

# Male and Female Politicians on Twitter: A Machine Learning Approach

Javier Beltran <sup>1</sup>

Aina Gallego <sup>2</sup>

Alba Huidobro <sup>3</sup>

Enrique Romero <sup>4</sup>

Lluís Padró <sup>5</sup>

## Abstract

How does the language of male and female politicians differ when they communicate directly with the public on social media? Do citizens address them differently? We apply Lasso logistic regression models to identify the linguistic features that most differentiate the language used by or addressed to male and female Spanish politicians. Male politicians use more words related to politics, sports, ideology, and infrastructure, while female politicians talk about gender and social affairs. The choice of emojis varies greatly across genders. In a novel analysis of tweets written by citizens, we find evidence of gender-specific insults, and note that mentions of physical appearance and infantilizing words are disproportionately found in text addressed to female politicians. The results suggest that politicians conform to gender stereotypes online and reveal ways in which citizens treat politicians differently depending on their gender.

**Keywords:** Twitter, gender differences, politicians, machine learning, social media

---

<sup>1</sup> Universitat de Barcelona. E-mail: [javier.beltran@ub.edu](mailto:javier.beltran@ub.edu)

<sup>2</sup> Institut Barcelona d'Estudis Internacionals and Institute of Political Economy and Governance. E-mail: [agallego@ibeio.org](mailto:agallego@ibeio.org)

<sup>3</sup> Institut Barcelona d'Estudis Internacionals and Universitat Pompeu Fabra. E-mail: [alba.huidobro@upf.edu](mailto:alba.huidobro@upf.edu)

<sup>4</sup> Universitat Politècnica de Catalunya. E-mail: [eromero@cs.upc.edu](mailto:eromero@cs.upc.edu)

<sup>5</sup> Universitat Politècnica de Catalunya. E-mail: [padro@cs.upc.edu](mailto:padro@cs.upc.edu)

# 1 Introduction

New social media, including Twitter, offer politicians an unprecedented opportunity to communicate directly with large numbers of citizens. When politicians have full control over the messages they send to the public, do they communicate in ways that conform to gender stereotypes? This research note examines differences in the topics discussed and the language used by male and female politicians on Twitter. Besides providing abundant data on the messages written by politicians, social media data also offers novel but still underutilized opportunities to study whether citizens communicate with male and female politicians differently. Hence, we also ask whether citizens speak to politicians of both genders in different ways and if so, how.

To address these questions, we collected around 121,000 messages written by Spanish national and regional politicians on Twitter for six months, and 207,000 messages addressed to them. We use a machine learning algorithm to identify the tokens most associated with gender. We refer to these tokens as ‘male-linked’ or ‘female-linked’ (Park et al. 2016). Our findings suggest that when politicians communicate with the public they reproduce gender stereotypes in both content and style. Stems relating to ideology, transportation, and infrastructure and factual emojis are male-linked, while emotional emojis and stems relating to gender and social issues are female-linked.

Our empirical approach, combined with data on responses to messages posted by politicians, allows us to make new discoveries about how citizens address male and female politicians differently. The analysis reveals extensive evidence of sexist speech addressed at female MPs: citizens are more likely to address them using infantilizing words or words about their physical appearance. Both genders receive insults and demeaning words, but the insults addressed at women are disproportionately related to their appearance or to their feminist positions.

This research note contributes to previous work on gender differences in candidates’ communication style on Twitter (Evans and Clark 2016) and the integration of Twitter into campaigns and its electoral effects (Wagner et al. 2017), which codes text manually. A growing body of literature applies machine learning to analyse text from other sources such as media portrayals (e.g., Aldering and Pas 2018). Until now, however, and in spite of its obvious potential, the use of machine learning methods to analyse gendered political communication on Twitter has been limited, as we discuss in more detail below.

Our work also adds to the growing political science literature that draws on Twitter data to study topics like online racism (Munger 2017), homophily in political discussion (Barberá 2015, Barberá et al. 2015), the determinants of political protest (Jost et al. 2018), and vote shifts across

districts and electoral results (Beauchamp 2017). It is also closely related to the recent studies by Wu (2018), who analyses differences in the words being used when talking about male and female economists in the Econ Job Rumors forum using Lasso logistic models, and by Bohren et al. (2018), who study whether comments posted on a mathematics forum vary depending on the user's gender. We use similar machine learning techniques but apply them to a corpus that is much more relevant to the study of politics.

Finally, the majority of existing studies about politicians on Twitter have been conducted in English-speaking countries and consequently the most widely used analytical tools are also developed in English. We analyse political texts written in Spanish, the language with the second-largest number of native speakers (after Chinese). By sharing the steps we took to preprocess and analyse the text, we hope to help expand the use of machine learning methods in political communication beyond English-speaking contexts.

## 2 Gender differences in political communication

Do male and female politicians talk differently when they communicate directly with the public? If so, how? Previous studies about these questions using social media data mostly confirmed that politicians talk about different issues and use different tones. Bailey & Nawara (2018) conducted a content analysis of every tweet, Facebook post, and Instagram image posted by Donald Trump and Hillary Clinton in the run-up to the 2016 presidential election. While Clinton posted more positive and negative content overall and talked more about gender issues, race, and immigration, Trump posted more about crime. Also focusing on Twitter, Just & Crigler (2014) found that female candidates displayed emotions of hope and enthusiasm in their messages, while men conveyed disgust and anger. Yarchi & Samuel-Azran (2018) report that female politicians are more likely to use Facebook to attack their opponents or to talk about gender issues and men are more likely to talk about security, economy, and welfare.

Most of this research has focused on electoral campaigns in the United States and it is unclear if the findings extend to other contexts. Moreover, most studies with social media data rely on hand coding and do not use automated processing of text and machine learning methods, which limits the amount of data that can be included in any given study. We contribute to this body of research by focusing on regular communication outside the campaign period on Twitter and in a context outside the United States, and by using machine learning methods.

Other relevant evidence on gender differences in online language use comes from outside the political domain. Research on gender differences in language use suggests that women use more

emotional, social, and warmer language, while men use more descriptive and impersonal language (Newman et al. 2008, Palomares 2008, Park et al. 2016). Studies on computer-mediated communication have found strong gender differences in the use of emojis (Wolf 2000, Chen et al. 2018, Hwang 2014), which motivates us to retain emojis when preprocessing text.

The second question we ask is whether citizens address male and female politicians differently on social media. Research into this issue is much scarcer. Previous studies suggest that voters view and evaluate candidates from the perspective of gender stereotypes (e.g. Dolan & Lynch 2014, Lawless 2004, Jamieson 1995, Herrnson et al. 2003, Ditonto et al. 2014, Bauer 2014, 2018, Cassese & Holman 2018, Teele et al. 2018), but they focus on attitudes and voting behaviour rather than actual communication from citizens to politicians. Actual communication is interesting because politicians experience it directly and it can possibly affect their communication decisions. Text data obtained from Twitter is a novel way to explore how citizens communicate with politicians.

Besides observing which issues and emotion-related words are linked to the gender of the receiver, one particular concern relates to studying the hostile experiences of politicians online, and whether these experiences vary by gender. The anonymity of Twitter may be particularly conducive to uncivil behaviour (Munger 2017). There is evidence that politicians frequently receive hostile messages (Theocharis et al. 2016) and that violence against women in politics is manifested through online threats (Krook 2017). We thus pay particular attention to differences in hostile, aggressive, and demeaning language. Such messages may or may not be disproportionately addressed to women politicians. Recent evidence on media coverage does not, in fact, support the view that women face a more hostile campaign environment than men (Lawless 2015, Rheault et al. 2019). An alternative possibility, one that is consistent with theories of benign sexism, is that women politicians may be less likely to be taken seriously and more likely to be addressed using condescending and infantilizing language.

### 3 Data and method

#### 3.1 Text corpus

To investigate gender differences in political speech, we searched for the Twitter profiles of all national and regional MPs in Spain and found 1,221 Twitter accounts, which represent 85% of the MPs for the Spanish parliament (*Congreso de los Diputados*), and 79% of regional parliamentarians. We then collected two corpora of text. First, we gathered 121,316 tweets posted by legislators between 18 December 2017 and 4 June 2018. Second, we collected 207,574

tweets that replied to these politicians during the same period through the Standard Search API. We ran the search weekly for 24 weeks.

We identify the gender of all politicians from their names. Table 1 shows descriptive statistics about the politicians included in this study and their tweets, grouped by the politicians' gender.

Table 1 about here

The data collection process was complicated by the fact that Spanish is not the only language spoken in Spain. But because multilingual interactions are extremely common (e.g. a Catalan parliamentary speaking in Catalan and an individual answering in Spanish) and because Spanish, Catalan, and Galician have plenty of common features, we decided to include all the tweets in our analysis.

### 3.2 Data processing

We preprocessed the text to reduce variability and transformed it into a database (Grimmer & Stewart 2013, Wilkerson & Casas 2017). This preprocessing used the NLTK library, which provides tools for processing Spanish. We first applied a tokenizer that splits tweets into a list of tokens. We then applied stemming to transform the tokens obtained into linguistic stems. Reducing the complexity of text is important when dealing with Spanish, where gender, number, and verbal mode, among other features, modify the stem of the word. For this step, we used the Snowball Stemmer from NLTK. We discarded certain types of tokens, namely Twitter mentions, hashtags, hyperlinks, and words that occur infrequently. Because the complexity of the model is proportional to the number of words represented in it, we only kept the 5000 most frequent stemmed tokens. We did not remove stop words from this list because they can be informative of gender differences.

The resulting text, which is a sequence of stems, was transformed into a bag-of-words representation as described by Grimmer and Stewart (2013). The whole corpus is represented as a matrix where each stemmed token appearing in the corpus is a column (a variable or feature) and each tweet is a row (an observation). The content of each cell corresponds to the frequency of a stemmed token in a tweet. One limitation of this approach is that the order in which the words appear in the text is lost. On the other hand, it allows us to interpret the model clearly because the features coincide with the stemmed tokens (including words and emojis). In order to evaluate the fit of our model, we reserved 20% (randomly sampled) of the corpus as a test set, and the rest is used for training the models.

### 3.3 Models

We fitted Lasso logistic regressions that regress the gender of the politician on the stemmed tokens written by them (including words and emojis), and a second set of models regressing gender on the tokens addressed to them. Intuitively, these models test if some words and other features such as emojis are disproportionately used by male or female politicians, that is if they are relatively more frequently used by males. This measure is more meaningful than merely testing the words used most frequently by both genders, because these are generally uninformative and ineffective in distinguishing between men and women. We choose to analyse the data using Lasso models rather than the topic models that are used more frequently in Political Science (Grimmer & Stewart 2013, Roberts et al. 2014, Lucas et al. 2015, Wilkerson & Casas 2017) because we are interested not only in the topics discussed, but also in other features such as emojis and emotional expression. A focus on words and other features allows us to examine nuances (e.g. the tone of insults addressed at male and female politicians) that would be lost in topic models.

Lasso logistic regression is a probabilistic classifier that estimates a coefficient for each of the words in the bag-of-words representation presented above. Unlike results from other machine learning methods, these coefficients are easy to interpret: the words with the largest coefficients are the most ‘predictive’ words or are linked to the male or female categories, depending on their sign. A regularizer is used to prevent overfitting, which is the tendency of the classifier to be sample-specific. The coefficients from Lasso logistic models are obtained as follows:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} J(\theta)$$
$$J(\theta) = \frac{1}{n} \sum_{i=1}^n \left( \frac{1}{2} x_i \theta - y_i \right)^2 + \frac{1}{C} \|\theta\|_1$$

where  $C$  is a hyperparameter that controls the amount of regularization applied by Lasso. It is proportional to the number of coefficients that Lasso sets to 0. To be clear, a positive coefficient is not informative about the absolute frequency with which male or female politicians use particular words (we examine absolute differences in the supporting information), but implies that a word is used relatively more frequently by one gender.

We perform 10-fold cross-validation on our training set to find the value of  $C$  that maximizes the predictive score. The gender  $y_i$  of tweet  $x_i$  can be predicted as:

$$y_i = \sigma(x_i \hat{\theta} + b) - 0.5$$

where  $\sigma$  is the logistic function, and  $b$  is the intercept. Our evaluation score is the balanced accuracy, which is the proportion of tweets correctly classified as male or female, weighted to make them equally significant in the measure. That is, if there are more male tweets, mistakes classifying female tweets weigh more.

To train the Lasso logistic classifier, we used the Python 3 library Scikit-learn 0.20.2. Training the model is not completely deterministic, and different runs with the same training data may yield slightly different coefficients. We overcome this instability by running 30 models and averaging the values of the coefficients.

## 4 Results for text written by politicians

We first analysed a series of tweets published by politicians. To examine whether male and female politicians communicate with the public in different ways, we regressed the gender of politicians on the preprocessed words and emojis included in their texts. We fixed the Lasso regularization parameter  $C$  at the value that yielded the best validation score. The best results were for  $C$  equal to 5, which yielded a balanced accuracy of 62.97% in the validation set (see details about the choice of regularization parameters and the overfitting curves in the Supporting Information). We report the words that were most linked to male or female senders.

Figure 1 plots the 50 linguistic features that are most male- and female-linked, that is, the words that are relatively more frequently used by each gender. The figure shows the coefficients and the English translations. A list with the 100 most linked words in Spanish and English, their coefficients, and the topic suggested by the word can be found in the Supporting Information.

Figure 1 about here

The results suggest that the differences in text written by male and female politicians are strongly consistent with gender stereotypes. The list of female-linked features includes many words relating to gender, including *transsexual*, *patriarchal*, *feminist*, *reproduction*, a form of gender-neutral writing<sup>6</sup>, *woman*, *prostitute*, *LGBTI*, *mistreatment* and so on. No gender-related

---

<sup>6</sup> In Spanish, the grammatical gender is present in nouns, adjectives, articles, and some pronouns (for instance, *hijo* means son and *hija* means daughter). The plural form is masculine when the group contains both males and females (for instance, a son and a daughter are *hijos*). To avoid this masculine form, some people use alternative letters or symbols to create a gender-neutral ending, like @ or x (*hij@s*, *hijxs*).

words appear in the list of male-linked words. The same is true of words about social issues: words like *human rights*, *dependant*, or *infancy* are among the 50 most female-linked words and *public health*, *rent*, *racism*, or *poverty* are among the 100 most female-linked words. Gender and social issues are clearly associated with women MPs.

The male-linked words are about classic political and ideological issues, such as *extremist*, *liberal*, *separatist*, *plurinational*, and *bipartisan*. Some sport-related words also appear on the list, like *basketball* or *athlete*, as do words about infrastructure and transportation, such as *vehicle*, *water transfer*, *airport* or the steam locomotive emoji.<sup>7</sup>

Emojis are present in both lists, but especially in the women's list. There are very clear differences in the tone and emotions expressed by these emojis. Men's are generally factual, such as the film projector, data chart, camera, and printed page. When they include faces, these tend to be emoticons, that is, visual representations of facial expressions using keyboard characters such as :), and ;). On the other hand, women's emojis show many faces reflecting a variety of both positive and negative emotions. Hearts in several colours also appear, as do roses.

## 5 Results for text addressed at politicians

In this section we examine tweets that are written by citizens in reply to politicians on Twitter, in order to examine if the public address male and female politicians differently on social media. We repeat the same procedures as in the previous section but using the second dataset. Figure 4 shows the list of 50 words that are most predictive of the gender of the receiver. Again, the best model is found at  $C = 5$ , with a balanced accuracy of 59.83%. Results for the rest of the models are available in the Supporting Information.

Two results stand out: first, there is some overlap between the words most associated to male or female politician authors and the words that the public addresses to them. The list of male-linked words estimated from the replies to politicians contains several words about politics, such as *Minister*, *President*, *monarchy*, words about sports, like *referee* and *soccer*, and about infrastructure, like *train*, *toll*, *highway*, or *bridge*. Many female-linked words are related to gender, such as *patriarchy*, *feminist*, *woman*, or *rapist*, and to social issues, like *xenophobic*.

---

<sup>7</sup> The list of the 100 most male-linked terms contains additional sports-related words like *football*, *coach*, *goal*, *player*, *league* (see supporting information) and words related to infrastructures like *train*, the name of the public rail company, RENFE, and *railway*.



Second, the results about the words used in replies reveal the insults or demeaning words that are most probably addressed at male and female politicians. Some offensive words are more associated with men, such as *clown*, *foolish*, or *stupid*. The list of the 100 most male-linked words includes other insults like *brownozer* and *puppet*. It also contains other words which can have critical or insulting connotations in Spanish depending on the context, such as *nice*, *champion*, or *streetwise*. And it contains a number of informal ways of addressing people such as *kid* or *guy*. Male-linked words also contain terms clearly associated with corruption such as *black* and *embezzlement*.

By contrast, two clearly offensive terms that appear in the list of female-linked words are the words *Hitlerian* and *totalitarian*, which can be an insult depending on who they are addressed to. In Spain, these two words are sometimes addressed to feminists (the word ‘feminazi’ is frequently used to discredit feminists), which would suggest that specifically female-linked insults are used to address women who defend gender issues. This interpretation is reinforced by the appearance of the words *fascist* and *intolerant* among the 100 most female-linked words). In addition, the female-linked words contain words that are clearly sexist, especially ones that are related to appearance and sexual objectification such as *pretty* (which appears in two forms, ‘guapa’ and the more informal ‘guapi’), *cute*, and *blonde*. The list of the 100 most female-linked words contains the quintessentially sexist insults *ugly*, *bitch*, and *stupid*. The contrast between the lack of emoticons in these sets of results about the text addressed to politicians and the abundance of gender-linked emoticons in the text written by politicians suggests that emoticons are predictive of the gender of the writer but not of the receiver.

Figure 2 about here

## 6 Conclusions

This research note employs a large sample of tweets to study differences in the language used by male and female politicians when communicating on Twitter, as well as by citizens when addressing them. In both sets of analyses, we document significant gender differences in both content and style.

The results suggest that politicians play an active role in reproducing gender stereotypes when they can directly communicate with the public online, and they are not treated by the public equally but in discriminatory ways. It would be beyond the scope of this research note to examine the causes driving gendered communications on Twitter. Male and female politicians may make

different use of language because they adhere to cultural norms, have different preferences, or because they fear that communicating anti-stereotypically would attract hostility or indifference from users. Gendered communication from citizens to politicians could be driven by different audiences self-selecting to follow male and female politicians, by the different content initiated by those, or by differential treatment due to gender. In any case, our descriptive results contribute to studies of gendered communication in politics in several respects.

First, the approach based on social media data used here can easily be extended to the study of other contexts. It is particularly useful when studying countries with proportional systems and closed party lists, where there are limited opportunities for establishing whether male and female politicians prioritize different issues in office due to the strong role of parties and party discipline in other observable behaviour, such as roll call votes, and there is limited other data about the speech of male and female politicians. Posting on social media is less constrained by political parties than other activities in which politicians engage.

Second, Lasso logistic regression are based on linguistic features and reveal more subtle gender differences in the content and the tone of political communication than other approaches appropriate for usage with big data. While the more frequently used topic models obtain a small number of topics from the text, our approach identifies differences in other respects as well, such as the gendered usage of insults or emoticons.

Third, we have broken new ground by focusing on the words that are addressed disproportionately to male and female politicians by other Twitter users. We find that they face different treatment online. A result that stands out is that both male and female politicians encounter hostile language that is more often addressed specifically at their gender (in addition to the hostile content which is addressed at both men and women and hence is not captured in our lists). We confirm the suspicion that the hostile talk addressed specifically at women is often related to criticisms of feminist positions and female MPs are more likely to receive words that are apparently positive but are in fact sexist as they relate to their physical appearance, along with condescending words that infantilize them.

The study has similar limitations to other studies using social media data. One common criticism is sample selection bias. In this case, a large majority of MPs do have a Twitter account, but they differ greatly in how active they are. We are not attempting to characterize all political communication between politicians and citizens, but only to describe differences on Twitter. Another limitation concerns the fact that it focuses on Spain, which may differ from other countries. In spite of these limitations, this study is an example of the possibilities of machine

learning and natural language processing techniques to study classic political communication and gender issues.

Substantively, our results support the view that descriptive representation is relevant because politicians with different individual characteristics are likely to discuss different issues and have different communication styles. The claim that descriptive and substantive representation are linked is one of the most important arguments in support of the political inclusion of women and other politically under-represented groups in proportion to their numbers in the population, and consequently for the design of measures such as gender quotas that foster progress towards this goal. Our finding that male and female politicians talk differently online supports the view that the gender composition of legislatures is relevant for the content of political communication as well.

Finally, the finding that male and female politicians have different experiences online has practical implications. An intriguing possibility is that the differential treatment from citizens to male and female politicians is an understudied mechanism that produces substantive representation. Words related to feminism are more likely to be addressed to female politicians, which may motivate female politicians to specialize on gender issues, creating a growing divergence over time in the topics discussed by male and female politicians. Second, some forms of differential treatment may constitute a barrier in the progress towards equal representation. Female politicians are more likely to receive comments related to their physical appearance, which may be uncomfortable in a culture that values beauty and may motivate some women to withdraw to avoid public exposure. The impact of such gendered online experiences on the decision to run as a politician is difficult to counter-act because of the impossibility to mandate equal treatment from citizens to politicians, or technical solutions, such as moderating words linked to the gender of the receiver, would collide with the principle of free speech.

## Acknowledgements

We gratefully acknowledge financial support from the Spanish Ministry for the Economy, Industry and Competitiveness (grant CSO2016-79569-P).

## References

- Aaldering, L., & Van Der Pas, D. J. (2018). Political leadership in the media: Gender bias in leader stereotypes during campaign and routine times. *British Journal of Political Science* 1-21.
- Bailey, M. B. & Nawara, S. P. (2018). Addressing Women and Minorities on Social Media by the 2016 U.S. Presidential Candidates. *Social Science Computer Review*. Available online at: <https://doi.org/10.1177/0894439318779342>
- Barberá, P. (2015). Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data. *Political Analysis* 23(1): 76–91.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A. & Bonneau, R. (2015). Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber?. *Psychological Science* 26(10): 1531–1542.
- Bauer, N. M. (2015). Emotional, sensitive, and unfit for office? Gender stereotype activation and support female candidates. *Political Psychology* 36(6): 691-708.
- Bauer, N. M. (2018). Running Local: Gender Stereotyping and Female Candidates in Local Elections. *Urban Affairs Review*. <https://doi.org/10.1177/1078087418770807>
- Beauchamp, N. (2017). Predicting and Interpolating State-Level Polls Using Twitter Textual Data. *American Journal of Political Science* 61(2): 490–503.
- Bohren, J. Aislinn and Imas, Alex and Rosenberg, Michael (2018). The Dynamics of Discrimination: Theory and Evidence. *Working Paper*. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3235376>
- Bystrom Theocharis, Y., Barberá, P., Fazekas, Z., Popa, S. A. & Parnet, O. (2016). A Bad Workman Blames His Tweets: The Consequences of Citizens Uncivil Twitter Use When Interacting With Party Candidates. *Journal of Communication* 66(6): 1007–1031.
- Cassese, E. C., & Holman, M. R. (2018). Party and gender stereotypes in campaign attacks. *Political Behavior* 40(3): 785-807.
- Ceron, A., Curini, L., & Iacus, S. (2017). Social media and elections: A meta-analysis of online-based electoral forecasts. In K. Arzheimer, J. Evans, & M. Lewis-Beck (eds), *Handbook of electoral behavior*. London: Sage.
- Chen, Z., Lu, X., Shen, S., Ai, W., Liu, X. & Mei, Q. (2018). Through a Gender Lens: An Empirical Study of Emoji Usage over Large-Scale Android Users. *Working Paper*. Available at *Computing Research Repository (CoRR)*: <https://arxiv.org/abs/1705.05546v2>
- Ditonto, T. M., Hamilton, A. J. & Redlawsk, D. P. (2014). Gender Stereotypes, Information Search, and Voting Behavior in Political Campaigns. *Political Behavior* 36(2): 335–358.
- Dolan, K., & Lynch, T. (2014). It takes a survey: Understanding gender stereotypes, abstract attitudes, and voting for women candidates. *American Politics Research* 42(4): 656-676.

- Evans, H. K., & Clark, J. H. (2016). "You tweet like a girl!" How female candidates campaign on Twitter. *American Politics Research* 44(2): 326-352.
- Grimmer, J. & Stewart, B. M. (2013). Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis* 21(03): 267–297.
- Herrnson, P. S., Lay, J. C. & Stokes, A. K. (2003). Women Running "as Women": Candidate Gender, Campaign Issues, and Voter-Targeting Strategies. *The Journal of Politics* 65(1): 244–255.
- Hwang, H. S. (2014). Gender Differences in Emoticon Use on Mobile Text Messaging: Evidence from a Korean Sample. *International Journal of Journalism & Mass Communication*.
- Jamieson, K. H. (1995). *Beyond the double bind: Women and leadership*. New York: Oxford University Press.
- Jost, J. T., Barberá, P., Bonneau, R., Langer, M., Metzger, M., Nagler, J., Sterling, J. & Tucker, J. A. (2018). How Social Media Facilitates Political Protest: Information, Motivation, and Social Networks. *Political Psychology* 39(S1): 85–118.
- Just, M. R. & Crigler, A. N. (2014). Gender and Self-Presentation in Social Media: An Analysis of the 10 Most Competitive 2012 U.S. Senate Races. *Working Paper*. Available at SSRN: <https://papers.ssrn.com/abstract=2453741>
- Krook, M. L. (2017). Violence against Women in Politics. *Journal of Democracy* 28(1): 74–88.
- Lawless, J. L. (2004). Women, War, and Winning Elections: Gender Stereotyping in the Post-September 11th Era. *Political Research Quarterly* 57(3): 479–490.
- Lawless, J. L. (2015). Female Candidates and Legislators. *Annual Review of Political Science* 18(1): 349–366.
- Lucas, C., Nielsen, R. A., Roberts, M. E., Stewart, B. M., Storer, A. & Tingley, D. (2015). Computer-Assisted Text Analysis for Comparative Politics. *Political Analysis* 23(2): 254–277.
- Munger, K. (2017). Dont @ Me: Experimentally Reducing Partisan Incivility on Twitter. *Working Paper*. Available online at: <http://cess.nyu.edu/wp-content/uploads/2017/02/Dont-@-Me.pdf>
- Newman, M. L., Groom, C. J., Handelman, L. D. & Pennebaker, J. W. (2008). Gender Differences in Language Use: An Analysis of 14,000 Text Samples. *Discourse Processes* 45(3): 211–236.
- Palomares, N. A. (2008). Explaining Gender-Based Language Use: Effects of Gender Identity Salience on References to Emotion and Tentative Language in Intra- and Intergroup Contexts. *Human Communication Research* 34(2): 263–286.
- Park, G., Yaden, D. B., Schwartz, H. A., Kern, M. L., Eichstaedt, J. C., Kosinski, M., Stillwell, D., Ungar, L. H. & Seligman, M. E. P. (2016). Women are Warmer but No Less Assertive than Men: Gender and Language on Facebook. *PLoS ONE* 11(5): e0155885.

- Rheault, L., Rayment, E. & Musulan, A. (2019). Politicians in the line of fire: Incivility and the treatment of women on social media. *Research & Politics* 6(1): 1-7.
- Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B. & Rand, D. G. (2014). Structural Topic Models for Open-Ended Survey Responses. *American Journal of Political Science* 58(4): 1064–1082.
- Teele, D. L., Kalla, J., & Rosenbluth, F. (2018). The Ties That Double Bind: Social Roles and Women's Underrepresentation in Politics. *American Political Science Review* 112(3): 525-541.
- Wagner, K. M., Gainous, J., & Holman, M. R. (2017). I am woman, hear me tweet! Gender differences in Twitter use among congressional candidates. *Journal of Women, Politics & Policy* 38(4): 430-455.
- Wilkerson, J. & Casas, A. (2017). Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges. *Annual Review of Political Science* 20(1): 529–544.
- Wolf, A. (2000). Emotional Expression Online: Gender Differences in Emoticon Use. *CyberPsychology & Behavior* 3(5): 827–833.
- Wu, A. H. (2018). Gendered Language on the Economics Job Market Rumors Forum. *AEA Papers and Proceedings* 108: 175–179.