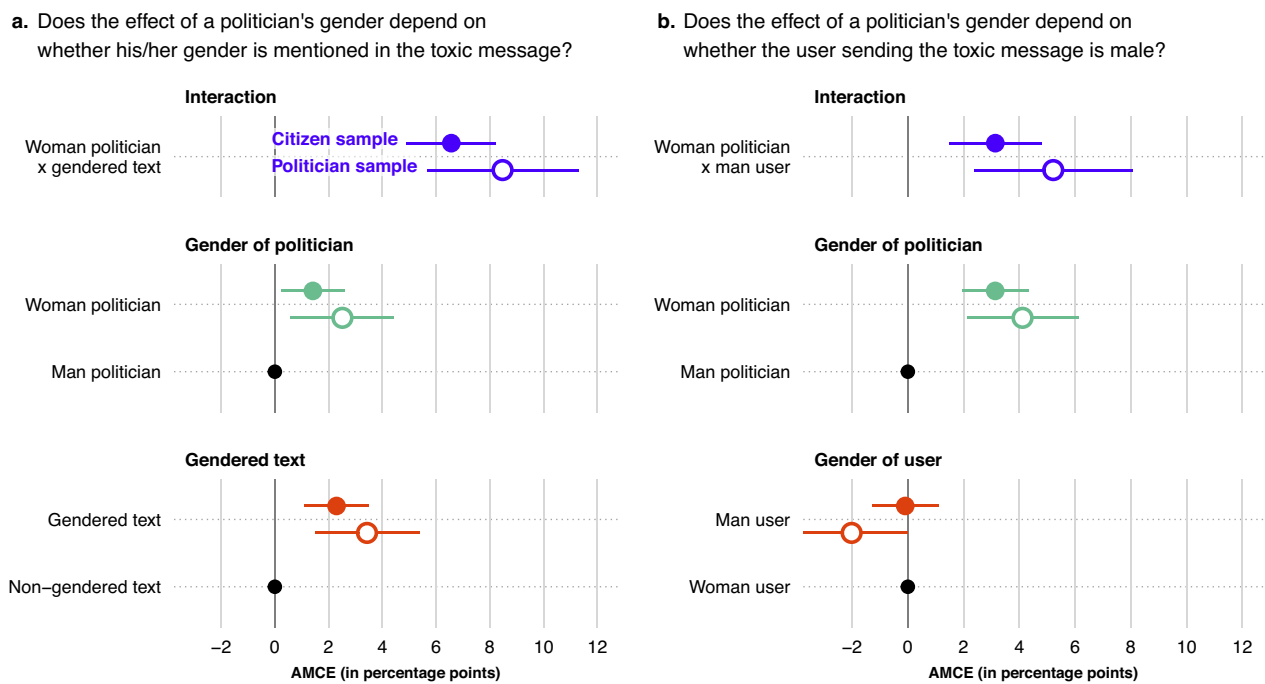
We begin by presenting findings from the paired conjoint design, where respondents select which of two pairs of images depicting social media interactions is more toxic. This allows us to investigate whether the gender of a politician and other attributes of a conversation cause differences in perceptions of the severity of the political hostility directed at a politician. In Figure 3, we present results separately for politician and citizen respondents. The figure shows the extent that a social media conversation is perceived as more toxic than another depending on the treatment conditions of interest. To conserve space, parameter estimates for

FIGURE 3. Effects of Politician’s Gender, User’s Gender, Co-Partisanship, and Gendered Text on Perceptions of Toxicity



Note: This figure displays point estimates and 95% confidence intervals for each attribute of interest, with robust standard errors clustered at the level of the respondent. Effects of each of the 16 toxic message texts and 20 politician message texts not shown. For complete regression results, see Section N of the Supplementary Material. Citizen sample observations = 53,630 (5,371 respondents); politician sample observations = 19,012 (2,153 respondents).

FIGURE 4. Effects of Politician’s Gender Conditional on the User’s Gender and Whether the Text Is Gendered



Note: This figure presents point estimates and 95% confidence intervals for each attribute of interest, and interactions between the politician’s and user’s gender (panel a), and the politician’s gender and whether the text is gendered (panel b). Robust standard errors are clustered at the level of the respondent. Coefficients for each of the 16 toxic message texts and 20 politician message texts not shown. For complete regression results, see Section N of the Supplementary Material. Citizen sample observations = 53,630 (5,371 respondents); politician sample observations = 19,012 (2,153 respondents).

each of the 20 texts from politicians and 16 toxic messages from users are not shown (complete results in Section N of the Supplementary Material). First, as expected from our hypotheses, Figure 3 shows that, independent of a message’s textual content, toxic messages sent to women politicians are perceived as more toxic than those sent to politicians who are men (H1). Pooling the politician and citizen samples, respondents are approximately 6 percentage points more likely to indicate that a conversation is the more toxic one between a pair of conversations if the targeted politician is a woman rather than a man ($p < 0.001$). The effect of the gender of the politician being attacked is stronger for respondents who are politicians—that is, the potential targets of those attacks—than ordinary citizens ($p < 0.05$) (H2).¹³

For completeness, we also show estimates for the other primary attributes in the conjoint experiment. In our preregistration, we did not explicitly make

hypotheses for the main effects of these attributes, because our focus is on how they moderate the effect of a politician’s gender. As the figure demonstrates, we find evidence that toxic messages that note the gender of a politician are perceived as more toxic than those that do not. Averaging across the politician and citizen samples, respondents are 7 percentage points more likely to indicate a conversation is the more toxic one if the message highlights the fact that the politician is a woman or a man ($p < 0.001$). Estimates of the effect of whether the perpetrator is a man (rather than a woman) is positive in both the politician and citizen samples, although only statistically significant among citizens ($p < 0.001$).¹⁴ We find no effect of whether the politician is a co-partisan of the respondent either among the politician ($p = 0.90$) or citizen ($p = 0.24$) sample. We note that the results are similar across countries, with the exception of the effect of a politician’s gender among ordinary citizens in Chile. Given the number of potential comparisons between treatment effects, we do not draw different conclusions specific to Chile, and leave this for further replication or future research.

¹³ In Section J of the Supplementary Material, we test whether the effect of a woman politician is larger the more conjoint tasks a respondent completes. This allows us to investigate whether respondents learn the purpose of the study throughout, and thus respond differently due to potential social desirability or demand effects. We find no evidence that the effect of a woman politician is different in the first conjoint task than that in any later task.

¹⁴ Estimates of the effect of the gender of the user for the politician and citizen sample are not significantly different. As for all tests in this article, statistical power is higher in the citizen sample due to differences in sample sizes.

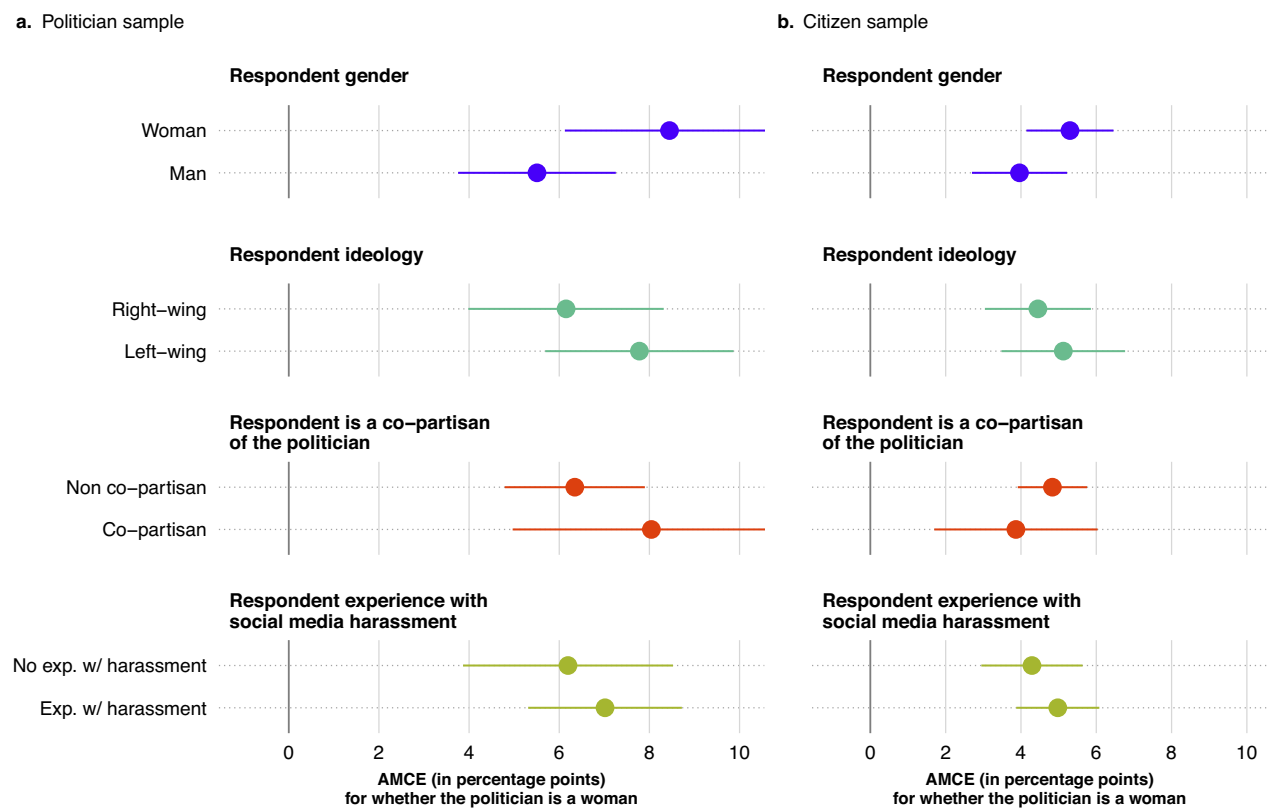
Finally, as a robustness check, in Section G of the Supplementary Material, we examine the main results with the single vignette design, using as an outcome a rating-based measure (0–10 scale) of how toxic a conversation is. Results are consistent with the preregistered paired conjoint design (see Model 1 in Table G14 in the Supplementary Material). In Section K of the Supplementary Material, we also show that the magnitude of the experimental effects for local-level and national-level politician respondents are neither substantively nor statistically significantly different from each other.

We now test hypotheses concerning two conditional relationships. First, we examine whether the effect showing that a message sent to a woman politician is perceived as more toxic than one sent to a man politician is stronger if that message notes the politician's gender (H3). Second, we examine whether the effect of a message to a woman politician is stronger if the perpetrator is a man rather than a woman user (H4). Results from two models testing these conditional effects are presented in Figure 4, which shows both the interaction and component terms. We find empirical support for both hypotheses. When a message is sent to a woman politician, it is roughly 8 percentage points more likely to

be perceived as the more toxic one when that message notes that the politician is a woman ($p < 0.001$) (panel a). Furthermore, toxic messages sent to a woman politician are deemed more toxic if it is sent by a user who is a man rather than a woman: messages sent to women politicians are roughly 4 percentage points more likely to be selected as the more toxic conversation if it is sent by a user who is a man rather than a woman ($p < 0.001$) (panel b). The estimated conditional relationships shown in panels a and b of Figure 4 are similar when examined by each country context separately (see Section E of the Supplementary Material).

Does the effect of whether a politician is a woman or man differ depending on the characteristics of politician and citizen respondents? To test our subgroup hypotheses, we calculate separate estimates for respondent subgroups separated by respondents' gender (H5); political ideology (H6); whether respondents are co-partisans of the politician being attacked (H7); and whether respondents indicate having themselves experienced harassment on social media (H8A/H8B). Results are presented in Figure 5. They show that the magnitude of the effect of whether a toxic message is sent to women instead of men politicians does not meaningfully differ depending on key

FIGURE 5. Effects of Politician's Gender on Perceptions of Toxicity by Respondent Subgroup



Note: This figure presents point estimates and 95% confidence intervals of the effect of the gender of a politician on perceptions of toxic behavior for respondent subgroup characteristics. Robust standard errors are clustered at the level of the respondent. Effects for each of the 16 toxic message texts and 20 politician message texts not shown. For complete regression results, see Section N of the Supplementary Material. Average politician sample observations per subgroup = 9,049 (1,167 respondents); avg. citizen sample observations per subgroup = 24,380 (2,690 respondents).

respondent characteristics (with similar results by subgroup when examined by country contexts, see Section E of the Supplementary Material). Understandings of differences in the severity of political hostility toward women and men politicians, in other words, are widely recognized. Finally, we note that we also test a number of more complex conditional relationships (using triple interaction terms), which for reasons of space, we present in Section I of the Supplementary Material.

Mechanisms

Our results above demonstrate that equivalent messages sent to women politicians are systematically considered more toxic than those sent to men politicians. In this section, we investigate why. To do so, we examine how politicians and citizens understand differences in the motivations that drive the perpetrators of political hostility. We ask, in other words, whether politicians and citizens use heuristics about toxic behavior toward politicians to make inferences about the motivations driving toxic behaviors.

We investigate this by using a second, single-image conjoint experiment. In this experiment, respondents were presented with single images of conversations between a politician and a user and asked questions designed to measure their perceptions of the motivations that drove the perpetrator to send a toxic message. This allows us to examine whether the gender of a politician affects assessments of an array of motivations associated with harmful intentions. As outcomes, we ask the extent that a respondent believes that the perpetrator is driven by prejudice or discrimination; by a desire to discourage the politician from being in politics; by opinion differences with the politician; by dislike of the politician's party; by dissatisfaction with own life; by a desire to get a reaction from the politician; by a desire to get a reaction from other users.

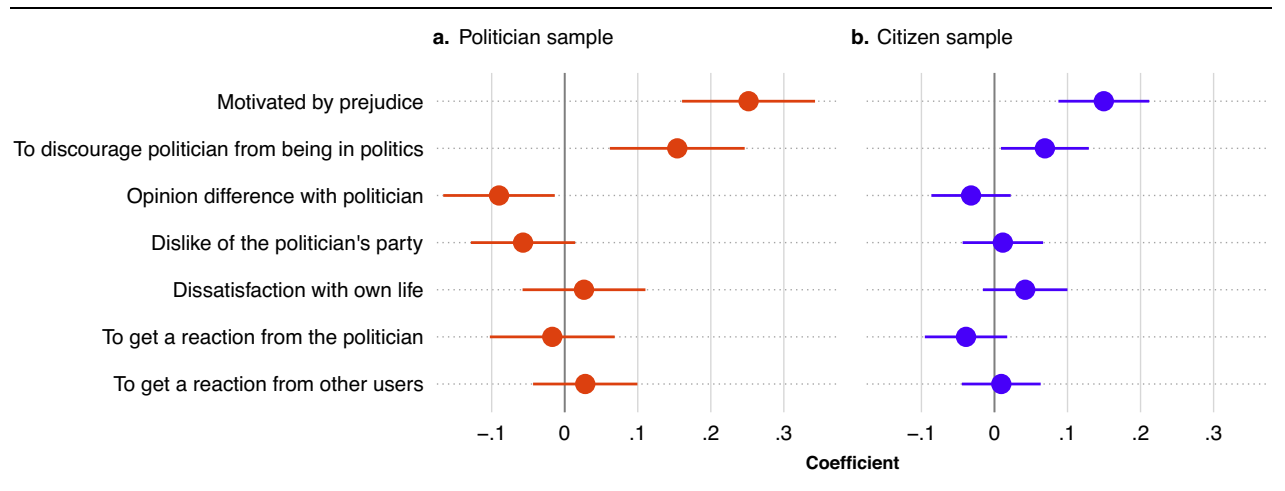
with his/her (the user's) own life; by a desire to get a reaction from the politician (i.e., "trolling"); and by a desire to get a reaction from other users.

As in the first, paired conjoint study, we include in our models all attributes that were randomized within the conversation. We fit an OLS regression model for each outcome, where our quantity of interest is the effect of a politician's gender on respondents' assessments of a perpetrator's motivations. Standard errors are clustered at the respondent level.

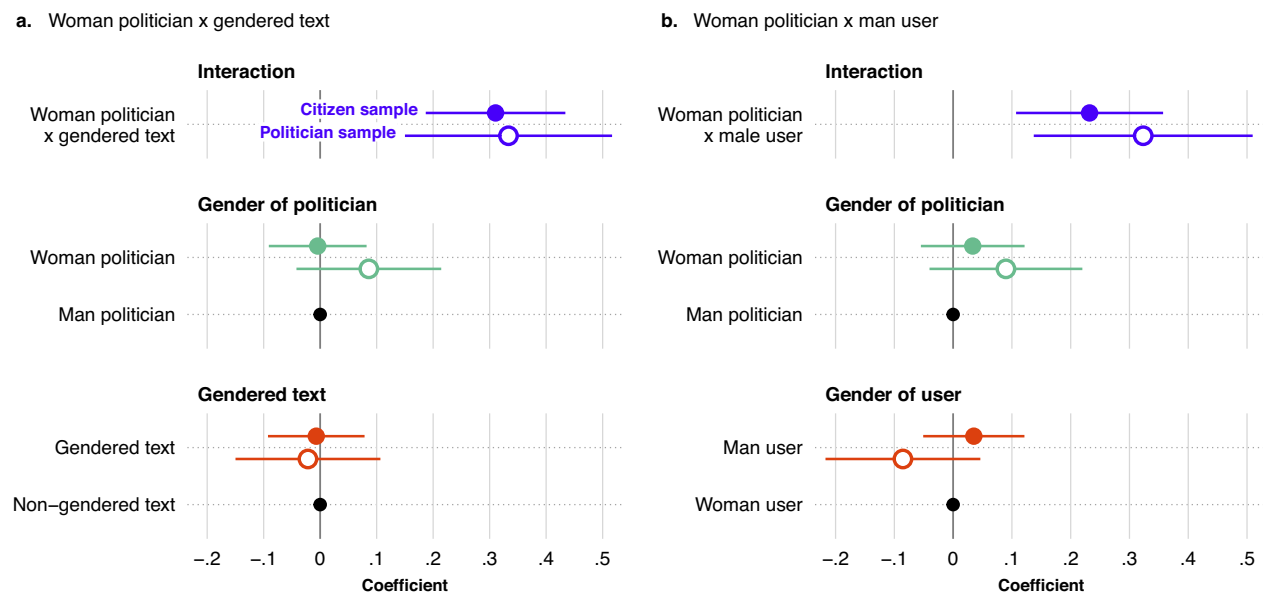
Results for the effect of a politician's gender on each outcome are presented in Figure 6. We calculate effect estimates separately for politician and citizen respondents. As the figure shows, assessments of the motivations of a user who sends a toxic message to a politician depends on whether the targeted politician is a woman or a man, independent of the message content. Users who send toxic messages to women politicians are perceived to be (1) more strongly motivated by prejudice ($p < 0.001$), (2) more driven by a desire to use harassment to discourage a politician from being in politics ($p < 0.01$, politicians; $p < 0.05$, citizens), and (3) less driven by policy differences ($p < 0.05$, politicians; $p = 0.25$, citizens). These results demonstrate that politicians and citizens both recognize that gendered prejudices, not policy differences, tend to drive toxic behavior toward women politicians. We show in Section E of the Supplementary Material that these results are similar when disaggregated by country context.

In Section H of the Supplementary Material, we also examine these results specifically for women and men politicians, that is, those who themselves face toxic behavior, and will interpret the motivations of perpetrators from the perspective of being a woman or man. We show that women politicians themselves interpret the motives of attacks on women politicians as driven

FIGURE 6. Effects of a Politician's Gender on Perceptions of the Motivations behind a Toxic Message



Note: This figure presents point estimates and 95% confidence intervals for the effect of a politician's gender on seven separate outcomes. Each point represents the effect of politician gender on a separate each mechanism, as estimated from seven separate models. Effects for each of the 16 toxic message texts and 20 politician message texts not shown. For complete regression results, see Section N of the Supplementary Material. Robust standard errors are clustered at the level of the respondent, with country fixed effects. Average citizen sample observations = 10,617 (5,331 respondents) per outcome; average politician sample observations = 4,350 (2,223 respondents) per outcome.

FIGURE 7. Effects of Politician's Gender on Perceptions of Whether a User Is Prejudiced Conditional on Whether the Message Is Gendered and the User's Gender

Note: This figure presents point estimates and 95% confidence intervals for each attribute of interest. Effects for each of the 16 toxic message texts and 20 politician message texts not shown. For complete regression results, see Section N of the Supplementary Material. Robust standard errors are clustered at the level of the respondent, with country fixed effect. Citizen sample observations = 10,618 (5,331 respondents); politician sample observations = 4,351 (2,223 respondents).

by prejudice and a desire to push them out of office. We show, furthermore, that the magnitude of the effect of a woman politician facing toxic behavior on perceptions that the perpetrator is driven by prejudices is twice that for women politician respondents than that for men politician respondents ($p = 0.06$). In other words, women politician respondents are more likely to perceive attacks on women are the result of prejudice than men politician respondents. In Section G of the Supplementary Material, we show further that assessments of a user sending a toxic message as being driven by prejudice or a desire to push a politician out of office are positively associated with the extent that respondents perceive a social media interaction as toxic (ratings of the severity of a toxic exchange).

Finally, we investigate these results further by examining whether the extent to which a message sent to a woman politician is seen as driven by prejudice depends on whether the toxic message notes the gender of the politician or is sent by a user who is a man. Results are presented in Figure 7. As would be expected, Panel a shows that prejudice is more likely to be perceived as the motivation of a toxic message that is sent to a woman politician when that message notes the gender of the politician ($p < 0.001$). Panel b shows similarly that the effect of a politician being a woman is magnified when the perpetrator is a man ($p < 0.001$). In Section F of the Supplementary Material, we present analogous results for whether toxic messages sent to women are more likely to be seen as driven by motives to push the politician out of office when sent by a perpetrator who

is a man or if the message notes the gender of the politician. These analyses present weaker evidence of conditional relationships.

CONCLUSION

Hostile political behavior toward politicians is an important societal concern, due to its broad effects on the quality of political discourse, but more critically because it can undermine the quality of representation. Politically toxic behavior is especially problematic if its harm is more severe for politicians who are currently underrepresented in politics, such as women (UN Women 2023). For instance, differences in understandings of toxic behavior might discourage women from entering politics altogether or have physical and mental health consequences for incumbent politicians who might leave politics.

Not surprisingly, researchers have therefore focused on differences in the frequency and content of hostility directed at women and men politicians (e.g., Collignon and Rüdiger 2021; Daniele, Dipoppa, and Pulejo 2023; Håkansson 2021; Herrick and Thomas 2021; Rheault, Rayment, and Musulan 2019, see also Yan and Bernhard 2024). This is an important and growing area of research. In this article, however, we argue that women are faced with a double burden from toxic behaviors. Not only may women politicians be exposed to higher rates of these behaviors, but they must also deal with what those attacks mean about their place as women in

politics. Consequently, the severity and potential costs of these behaviors are a function not only of their content and frequency, but also of the perceived motives and underlying prejudices that fuel them. Much like in many legal regimes where prejudice-driven psychological and physical violence is considered an aggravating factor in assessing the severity of a crime (Jeness 2007), so too do the politicians who may be exposed to it recognize that prejudices aggravate the harms of political toxicity on women politicians. To understand the severity of toxic political behavior, it is thus important to consider how these behaviors are understood both by those who actually experience them—elected representatives—and by members of the wider public who witness them and may consider entering politics.

We investigate this by unpacking how elected politicians and citizens understand toxic behavior toward elected representatives. Using image-based conjoint experiments in the US, Denmark, Belgium, and Chile, we show that when women politicians receive toxic messages, they are widely perceived to be more severe than equivalent messages sent to men. This finding holds regardless of whether respondents are politicians or citizens, with stronger effects among politicians, that is, those best positioned to understand the consequences of these behaviors. These differences in gendered understandings of political hostility are also widely shared across various population subgroups, including those separated by gender, political affiliation, and co-partisanship. For instance, we find no evidence that respondents are more sensitive to attacks on women politicians who might be seen as part of their potential in-group, for example, because they share the harassed politician's gender or political party. Moreover, the difference in the perceived severity of toxic messages sent to women and men politicians is greater when the message is sent by a perpetrator who is a man or references a politician's gender.

Our results underline how the inferred motivations of perpetrators affect assessments of a toxic behavior's severity, regardless of its content. We show that a key reason why such behavior directed at women politicians is perceived as more severe than equivalent behavior toward men is that it is often perceived as being motivated by gender biases rather than by policy disagreements. Empirically, we demonstrate that social media users who send toxic messages to women politicians are perceived as more likely (1) to be driven by prejudice, and (2) to use political hostility as a weapon to push women out of office. Moreover, toxic interactions targeting women are understood as less likely to be driven by policy. Overall, our results illustrate how the motivations that underlie toxic behaviors amplify understandings among politicians and citizens of their harm.

These findings have important implications for understanding the consequences of political toxicity for women politicians. If these perceptions are similar in real-world toxic interactions, it follows that political hostility will have more pronounced detrimental effects on women politicians than their counterparts who are

men. There is substantial anecdotal documentation to suggest this, with numerous instances of women politicians citing the toxic environment associated with being a politician as a reason for their unwillingness to run, or rerun, for office (e.g., Astor 2018; Camut 2023; Hunt, Evershed, and Liu 2016; Lamensch 2021; Mekouar 2019; Morgan 2020; Specia 2019).

Our results also speak to the mechanisms at work in other research. Recent work provides systematic evidence, for example, that facing political violence has differential consequences for women and men. It shows that, among politicians who are subjected to attacks, women mayors are less likely to rerun for elections than mayors who are men and are also the victims of an attack (Daniele, Dipoppa, and Pulejo 2023). The authors, furthermore, present suggestive evidence that the perpetrators of attacks on women mayors are likely driven by gender-based motivations, an observational finding consistent with the mechanisms in our experimental results. Similarly, Yan and Bernhard (2024) show that women who engage in politics are more likely to face toxic behavior than men, and speculate that the potential mechanisms driving these differences in behavior may be related to the internalized sexist attitudes of perpetrators. Our findings add to these studies by tapping into the mechanisms driving such behaviors and showing that the interpretation of the motivations behind them by politicians can exacerbate the severity of these behaviors, even if their frequency or content were similar.

We note that while our research is an important contribution to understanding how politicians themselves interpret the actions of those engaged in toxic behavior, our experiments do not directly estimate the effects of perceptions on the wide array of potential downstream outcomes. Although perceptions of the severity of toxic behavior and the prejudices attached to them may decrease willingness to remain in or run for office, they could also, for example, increase sympathy among the public for those affected. This aligns with evidence from a recent meta-study indicating that women running for office may be advantaged—or, at least, not disadvantaged—compared to men (Schwarz and Coppock 2022). There is scope in future research, therefore, to consider the pathways through which perceptions of toxicity affect those involved. Such work should look into how politicians' interpretations of the motives behind online and offline psychological and physical hostility affect the employment of women politicians and their political decision-making, both quantitatively and qualitatively (see, e.g., Clayton et al. 2020). Research would also benefit from examining the accuracy of politicians' perceptions of toxic behaviors, including whether they underestimate or overestimate the extent to which the behaviors faced by male and female politicians are differentially motivated by gender-based prejudices.

Additionally, there is an opportunity to explore whether animosity aimed at politicians who belong to underrepresented groups other than women (e.g., racial or religious groups) is influenced by comparable prejudices. For example, in Section L of the

Supplementary Material, we show suggestive evidence in our sample of politicians that otherwise equivalent attacks on politicians of color are seen as more toxic than those on white politicians. Moreover, although we find that respondent-level characteristics, such as gender, partisanship, and ideology do not substantially moderate the effect of a politician's gender on perceptions of the severity of toxic behavior, there is also scope to examine whether additional characteristics do. Egalitarian attitudes, often associated with the ideological left, or benevolent sexism, often associated with the ideological right, may nevertheless moderate the effects presented in this article. At the same time, such attitudes should be highly correlated with ideology, which we do not find is a meaningful moderator. Finally, research could test the robustness of our findings with alternative outcome measures beyond disrespect to investigate how the gendered effects of witnessing online toxic behavior compare to other forms of attacks on politicians, such as threats, demonstrations, or acts of physical violence.

In sum, our findings have substantial significance for the study of toxic behavior toward politicians from groups who are underrepresented in politics, and on politically toxic behavior more broadly. They highlight the importance of assessing the quality of democratic dialogue and its darker manifestations by moving beyond analyzing its content and frequency. By considering how perceptions of political interactions between politicians and citizens vary depending on the characteristics of those involved, and their inferred motives, we can make significant strides in this endeavor. Doing so deepens our understanding of the repercussions of political toxicity for underrepresented groups and the potential downstream consequences for political representation.

SUPPLEMENTARY MATERIAL

The supplementary material for this article can be found at <https://doi.org/10.1017/S0003055424001205>.

DATA AVAILABILITY STATEMENT

Research documentation and data that support the findings of this study are openly available at the American Political Science Review Dataverse: <https://doi.org/10.7910/DVN/B3UBXR>. The experimental design, hypotheses, and primary analytical strategies were pre-registered at: <https://osf.io/q3p2a>.

ACKNOWLEDGMENTS

We thank Daniel Cruz, as well as Zeljko Poljak and Karolin Soontjens for assisting with our elite surveys in Chile and Belgium, respectively, and Nathan Lee and his team at CivicPulse for implementing our U.S. elite survey. We are also grateful to Dynata for project

assistance in running our citizen surveys. Earlier versions of this article were presented at the 2022 Annual Meeting of the Danish Political Science Association, the 2023 Political Economy departmental seminar at King's College London, and invited presentations at the University of Aarhus (2022), University Carlos III in Madrid (2023), the University of Gothenburg (2023), the University of Oslo (2022), and Queen's University Belfast (2024). We thank all participants, particularly Elin Allern, Michael Bang Petersen, Simon Chauchard, Michele Crepaz, Alberto Díaz-Cayeros, Marco Giani, Anna Gwiazda, Felix Haass, Vibeke Wøien Hansen, Shaun Hargreaves Heap, Christel Koop, Krzysztof Krakowski, Sandra León, Julian Limberg, Jakob Nyrup, Elias Markstedt, Juan Mayoral, Mikael Persson, Pedro Riera, Jesper Rasmussen, Lise Rødland, Moses Shayo, Rasmus Skytte, Anders Sundell, Bouke Klein Teeselink, as well as Lior Scheffer, the APSR editors, and four anonymous reviewers for their valuable feedback and suggestions.

FUNDING STATEMENT

The experimental survey was funded by the Danish Research Council for Independent Research to the authors under grant number 0133-00034B.

CONFLICT OF INTEREST

The authors declare no ethical issues or conflicts of interest in this research.

ETHICAL STANDARDS

The authors declare the human subjects research in this article was reviewed and approved by the Ethics Committee of the Department of Political Science, Copenhagen University under protocol number 2022-05. The authors affirm that this article adheres to the principles concerning research with human participants laid out in APSA's Principles and Guidance on Human Subject Research (2020).

REFERENCES

- ACE Electoral Knowledge Network. 2023. "Barriers to Women's Effective Participation in the Electoral Process." <https://aceproject.org>.
- Adida, Claire L., Adeline Lo, and Melina R. Platas. 2018. "Perspective Taking Can Promote Short-Term Inclusionary Behavior toward Syrian Refugees." *Proceedings of the National Academy of Sciences* 115 (38): 9521–6.
- Amnesty International. 2018. "Toxic Twitter: A Toxic Place for Women." Amnesty International Report.
- Astor, Maggie. 2018. "For Female Candidates, Harassment and Threats Come Every Day." *New York Times*, August 24.
- Ballington, Julie. 2018. "Turning the Tide on Violence against Women in Politics: How Are We Measuring Up?" *Politics & Gender* 14 (4): 695–701.

- Banwart, Mary Christine. 2010. "Gender and Candidate Communication: Effects of Stereotypes in the 2008 Election." *American Behavioral Scientist* 54 (3): 265–83.
- Bardall, Gabrielle, Elin Bjarnegård, and Jennifer M. Piscopo. 2020. "How is Political Violence Gendered? Disentangling Motives, Forms, and Impacts." *Political Studies* 68 (4): 916–35.
- Bauer, Nichole M. 2019. "Shifting Standards: How Voters Evaluate the Qualifications of Female and Male Candidates." *Journal of Politics* 82 (1): 1–12.
- Berevoescu, Ionica, and Julie Ballington. 2021. "Women's Representation in Local Government: A Global Analysis." UN Women Report.
- Bernhard, Rachel, Shauna Shames, and Dawn Langan Teele. 2021. "To Emerge? Breadwinning, Motherhood, and Women's Decisions to Run for Office." *American Political Science Review* 115 (2): 379–94.
- Bjarnegård, Elin. 2013. *Gender, Informal Institutions and Political Recruitment: Explaining Male Dominance in Parliamentary Representation*. London: Palgrave Macmillan.
- Bjarnegård, Elin. 2023. "The Continuum of Election Violence: Gendered Candidate Experiences in the Maldives." *International Political Science Review* 44 (1): 107–21.
- Bolsen, Toby, James N. Druckman, and Fay Lomax Cook. 2014. "The Influence of Partisan Motivated Reasoning on Public Opinion." *Political Behavior* 36 (2): 235–62.
- Bor, Alexander, and Gabor Simonovits. 2021. "Empathy, Deservingness, and Preferences for Welfare Assistance: A Large-Scale Online Perspective-Taking Experiment." *Political Behavior* 43: 1247–64.
- Borah, P. 2013. "Interactions of News Frames and Incivility in the Political Blogosphere: Examining Perceptual Outcomes." *Political Communication* 30 (3): 456–73.
- Camut, Nicolas. 2023. "Top Dutch Minister Steps Down from Party Leadership over 'Intimidation' and 'Threats'." *Politico*, July 13.
- Carnevale, Anthony P., Nicole Smith, and Kathryn Peltier Campbell. 2019. "May the Best Woman Win? Education and Bias against Women in American Politics." Georgetown University Center on Education and the Workforce (CEW) Report.
- Clayton, Amanda, Amanda Lea Robinson, Martha C. Johnson, and Ragnhild Muriaas. 2020. "(How) Do Voters Discriminate against Women Candidates? Experimental and Qualitative Evidence from Malawi." *Comparative Political Studies* 53 (3–4): 601–30.
- Collignon, Sofia, Rosie Campbell, and Wolfgang Rüdiger. 2022. "The Gendered Harassment of Parliamentary Candidates in the UK." *The Political Quarterly* 93 (1): 32–8.
- Collignon, Sofia, and Wolfgang Rüdiger. 2021. "Increasing the Cost of Female Representation? The Gendered Effects of Harassment, Abuse and Intimidation towards Parliamentary Candidates in the UK." *Journal of Elections, Public Opinion and Parties* 31 (4): 429–49.
- Craig, Stephen C., and Paulina S. Cossette. 2022. "Eye of the Beholder: Partisanship, Identity, and the Politics of Sexual Harassment." *Political Behavior* 44 (2): 749–77.
- Daniele, Gianmarco, Gemma Dipoppa, and Massimo Pulejo. 2023. "Attacking Women or their Policies? Understanding Violence against Women in Politics." CEPR Press Discussion Paper No. 18333.
- Darcy, Robert, Susan Welch, and Janet Clark. 1994. *Women, Elections & Representation*, 2nd edition. Lincoln: University of Nebraska Press.
- Dardenne, Benoit, Muriel Dumont, and Thierry Bollier. 2007. "Insidious Dangers of Benevolent Sexism: Consequences for Women's Performance." *Journal of Personality and Social Psychology* 93 (5): 764.
- Dhrodia, Azmina. 2017. "Unsocial Media: Tracking Twitter Abuse against Women MPs." Amnesty International Global Insights Report.
- Dittmar, Kelly, Kira Sanbonmatsu, and Susan J. Carroll. 2018. *A Seat at the Table: Congresswomen's Perspectives on Why Their Presence Matters*. Oxford: Oxford University Press.
- Dolan, Kathleen. 1998. "Voting for Women in the 'Year of the Woman'." *American Journal of Political Science* 42 (1): 272–93.
- Dolan, Kathleen. 2008. "Is There a 'Gender Affinity Effect' in American Politics? Information, Affect, and Candidate Sex in U.S. House Elections." *Political Research Quarterly* 61 (1): 79–89.
- Dover, Tessa L., Jeffrey M. Hunger, and Brenda Major. 2021. "Health Consequences of Prejudice and Discrimination." In *The Wiley Encyclopedia of Health Psychology*, eds. Kate Sweeny, Megan L. Robbins, and Lee M. Cohen, 231–8. Hoboken, NJ: John Wiley & Sons.
- Eady, Gregory, and Anne Rasmussen. 2024. "Replication Data for: Gendered Perceptions and the Costs of Political Toxicity: Experimental Evidence from Politicians and Citizens in Four Democracies." Harvard Dataverse. Dataset. <https://doi.org/10.7910/DVN/B3UBXR>.
- Every-Palmer, Susanna, Justin Barry-Walsh, and Michele Pathé. 2015. "Harassment, Stalking, Threats and Attacks Targeting New Zealand Politicians: A Mental Health Issue." *Australian & New Zealand Journal of Psychiatry* 49 (7): 634–41.
- Fox, Richard L. 2005. "Congressional Elections: Where Are We on the Road to Gender Parity?" In *Gender and Elections: Shaping the Future of American Politics*, eds. Susan J. Carroll and Richard L. Fox, 97–116. New York: Cambridge University Press.
- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22 (1): 1–30.
- Håkansson, Sandra. 2021. "Do Women Pay a Higher Price for Power? Gender Bias in Political Violence in Sweden." *Journal of Politics* 83 (2): 515–31.
- Heath, Roseanna Michelle, Leslie A. Schwindt-Bayer, and Michelle M. Taylor-Robinson. 2005. "Women on the Sidelines: Women's Representation on Committees in Latin American Legislatures." *American Journal of Political Science* 49 (2): 420–36.
- Herrick, Rebekah, and Lori D. Franklin. 2019. "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): 2047–58.
- Herrick, Rebekah, and Sue Thomas. 2021. "Not Just Sticks and Stones: Psychological Abuse and Physical Violence among U.S. State Senators." *Politics & Gender* 18 (2): 422–47.
- Herrick, Rebekah, Sue Thomas, Lori Franklin, Marcia L. Godwin, Eveline Gnabasi, and Jean Reith Schroedel. 2019. "Physical Violence and Psychological Abuse against Female and Male Mayors in the United States." *Politics, Groups, and Identities* 9 (4): 681–98.
- Horiuchi, Yusaku, Zachary Markovich, and Teppei Yamamoto. 2022. "Does Conjoint Analysis Mitigate Social Desirability Bias?" *Political Analysis* 30 (4): 535–49.
- Hunt, Elle, Nick Evershed, and Ri Liu. 2016. "From Julia Gillard to Hillary Clinton: Online Abuse of Politicians around the World." *The Guardian*, June 27.
- Inter-Parliamentary Union. 2016. "Sexism, Harassment and Violence against Women Parliamentarians." Inter-Parliamentary Union Issues Brief.
- Inter-Parliamentary Union. 2024. "Monthly Ranking of Women in National Parliaments." Inter-Parliamentary Union Report, January.
- Jenness, Valerie. 2007. "The Emergence, Content, and Institutionalization of Hate Crime Law: How a Diverse Policy Community Produced a Modern Legal Fact." *Annual Review of Law and Social Science* 3: 141–60.
- Karpowitz, Christopher, and Tali Mendelberg. 2014. *The Silent Sex: Gender, Deliberation, and Institutions*. Princeton, NJ: Princeton University Press.
- Kenski, Kate, Kevin Coe, and Stephen A. Rains. 2017. "Perceptions of Uncivil Discourse Online: An Examination of Types and Predictors." *Communication Research* 47 (6): 795–814.
- King, Amy, and Andrew Leigh. 2010. "Bias at the Ballot Box? Testing Whether Candidates' Gender Affects Their Vote." *Social Science Quarterly* 91 (2): 324–43.
- Krook, Mona Lena. 2020. *Violence against Women in Politics*. New York: Oxford University Press.
- Krook, Mona Lena, and Juliana Restrepo Sanín. 2019. "The Cost of Doing Politics? Analyzing Violence and Harassment against Female Politicians." *Perspectives on Politics* 18 (3): 740–55.
- Kubn, Emily, Curtis Puryear, Chelsea Schein, and Kurt Gray. 2021. "Personal Experiences Bridge Moral and Political Divides Better than Facts." *Proceedings of the National Academy of Sciences* 118 (6): 1–9.
- Lamensch, Marie. 2021. "When Women are Silenced Online, Democracy Suffers." Centre for International Governance Innovation, April 21.

- Lawless, Jennifer L., and Richard L. Fox. 2005. *It Takes a Candidate: Why Women Don't Run for Office*. New York: Cambridge University Press.
- Lawless, Jennifer L., and Richard L. Fox. 2010. *It Still Takes a Candidate: Why Women Don't Run for Office*. New York: Cambridge University Press.
- Lazarus, Jeffrey, and Amy Steigerwalt. 2018. *Gendered Vulnerability: How Women Work Harder to Stay in Office*. Ann Arbor: University of Michigan Press.
- Lees, Alyssa, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Gupta, Donald Metzler, and Lucy Vasserman. 2022. "A New Generation of Perspective API: Efficient Multilingual Character-Level Transformers." Perspective API Technical Report, February 22.
- Maisel, L. Sandy. 2012. "The Negative Consequences of Uncivil Political Discourse." *PS: Political Science & Politics* 45 (3): 405–11.
- McDevitt, Jack, Jennifer Balboni, Luis Garcia, and Joann Gu. 2001. "Consequences for Victims: A Comparison of Bias- and Non-Bias-Motivated Assaults." *American Behavioral Scientist* 45 (4): 697–713.
- McGregor, Shannon C. 2020. "Taking the Temperature of the Room": How Political Campaigns Use Social Media to Understand and Represent Public Opinion." *Public Opinion Quarterly* 84 (S1): 236–56.
- Mekouar, Dora. 2019. "For Women in Politics, Being Terrorized Comes with the Job." *Voice of America*, August 23.
- Mo, Cecilia Hyunjung. 2015. "The Consequences of Explicit and Implicit Gender Attitudes and Candidate Quality in the Calculations of Voters." *Political Behavior* 37 (2): 357–95.
- Morgan, Dafydd. 2020. "Women in Politics Face 'Daily' Abuse on Social Media." *BBC News*, May 26.
- Mummolo, Jonathan, and Erik Peterson. 2019. "Demand Effects in Survey Experiments: An Empirical Assessment." *American Political Science Review* 113 (2): 517–29.
- Norris, Pippa, and Ronald Inglehart. 2019. *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*. Cambridge: Cambridge University Press.
- OECD. 2024. "Violence against Women Indicators." Organization for Economic Co-operation and Development.
- Okimoto, Tyler G., and Victoria L. Brescoll. 2010. "The Price of Power: Power Seeking and Backlash against Female Politicians." *Personality and Social Psychology Bulletin* 36 (7): 923–36.
- Pascoe, Elizabeth A., and Laura Smart Richman. 2009. "Perceived Discrimination and Health: A Meta-Analytic Review." *Psychological Bulletin* 135 (4): 531.
- Plutzer, Eric, and John F. Zipp. 1996. "Identity Politics, Partisanship, and Voting for Women Candidates." *Public Opinion Quarterly* 60 (1): 30–57.
- Pratto, Felicia, Lisa M. Stallworth, and Jim Sidanius. 1997. "The Gender Gap: Differences in Political Attitudes and Social Dominance Orientation." *British Journal of Social Psychology* 36 (1): 49–68.
- Rasmussen, Jesper. 2024. "The (Limited) Effects of Target Characteristics on Public Opinion of Hate Speech Laws." Unpublished Manuscript.
- Rheault, Ludovic, Erica Rayment, and Andreea Musulan. 2019. "Politicians in the Line of Fire: Incivility and the Treatment of Women on Social Media." *Research & Politics* 6(1). <https://doi.org/10.1177/2053168018816228>
- Sanbonmatsu, Kira. 2002. "Gender Stereotypes and Vote Choice." *American Journal of Political Science* 46 (1): 20–34.
- Sandberg, Linn, and Anamaria Duceac Segesten. 2022. "Online Political Hostility." In *Elgar Encyclopedia of Technology and Politics*, ed. Andrea Ceron, 244–8. Cheltenham, UK: Edward Elgar Publishing.
- Schmader, Toni, Alyssa Croft, Marchelle Scarnier, Brian Lickel, and Wendy Berry Mendes. 2012. "Implicit and Explicit Emotional Reactions to Witnessing Prejudice." *Group Processes & Intergroup Relations* 15 (3): 379–92.
- Schwarz, Susanne, and Alexander Coppock. 2022. "What Have We Learned about Gender from Candidate Choice Experiments? A Meta-Analysis of 42 Factorial Survey Experiments." *Journal of Politics* 84 (2): 655–68.
- Smith, Eliot R. 1993. Social Identity and Social Emotions: Toward New Conceptualizations of Prejudice. In *Affect, Cognition and Stereotyping*, eds. Diane M. Mackie and David L. Hamilton, 297–315. San Diego, CA: Academic Press.
- Smith, Jessi L., David Paul, and Rachel Paul. 2007. "No Place for a Woman: Evidence for Gender Bias in Evaluations of Presidential Candidates." *Basic and Applied Social Psychology* 29 (3): 225–33.
- Specia, Megan. 2019. "Threats and Abuse Prompt Female Lawmakers to Leave U.K. Parliament." *New York Times*, November 1.
- Stryker, Robin, Bethany Anne Conway, and J. Taylor Danielson. 2016. "What Is Political Incivility?" *Communication Monographs* 83 (4): 535–56.
- Tajfel, Henri. 1981. *Human Groups and Social Categories*. Cambridge: Cambridge University Press.
- Tajfel, Henri, ed. 1978. *Differentiation between Social Groups: Studies in the Social Psychology of Intergroup Relations*. London: Academic Press.
- Teele, Dawn Langan, Joshua Kalla, and Frances Rosenbluth. 2018. "The Ties that Double Bind: Social Roles and Women's Underrepresentation in Politics." *American Political Science Review* 112 (3): 525–41.
- Terris, Ben. 2016. "Youngest Female Lawmaker Says Social Media is Just Part of the Job." *New York Times*, September 13.
- Theocharis, Yannis, Pablo Barberá, Zoltán Fazekas, Sebastian Adrian Popa, and Olivier Parnet. 2016. "A Bad Workman Blames His Tweets: The Consequences of Citizens' Uncivil Twitter Use When Interacting with Party Candidates." *Journal of Communication* 66 (6): 1–25.
- Theocharis, Yannis, Pablo Barberá, Zoltán Fazekas, Sebastian Adrian Popa, and Olivier Parnet. 2020. "The Dynamics of Political Incivility on Twitter." *Sage Open* 10 (2): 1–15.
- Thomsen, Danielle M. 2020. "Ideology and Gender in U.S. House Elections." *Political Behavior* 42 (2): 415–42.
- Thomsen, Danielle M., and Aaron S. King. 2020. "Women's Representation and the Gendered Pipeline to Power." *American Political Science Review* 114 (4): 989–1000.
- Tromble, Rebekah. 2018. "Thanks for (Actually) Responding! How Citizen Demand Shapes Politicians' Interactive Practices on Twitter." *New Media & Society* 20 (2): 676–97.
- Tucker, Joshua A., Yannis Theocharis, Margaret E. Roberts, and Pablo Barberá. 2017. "From Liberation to Turmoil: Social Media and Democracy." *Journal of Democracy* 28 (4): 46–59.
- UN Women. 2023. "Facts and Figures: Women's Leadership and Political Participation." UN Women.
- UNDP. 2024. "Human Development Reports." United Nations Development Programme.
- Vecchiato, Alessandro, and Kevin Munger. 2022. "Validating the Visual Conjoint, with an Application to Candidate Evaluation on Social Media." Unpublished Manuscript.
- Vial, Andrea C., Jaime L. Napier, and Victoria L. Brescoll. 2016. "A Bed of Thorns: Female Leaders and the Self-Reinforcing Cycle of Illegitimacy." *The Leadership Quarterly* 27 (3): 400–14.
- Wulczyn, Ellery, Nithun Thain, and Lucas Dixon. 2017. "Ex Machina: Personal Attacks Seen at Scale." International World Wide Web Conference Committee (IW3C2), April 3–7, Perth, Australia.
- Yan, Alan N., and Rachel Bernhard. 2024. "The Silenced Text: Field Experiments on Gendered Experiences of Political Participation." *American Political Science Review* 118 (1): 481–7.