Policy Research Working Paper 10769

Gender, Social Support, and Political Speech

Evidence from Twitter

Rachel Heath
Roy van der Weide

WORLD BANK GROUP
Development Economics
Development Research Group
May 2024

Policy Research Working Paper 10769

Abstract

Despite evidence that women's political preferences differ from those of men, women are less likely to participate in political and social discussions on Twitter and other social media. Following recent evidence that in-person social support matters for women's political participation, women are hypothesized to form similarly supportive communities online. This paper tests this hypothesis using data from Twitter. The collected data comprises 451 hashtags on a broad range of (non-mutually exclusive) topics: social, gender, racial, LGBTQ, religion, youth, education, economic, health, COVID, climate, political, security, entertainment and lifestyle, and the Middle East and Northern Africa. The empirical results indicate that women are more likely to participate when the debate(s) feature female influential voices. This finding supports the potential role of mutual support in bolstering women's participation in important debates.

This paper is a product of the Development Research Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The authors may be contacted at rvanderweide@worldbank.org.

Gender, Social Support, and Political Speech: Evidence from Twitter

Rachel Heath[*] Roy Van Der Weide[†]

[*]Department of Economics, University of Washington; rmheath@uw.edu

[†]World Bank; rvanderweide@worldbank.org

[‡]The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. The authors gratefully acknowledge the financial support from the Office of the Chief Economist for the Middle East and North Africa (MNACE) under the regional Labor and Gender Research Programs. The authors would also like to thank Rwan Adil Osman Mahmoud for excellent research assistance and Jed Friedman and Hans Hoogeveen for their generous comments.

1 Introduction

Despite evidence that women's political preferences differ from those of men, women are less likely to participate in political and social discussions on Twitter and other social media. Given evidence that social media can affect offline behavior (Alatas et al., 2019, Bianchi et al., 2023, Müller and Schwarz, 2020), women's lack of participation on Twitter may have important ramifications for policy, especially around gender. Following recent evidence that in-person social support (in the form of self-help groups) matters for women's political participation (Prillaman, 2023), we hypothesize that women can form similarly supportive communities online.

We test this hypothesis using data from Twitter. We collected data on 451 hashtags, which we categorize as relating to various (non mutually exclusive) topics: social, gender, racial, LGBTQ, religion, youth, education, economic, health, COVID, climate, political, security, entertainment and lifestyle, and the Middle East and Northern Africa (MENA). We use data on interactions (retweets, quote tweets, and likes) between users to construct two key network statistics that proxy the support women tweeters might benefit from: clustering (the fraction of closed triangles, i.e. how often are two of someone's friends themselves friends?) and degree (the number of connections).

We find that women are more likely to participate and more likely to be influential when their networks are larger and more clustered. This finding supports the potential role of mutual support in bolstering women's participation in important debates.

2 Related Literature

2.1 Gender, Social Support, and Political Speech

Research has pointed out that women have different political preferences from men. There are two types of evidence that contribute to this consensus. First, when there is plausibly exogenous variation in the gender of a leader, different policies are enacted. Political reservations for women in India lead to differential provision of public goods valued by women, as measured by their complaints to village councils (Chattopadhyay and Duflo, 2004). When a woman just barely wins a close election, regression discontinuity results point to higher spending on health in US legislatures (Rehavi, 2007), more spending on childcare in Germany (Baskaran and Hessami, 2019), higher spending on education and a higher likelihood of girls in urban India completing primary education (Clots-Figueras, 2011, 2012), and less hiring of temporary public sector workers during electoral years (Brollo and Troiano, 2016).[1]

A second body of evidence finds causal impacts in history of women gaining the right to vote. In the United States, women's suffrage led to more liberal policies and higher government spending (Lott and Kenny, 1999), and in particular, increased spending on public health and lower child mortality (Miller, 2008).

Despite the fact that women frequently have different political preferences than do men, there is also evidence that women are less likely to participate politically. Worldwide, only 22% of legislators are women, representing 25% of legislators in Europe, 19% in Asia, 22% in Sub-Saharan Africa, and 27% in the Americas (Kalandadze, 2015). This disparity is not limited to elected offices. Women in every continent are less likely than men to be members of political parties (World Bank, 2011) and in data pooled across

[1]This finding is not universal; Ferreira and Gyourko (2014) find no evidence of effects of women mayors in the US on size of local government, the composition of municipal spending and employment, or crime rates in the short or long run.

African countries, women are 12% less likely to vote than men (Isaksson, Kotsadam and Nerman, 2014). Women are also less likely to contact government officials; women in the US are 21 percentage points less likely to contact a government official (Burns, Schlozman and Verba, 2021) and in rural India, women were 47 percentage points less likely to say they had attended a village assembly meeting and 36 percentage points less likely to say they had contacted the local leader (Prillaman, 2023).

Prillaman (2023) points to a strategy for decreasing this gender participation gap. In particular, she finds that women who were assigned in a quasi-random way to be in an all-women self-help group were more than twice as likely than other women to attend village assembly meetings and make claims on local leaders. She argues that the channel is collective action; women acted jointly to "demand representation" and "combat social sanctioning."

This important role of networks in political speech contrasts with other settings in which women appear to be disadvantaged by the role of networks. A classical case is the labor market; women in the US are less likely than men to find jobs through social networks (Bradshaw, 1973, Ports, 1993, Smith, 2000, Ioannides and Loury, 2004). Women also earn less when they do get a job through a social network (Loury, 2006). A possible mechanism for this is that women's contacts are other women who are less likely than men to be influential in hiring for good jobs. Accordingly, when women are in networks with men, they benefit just as much in the labor market as do men (McDonald, 2011).

2.2 Twitter and Other Social Media

A growing social science literature finds that social media such as Twitter reflects offline behaviors and beliefs, including reporting of sex crimes (Levy and Mattsson, 2023), anti-Muslim hate crimes in the UK (Ala' Alrababa'H et al., 2021), perceptions of economic uncertainty (Altig et al., 2020, Baker et al., 2021), utility from weather (Baylis, 2020), political polling data (Beauchamp, 2017), Arab spring protests (Acemoglu, Hassan and Tahoun, 2018), and which politicians are influential (Mankad and Michailidis, 2015).[2] Accordingly, social scientists frequently use social media to assemble data more quickly and at a higher frequency than is typically available in standard datasets.

Moreover, Twitter does not only reflect offline attitudes, there is also evidence that exposure to certain material on Twitter can affect offline behavior. Alatas et al. (2019) worked with celebrities in Indonesia willing to endorse vaccines; they find that a Twitter endorsement affects vaccine takeup and knowledge of study participants' social networks. Donald Trump's tweets about the federal reserve affected the interest rates and other macroeconomic variables (Bianchi et al., 2023). This influence is not limited to material directly attributable to politicians and celebrities; quasi-random exposure to Twitter (driven by the home counties of South by Southwest festival attendees in 2007) led to anti-Muslim hate crimes during Donald Trump's campaign (Müller and Schwarz, 2020).

This potential to influence offline behavior underscores the importance of understanding the determinants of participation in political and social discussions on Twitter. Indeed, there is evidence that women use social media differently than men. Men are followed and retweeted by both men and women in greater numbers in medicine (Zhu et al., 2019) and journalism (Usher, Holcomb and Littman, 2018). Bamman, Eisenstein and Schnoebelen (2014) find that men and women use different linguistic styles on Twitter; men were more likely to use numbers, quantifiers and technology words while women were more likely to use pronouns, emotional terms and family terms. Gender

[2]Though this finding is not universal; Zeitzoff, Kelly and Lotan (2015) study discussion on the Iran-Israel conflict on Twitter. They find that English and Farsi discussion correlates with offline behavior, but Hebrew discussion does not.

differences persist when focusing on one particular topic; Holmberg and Hellsten (2015) find that men are more likely to express skepticism than women when tweeting about climate change. Cunha et al. (2014) and Shapp (2014) find that when women use hashtags, they are more likely to show a personal involvement (e.g. their own vote), whereas men use hashtags to show a persuasive strategy, for instance by expressing a command (e.g. vote for ...). Moreover, even when women participate in Twitter, their participation sometimes is used instrumentally in patriarchal organizations; Nielsen (2020) studies the case of female preachers in the Islamist Salafi movement.

3 Data and Methods

3.1 Hashtag Selection and Attributes

To identify a cross-section of hashtags that covers public discourses with broad participation across an inclusive set of topics, we appeal to a variety of resources including: (a) trending hashtag lists compiled by Twitter and third parties, and (b) news articles, op-eds and studies that document discourses on Twitter that have shaped public narratives both online and offline (and in some cases public policy) and/or raised public awareness of the issues at hand, think of *#metoo* and *#BlackLivesMatter*. Typical headlines of such articles, which are published on a regular basis, include: "These 10 Twitter hashtags changed the way we talk about social issues", "8 Massive Moments When Hashtag Activism Really Worked", "10 Twitter Hashtags That Shaped History", "10 years of hashtags that changed Twitter", "How the humble hashtag changed world politics", "6 Influential Hashtag Movements Across the World", and "Popular hashtags for social-justice on Twitter".

The selection of hashtags consists of: (i) socio-economic hashtags corresponding to specific movements, e.g., *#metoo*, *#BlackLivesMatter*, *#IceBucketChallenge*,

#LoveWins, *#Flygskam*, *#Ferguson*, *#ArabSpring*, *#StopFundingHate*, *#Brexit*, *#Cancelstudentdebt*, *#BringBackOurGirls*, *#COVID*19, (ii) more generic socio-economic hashtags, e.g., *#ClimateChange*, *#GunViolence*, *#abortion*, *#corruption*, *#democracy*, *#inequality*, *#education*, *#immigration*, *#security*, (iii) generic hashtags on non-socio-economic issues, e.g., *#art*, *#family*, *#friends*, *#fashion*, *#love*, *#travel*, *#vegan*, *#championsleague*, and (iv) more specific hashtags on non-socio-economic issues, e.g. *#FollowFriday*, *#ShareYourEars*, *#CinnamonChallenge*, *#TheDress*, and *#CupforBen*.

The hashtags cover a wide variety of topics, ranging from social issues (further divided into e.g., gender, racial, and LGTBQ), economic, political, health, security, entertainment etc. This allows us to inspect how female participation varies across these different issues, and in particular, establish whether female participation is higher in public debates on issues that concern them. To that end, we created a set of binary variables that record for each hashtag the issue(s) it is associated with. Additionally, we created a binary variable that indicates whether the discourse is focused on Northern Africa and the Middle East (MENA). Female participation is conjectured to be lower in areas where women are more at risk of harassment, either online and offline, (and where female economic/political participation are lower). While the MENA region is not the only area where these conditions apply, we will use it as a proxy. Table 1 shows the full list of hashtag attributes along with the number of hashtags in our database that fall into each of these categories. Note that a given hashtag could be assigned to multiple categories. For example, the hashtag *#MahsaAmini* concerning gender based violence in the Islamic Republic of Iran is associated both with "gender" and "security".

The majority of hashtags are in English, which is the common choice of language for hashtags. Often, tweets in languages other than English would still feature an

English hashtag. Hashtags in other languages are also accounted for, albeit in smaller numbers. One prominent example is the hashtag *#Flygskam*, which was at the center of a large-scale movement to discourage people from flying in order to lower carbon emissions. (Accordingly, this hashtag is assigned to "climate".) An additional effort was made to screen for hashtags in Arabic, for the same reason we created the MENA variable. In total, our database counts 320 hashtags in English, 111 in Arabic, and 20 in other languages (and 451 hashtags in total).

3.2 Compilation of Twitter Data

The Twitter Enterprise Full-Archive search API is used to ingest the tweets that feature the selected hashtags. This API granted us access to the full repository of tweets dating back to the inception of the platform. There is however a bandwidth constraint that prevented us from downloading the entire repository. Accordingly, we drew a stratified random sample of tweets.

We stratified by hashtag-year-month for a period of 10 years ranging from January 2014 to June 2023. The ingestion of tweets is done iteratively in batches of up to 5000 tweets each. We start in June 2023 and then work our way back in time towards January 2014, one batch per hashtag-year-month at a time. Upon reaching January 2014 for any given hashtag, we return to June 2023 or the most recent year-month for which our sample of tweets does not yet cover the full set of available tweets, and again work our way back to 2014; one batch for each hashtag-year-month at a time until our bandwidth constraint is exhausted.

The resulting database is an unbalanced panel as different hashtags come in and out of existence at different points in time. Starting in 2014 gives a sufficient amount of data prior to the birth of the *#metoo* hashtag in 2017, which may mark a structural break in female participation on Twitter, particularly in public discourses that concern them. Select summary statistics derived from the database are shown in Table 2. (Note that for the year 2023, the database covers the period January through June.) In total, the database counts about 90 million tweets.

3.3 Identifying Gender of Users

Twitter does not indicate whether a user account belongs to an individual or an organization, and in the case of individuals, whether the user is female or male.[3] Accordingly. we have to determine these user attributes on our end. There are a variety of existing methods that have been developed for this task. We proceed in two steps. First we classify users as either individuals or organizations. Second, for users identified as individuals, we impute their gender. These two steps involve separate imputation methods that have been optimized for these individual tasks.

For the first step, we adopt the approach put forward by (Cetinkaya et al., 2023), which classifies accounts as either organizations or individuals. The predictors include the user account's profile metadata, such as name, bio description, location, whether the bio includes an URL, time on Twitter, as well as Twitter account statistics, such as the number of tweets, retweets, following count, followers count, media count, and the number of likes.

The publicly available "Demographer" dataset is used by the authors to train the model and evaluate out of sample performance. These data contain both Twitter user identifiers and a variable indicating whether the account belongs to an individual or an organization for a total of 214,236 user accounts. The study conducts a 7-fold cross validation test, where $6/7th$ of the data is used for training and the remaining $1/7th$ of observations are used for out of sample prediction. The reported accuracy is about 97% (i.e., only 3 percent of user accounts are incorrectly classified). We also applied

[3]For simplicity, we do not consider non-binary genders.

their model to a sample of 300 user accounts from our own database. The out of sample accuracy in that case is 91%.

For the second step, which is concerned with gender classification, we use the Twitter Demographer package, which offers a user-friendly API. The underlying prediction method is built on the M3 classifier introduced by (Wang et al., 2019). The predictors in this case include the user's name, screen name, bio description as well as the user's profile image (which involves image recognition analysis). Three distinct datasets are used, each serving a specific purpose, for training and out of sample performance testing of the M3 classifier:

1. A Twitter dataset for the period 2014 to 2017 for which gender and age of users have been identified.

2. An images dataset, comprising a total of 523,051 images. This dataset is derived from head-shots of actors sourced from IMDB and profile pictures extracted from Wikipedia. These images bolster the classifier's ability to incorporate visual cues in its predictions.

3. A dataset derived from crowd-sourcing efforts where user gender and age are collected across 32 different languages.

The reported out of sample accuracy of the gender classifier is 91.8% (Wang et al., 2019). When we apply the model to a sample of 200 individual users from our database, we obtain an out of sample accuracy of 88%.

The gender classification is computationally expensive, in large part stemming from the image recognition component that is utilized to infer the gender of the user. For this reason, we first draw a simple random sample of 1000 user accounts for each hashtag-year-month and group these into individuals and organizations (the first step). Among the users that are identified as individuals, the gender classifier is applied to a simple random sample of 50 individuals (for each hashtag-year-month). Accordingly, the share of female users (the dependent variable in our study) is derived from a random sample of up to 600 individuals for each hashtag-year (the unit of observation in our study).

Figure 1 shows the distribution of the share of female participation across hashtags for the period 2021 to 2023. Average female participation observed in our dataset for this period is 33 percent. Put differently, female participation is on average half of male participation. We could not find many sources to compare this estimate against. Two sources we could find report estimates of 35 and 43 percent ((DataReportal, 2023); (Statista, 2022)). Female participation is also seen to vary considerably across different issues (Figure 1). Yet, we rarely observe an even participation between men and women. For 90 percent of observations, female participation ranges from 11 to 58 percent.

3.4 Network Measures

The clustering coefficient (Watts and Strogatz, 1998) measures the fraction of closed triangles in a network. That is, if someone has two network connections, what is the probability that those two connections are themselves connected. We take it as a measure of social support, because clustered networks mean that mutual acquaintances are likely to know each other and thus be able to work together to support a fellow friend.

The local clustering coefficient (i.e. the clustering coefficient for a specific network node) for a given graph $G = (V, E)$ consisting of edges E and vertices V is calculated by summing the number of actual connections between nodes j and k, scaled by the number of potential connections $k_i(k_i - 1)$, where k_i is the number of nodes and N_i in the network of an individual:

$$C_i = \frac{|\{e_{jk} : v_j, v_k \in N_i, e_{jk} \in E\}|}{k_i(k_i - 1)}$$

We then sum this across all the individuals in a network to construct a network-wide measure.

The degree of an individual in a network is the fraction of the total nodes in the network that it is connected to.

4 Results

4.1 Women's Participation and Influence on Twitter

We begin by exploring the rate of participation of women in the selected hashtags over our study period. Over the course of the study period, an average of 36% of individual users in the selected hashtags are women. Figure 2 shows that this rate is decreasing over time: 38% of individuals in 2014 are by women, whereas by 2023, only 33% are. This is striking in light of a the MeToo hashtag that began in 2017 and highlighted the role of women's participation in social media debates about gender and other social issues. Instead, 2017 is actually approximately the time that women's participation began decreasing, which could reflect a backlash against women who speak up on social media and reinforce our hypothesis about the importance of social support. Moreover, panel B examines tweets about gender separately from tweets about other topics. The decline in women's participation is actually stronger in tweets about gender: the fraction of tweets about gender by women peaked at 50% in 2017, but had fallen to 42% in 2022 and 2023.

Given the overall lower rate of participation of women, we explore the determinants of women's participation on Twitter. We hypothesize that, in addition to the gender differences in the nature of Twitter participation discussed in section 2.2, women are more likely to participate in discussions about gender. Column 1 of table 3 confirms that, across the study years, the share of women in tweets about gender-related hashtags is 13.6 percentage points higher when compared to non-gender-related hashtags. Notably, most tweets even about gender-related hashtags are by men: the constant term of 0.34 (the share of tweets from women about non-gender-related hashtags) implies that a mean share female of 48% of tweets about gender-related hashtags come from women. Column 2 indicates that the relationship between gender and the share of tweets from women is constant across the survey period. Column 3 tests whether women are more likely to tweet about popular hashtags. This effect is present, though relatively small in magnitude: a 10,000 increase in the number of tweets worldwide (the median in our data is 5393 tweets) represents a 0.0167 increase in the share female of tweets with that hashtag.

Figure 3 explores the relationship between the share of female users and the subject of discourse across an inclusive set of topics. Women are approximately equally likely to engage on racial issues (40% of users) and LGBTQ issues (50% of users) as they are on gender issues. Women are much less represented in tweets focusing on MENA (22% of users) and religion (30% of users), and political issues (29% of users).

In table 4, we turn to the question of influence of women versus men, as measured by retweets. In particular, we examine the outcome of the logged mean retweets by female tweeters compared to the logged mean retweets by male tweeters. Note that this outcome is negative on average, suggesting that men are more influential, though this effect is not huge (1.28 retweets for the average tweet by a woman, compared to 1.11 retweets for the average tweet by a man). We find that women are more influential in conversations where more women participate; a 10% increase in the share of women participating in a conversation decreases the gender gap in influence by 17% percent. Women are also more influential in conversations about gender; there's a 23% decrease in the gender influence gap on tweets about gender. By contrast, column 2 indicates that women are less influential (conditional on the number of women participating) in conversations about race, youth, and health, but more influential in conversations about MENA. The fact that few women participate in hashtags about MENA – but those that do are quite influential – highlights the importance of studying the conditions under which women do participate and the role for social support, which we turn to next.

4.2 Gender and Social Support

We now turn to assessing the role of our measures of social support. Table 5 considers as outcome variables the logged average of retweets by women and men. Column 1 shows that women are retweeted more when networks are more tightly clustered: a one-standard deviation increase in the clustering coefficient (measured across all hashtags in the study period) increases the number of retweets of women by 2.1%. Column 2 indicates that this is entirely due to the clustered nature of women's networks, which increase women's influence by even more: a one-standard deviation increase in the mean clustering coefficient of women increases the number of retweets of women by 4.0%. By contrast, clustered men's networks decreases women's influence: a one-standard deviation increase in the mean clustering coefficient of men decreases the number of retweets of women by 1.5%.

In Column 3, we turn to the number of followers as another measure of support for women who tweet about a topic. Women have more retweets when women have larger networks, but actually have fewer retweets when men have larger networks: a one standard deviation in women's network size increases the number of retweets of women by 13%, whereas a one standard deviation in men's network size decreases the number of retweets of women by 6.0%. So, it is not a simple story where more people tweeting about something means more retweets, because larger male networks decrease women's influence. Rather, it appears that large male networks decrease women's participation, plausibly through reduced social support.

Columns 4 through 6 examine the retweets by men. Men also retweet more in tightly clustered networks, though the effect is smaller in magnitude than the effect on women. Column 5 shows a striking asymmetry with the results for women in column 2: tightly clustered networks of women do not decrease the number of retweets that men garner. Tightly clustered men's networks do increase the number of retweets garnered by men, but the effect is smaller in magnitude: a one-standard deviation increase in the mean clustering coefficient of men increases the number of retweets of men by 1.4%. Column 6 does show a similar pattern for network size as women: men get more retweets when men have bigger networks, but fewer when women have bigger networks.

Table 6 assesses the role of these measures of social support on women's participation across different hashtags. In addition to women being more influential in hashtags in which conversations are more tightly networked, women are also more likely to participate: a one standard deviation increase in the overall clustering coefficient increases the share of women participating by 2.2 percentage points. This effect is conditional on both gender, which increases participation, and MENA, which decreases participation. Column 2 shows that this effect is again driven by the clustering coefficient of women: a one standard deviation increase in women's clustering coefficient increases the share of women participating by 2.6 percentage points, while there is no average effect of men's clustering coefficient. Finally, women are also more likely to participate in hashtags where fellow women who participate have larger networks but less likely when men have larger networks – the share female is 3.5 percentage points higher when there is a one standard deviation increase in women's network size, while it is 3.3 percentage points lower when there is a one standard deviation increase in men's network size.

5 Discussion

We examine the relationship between two measures of social support – the clustering and size of networks – on women's participation and influence on Twitter. We find that women are more influential, and participate more, when they have more social support, as measured by tightness of networks (particularly women's). We interpret this finding as indicating that women participate in political debates when fellow women support them. Note that our findings on participation go against a self-selection story in which only very determined and influential women participate given the obstacles they face on social media; this story would predict that our network measures are associated with lower (not higher) female participation.

Given evidence that Twitter and other social media can influence offline events, these patterns affect the extent to which women participate in important debates in society. An important policy implication is that fora for public opinion should be designed in ways to both allow women to access other women for support and to minimize the role for unwelcome feedback from male participants.

References

Acemoglu, Daron, Tarek A Hassan, and Ahmed Tahoun. 2018. "The power of the street: Evidence from Egypt's Arab Spring." *The Review of Financial Studies*, 31(1): 1–42.

Alatas, Vivi, Arun G Chandrasekhar, Markus Mobius, Benjamin A Olken, and Cindy Paladines. 2019. "When celebrities speak: A nationwide twitter experiment promoting vaccination in Indonesia." National Bureau of Economic Research.

Ala' Alrababa'H, Marble, William, Salma Mousa, Alexandra A Siegel, et al. 2021. "Can exposure to celebrities reduce prejudice? The effect of Mohamed Salah on islamophobic behaviors and attitudes." *American Political Science Review*, 115(4): 1111–1128.

Altig, Dave, Scott Baker, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Scarlet Chen, Steven J Davis, Julia Leather, Brent Meyer, Emil Mihaylov, et al. 2020. "Economic uncertainty before and during the COVID-19 pandemic." *Journal of Public Economics*, 191: 104274.

Baker, Scott R, Nicholas Bloom, S Davis, and Thomas Renault. 2021. "Twitter-derived measures of economic uncertainty."

Bamman, David, Jacob Eisenstein, and Tyler Schnoebelen. 2014. "Gender identity and lexical variation in social media." *Journal of Sociolinguistics*, 18(2): 135–160.

Baskaran, Thushyanthan, and Zohal Hessami. 2019. "Competitively elected women as policy makers." CESifo Working Paper.

Baylis, Patrick. 2020. "Temperature and temperament: Evidence from Twitter." *Journal of Public Economics*, 184: 104161.

Beauchamp, Nicholas. 2017. "Predicting and interpolating state-level polls using Twitter textual data." *American Journal of Political Science*, 61(2): 490–503.

Bianchi, Francesco, Roberto Gómez-Cram, Thilo Kind, and Howard Kung. 2023. "Threats to central bank independence: High-frequency identification with Twitter." *Journal of Monetary Economics*, 135: 37–54.

Bradshaw, Thomas F. 1973. "Jobseeking methods used by unemployed workers." *Monthly Lab. Rev.*, 96: 35.

Brollo, Fernanda, and Ugo Troiano. 2016. "What happens when a woman wins an election? Evidence from close races in Brazil." *Journal of Development Economics*, 122: 28–45.

Burns, Nancy, Kay Lehman Schlozman, and Sidney Verba. 2021. "The private roots of public action." In *The Private Roots of Public Action.* Harvard University Press.

Cetinkaya, Yusuf, Mesut Gurlek, Ismail Toroslu, and Pinar Karagoz. 2023. "Twitter account classification using account metadata: organization vs. individual." *Turkish Journal of Electrical Engineering and Computer Sciences*, 30(4): 1404–1418.

Chattopadhyay, Raghabendra, and Esther Duflo. 2004. "Women as policy makers: Evidence from a randomized policy experiment in India." *Econometrica*, 72(5): 1409–1443.

Clots-Figueras, Irma. 2011. "Women in politics: Evidence from the Indian States." *Journal of public Economics*, 95(7-8): 664–690.

Clots-Figueras, Irma. 2012. "Are female leaders good for education? Evidence from India." *American economic journal: applied economics*, 4(1): 212–244.

Cunha, Evandro, Gabriel Magno, Marcos André Gonçalves, César Cambraia, and Virgilio Almeida. 2014. "He votes or she votes? Female and male discursive strategies in Twitter political hashtags." *PloS one*, 9(1): e87041.

DataReportal. 2023. "Twitter users, stats, data and trends." May 11.

Ferreira, Fernando, and Joseph Gyourko. 2014. "Does gender matter for political leadership? The case of US mayors." *Journal of Public Economics*, 112: 24–39.

Holmberg, Kim, and Iina Hellsten. 2015. "Gender differences in the climate change communication on Twitter." *Internet research*, 25(5): 811–828.

Ioannides, Yannis M, and Linda Datcher Loury. 2004. "Job information networks, neighborhood effects, and inequality." *Journal of economic literature*, 42(4): 1056–1093.

Isaksson, Ann-Sofie, Andreas Kotsadam, and Måns Nerman. 2014. "The gender gap in African political participation: Testing theories of individual and contextual determinants." *Journal of Development Studies*, 50(2): 302–318.

Kalandadze, Nana. 2015. "Women in politics – achieving key milestones." International Institute for Democracy and Electoral Assistance (International IDEA).

Levy, Ro'ee, and Martin Mattsson. 2023. "The effects of social movements: Evidence from# MeToo." *Available at SSRN 3496903*.

Lott, Jr, John R, and Lawrence W Kenny. 1999. "Did women's suffrage change the size and scope of government?" *Journal of political Economy*, 107(6): 1163–1198.

Loury, Linda Datcher. 2006. "Some contacts are more equal than others: Informal networks, job tenure, and wages." *Journal of Labor Economics*, 24(2): 299–318.

Mankad, Shawn, and George Michailidis. 2015. "Analysis of multiview legislative networks with structured matrix factorization: Does Twitter influence translate to the real world?" *The Annals of Applied Statistics*, 1950–1972.

McDonald, Steve. 2011. "What's in the "old boys" network? Accessing social capital in gendered and racialized networks." *Social networks*, 33(4): 317–330.

Miller, Grant. 2008. "Women's suffrage, political responsiveness, and child survival in American history." *The Quarterly Journal of Economics*, 123(3): 1287–1327.

Müller, Karsten, and Carlo Schwarz. 2020. "From hashtag to hate crime: Twitter and anti-minority sentiment." *Available at SSRN 3149103*.

Nielsen, Richard A. 2020. "Women's authority in patriarchal social movements: the case of female Salafi preachers." *American Journal of Political Science*, 64(1): 52–66.

Ports, Michelle Harrison. 1993. "Trends in job search methods, 1970-92." *Monthly Labor Review*, 116(10): 63–68.

Prillaman, Soledad Artiz. 2023. "Strength in numbers: how women's groups close India's political gender gap." *American Journal of Political Science*, 67(2): 390–410.

Rehavi, M Marit. 2007. "Sex and politics: Do female legislators affect state spending?" mimeo.

Shapp, Allison. 2014. "Variation in the use of Twitter hashtags." *New York University*, 1–44.

Smith, Sandra S. 2000. "Mobilizing social resources: Race, ethnic, and gender differences in social capital and persisting wage inequalities." *The Sociological Quarterly*, 41(4): 509–537.

Statista. 2022. "Distribution of Twitter users worldwide as of January 2022, by gender." December 5.

Usher, Nikki, Jesse Holcomb, and Justin Littman. 2018. "Twitter makes it worse: Political journalists, gendered echo chambers, and the amplification of gender bias." *The international journal of press/politics*, 23(3): 324–344.

Wang, Zijian, Scott Hale, David Adelani, Prze-myslaw Grabowicz, Timo Hartman, Fabian Flock, and David Jurgens. 2019. "Demographic inference and representative population estimates from multi-lingual social media data." *World Wide Web Conference*, 2056—-2067.

Watts, Duncan J, and Steven H Strogatz. 1998. "Collective dynamics of 'small-world'networks." *nature*, 393(6684): 440–442.

World Bank. 2011. *World development report 2012: Gender equality and development.* The World Bank.

Zeitzoff, Thomas, John Kelly, and Gilad Lotan. 2015. "Using social media to measure foreign policy dynamics: An empirical analysis of the Iranian–Israeli confrontation (2012–13)." *Journal of Peace Research*, 52(3): 368–383.

Zhu, Jane M, Arthur P Pelullo, Sayed Hassan, Lillian Siderowf, Raina M Merchant, and Rachel M Werner. 2019. "Gender differences in Twitter use and influence among health policy and health services researchers." *JAMA internal medicine*, 179(12): 1726–1729.

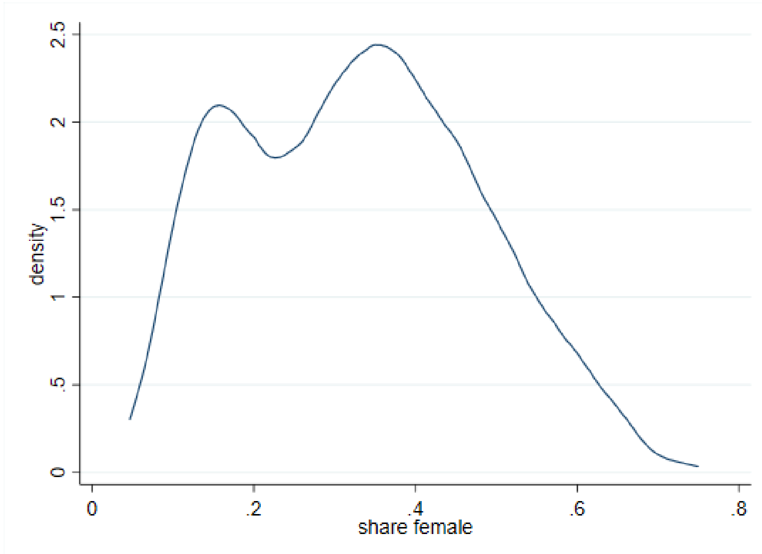

Figure 1: Kernel density of share of female participation for the years 2021 to 2023

Figure 2: Fraction of female users over time

Panel A: Overall

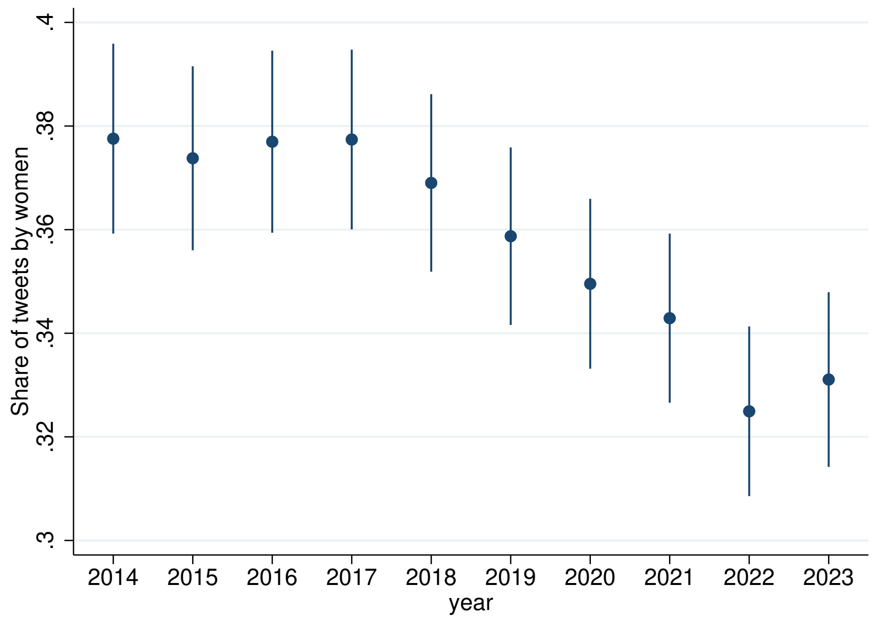

Panel B: Tweets on Gender versus other Tweets

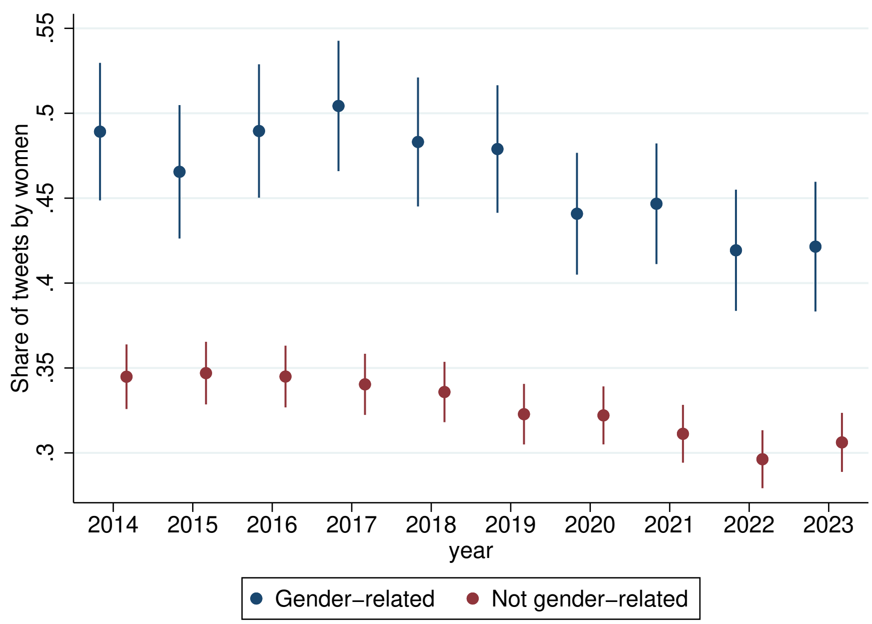

Figure 3: Fraction of tweets by women, by topic

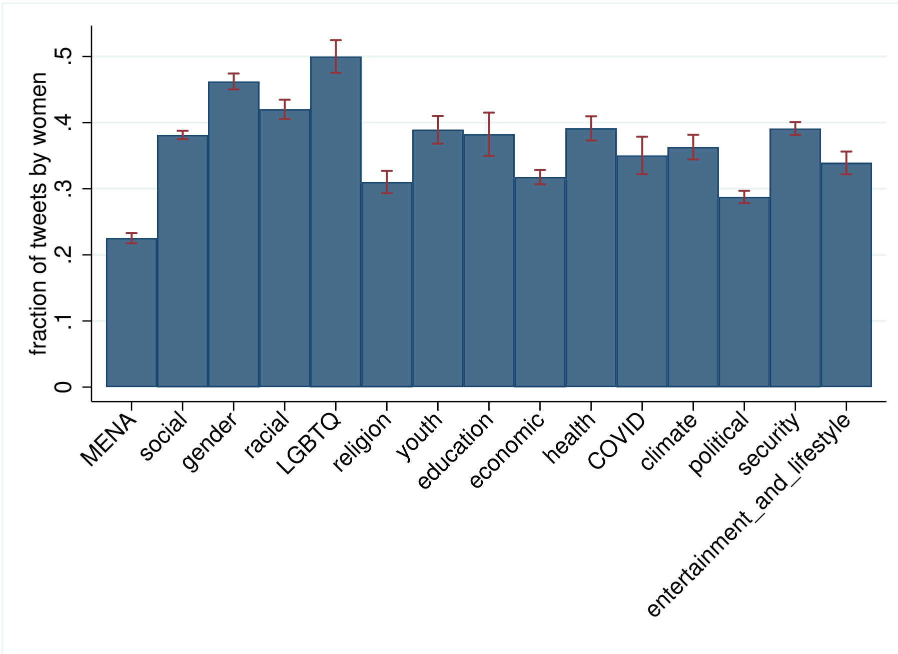

	nr. of hashtags
MENA	171
social	359
gender	110
racial	50
LGTBQ	17
religion	47
youth	27
education	14
economic	98
health	36
climate	14
political	110
security	155
entertainment	41
total	451

Table 1: Number of hashtags across topics

	nr. tweets	nr. orig tweets	share orig tweets
2014	4685619	2536352	0.54
2015	7036725	3243939	0.46
2016	5212028	2451700	0.47
2017	5071581	2083774	0.41
2018	5967553	2083349	0.35
2019	6714106	2280826	0.34
2020	14209366	4661713	0.33
2021	10478835	3776035	0.36
2022	14788522	5738745	0.39
2023	14926785	5197866	0.35
Total	89091120	34054299	0.40

Table 2: Summary statistics on volume of tweets and share of original tweets versus retweets

Table 3: Share of tweets by women, by topic

Dependent Variable: Share of Tweets by Women			
Gender	0.136***	0.115***	0.139***
	(0.00622)	(0.0198)	(0.00618)
Year = 2014 × Gender		0.0291	
		(0.0289)	
Year = 2015 × Gender		0.00327	
		(0.0285)	
Year = 2016 × Gender		0.0293	
		(0.0284)	
Year = 2017 × Gender		0.0487*	
		(0.0281)	
Year = 2018 × Gender		0.0320	
		(0.0280)	
Year = 2019 × Gender		0.0409	
		(0.0279)	
Year = 2020 × Gender		0.00424	
		(0.0273)	
Year = 2021 × Gender		0.0202	
		(0.0271)	
Year = 2022 × Gender		0.00778	
		(0.0272)	
Year = 2023 × Gender		.	
		.	
Global users (× 10,000)			0.0167***
			(0.00214)
Mean Dependent Variable	0.357	0.357	0.357
Adjusted R2	0.114	0.114	0.127
N	4059	4059	4059

*Notes: The unit of observation is a hashtag-year. Regression includes data from years 2014 to 2023. Standard errors in parentheses, clustered at the level of the hashtag: *** p<0.01, ** p<0.05, * p<0.1*

Table 4: Retweets of Women vs Men, by Topic

Dependent Variable: ln(mean retweets of women) - ln(mean retweets of men)		
Share female	1.693***	2.025***
	(0.126)	(0.167)
gender	0.229***	0.176***
	(0.0486)	(0.0556)
MENA		0.121**
		(0.0525)
social		0.00471
		(0.0535)
racial		-0.144**
		(0.0680)
LGTBQ		0.121
		(0.107)
religion		-0.0229
		(0.0692)
youth		-0.165*
		(0.0905)
education		-0.00236
		(0.123)
economic		0.0742
		(0.0514)
health		-0.239**
		(0.0934)
COVID		0.0948
		(0.141)
climate		-0.0926
		(0.120)
political		-0.0138
		(0.0512)
security		-0.0336
		(0.0464)
entertainment_and_lifestyle		-0.0717
		(0.0735)
Mean Dependent Variable	-0.257	-0.257

Table 5: Retweets of Women vs Men, by Measures of social support

Dep Var	Ln(retweets of women)			Ln(retweets of men)		
clustering (all)	0.0209***			0.0122***		
	(0.00433)			(0.00448)		
clustering (women)		0.0455***			0.00578	
		(0.00825)			(0.00718)	
clustering (men)		-0.0177**			0.0166**	
		(0.00871)			(0.00758)	
degree (women)			0.0670***			-0.0210***
			(0.00416)			(0.00365)
degree (men)			-0.0287***			0.0515***
			(0.00418)			(0.00367)
Mean Dependent Variable	-1.983	-1.983	-1.983	-1.974	-1.976	-1.976
Adjusted R2	0.0308	0.0376	0.126	0.0235	0.0287	0.103
N	3969	3942	3942	4032	3942	3942

*Notes: The unit of observation is a hashtag-year. Regression includes data from years 2014 to 2023. Standard errors in parentheses, clustered at the level of the hashtag: *** p<0.01, ** p<0.05, * p<0.1*

Table 6: Women's participation and social support

Dep Var	Share of women participating		
gender	0.119***	0.124***	0.123***
	(0.00509)	(0.00447)	(0.00455)
MENA	-0.190***	-0.186***	-0.193***
	(0.00449)	(0.00395)	(0.00403)
clustering (all)	0.0223***		
	(0.00220)		
clustering (women)		0.0295***	
		(0.00376)	
clustering (men)		-0.00550	
		(0.00400)	
degree (women)			0.0182***
			(0.00200)
degree (men)			-0.0158***
			(0.00201)
Mean Dependent Variable	0.357	0.361	0.361
Adjusted R2	0.410	0.484	0.477
N	4059	3942	3942

*Notes: The unit of observation is a hashtag-year. Regression includes data from years 2014 to 2023. Standard errors in parentheses, clustered at the level of the hashtag: *** p<0.01, ** p<0.05, * p<0.1*