Supporting Information for #polisci Twitter: A Descriptive Analysis of how Political Scientists Use Twitter in 2019

James Bisbee*, Jennifer M. Larson†, Kevin Munger‡

Contents

*Postdoc, Princeton Niehaus. jhb362@nyu.edu. Corresponding author.
†Associate Professor, Vanderbilt. jennifer.larson@vanderbilt.edu
‡Assistant Professor of Political Science and Social Data Analytics, Pennsylvania State University. kevinmunger@gmail.com

A Technical Details from Main Paper

This section contains a more detailed explanation of the statistical procedures employed in the body of the text, excised for readability.

A.1 Figure 9

We run an interacted regression with the following specification:

$$\begin{aligned} m_{i,j} = \beta_0 + \beta_1 G_i + \beta_2 G_j + \rho_1 G_i \times G_j + \\ + \beta_3 T_i + \beta_4 T_j + \rho_2 T_i \times T_j + \\ + \beta_5 I_i + \beta_6 I_j + \rho_3 I_i \times I_j + \mathbf{S}_{i,j} + \epsilon_{i,j} \end{aligned} \tag{1}$$

where $m_{i,j}$ is the number of mentions of alter j by ego i, G is a gender indicator, T is a tenure indicator, and I is an ideology indicator. $\mathbf{S}_{i,j}$ is a vector of school-level measures (tuition, enrollment, rank) that are measured as the absolute difference between ego i's school and alter j's school. The main coefficients of interest are the ρ's which capture the interaction effect of sharing an identity with another scholar. We implement dyad cluster-robust standard errors via multiway decomposition, as described in Aronow, Samii and Assenova (2015).

A.2 Figure 10

We bootstrap sample our egos and calculate the Euclidean distance between their vector of "friends" (accounts that they follow) and the friend vectors for every other ego in the data. We then subset the data to the bottom quartile of these Euclidean distances and estimate the same interacted regression described above in Equation 1.

A.3 Novel Measures of Research Influence

To fix notation, define each user on #polisci Twitter as $i \in \mathcal{N}$ and each tweet containing original research as t_i.

Other Twitter users $j \neq i$ can engage with t_i in one of four ways.

- First, they can re-tweet t_i: $rt_j^{t_i}$.
- Second, they can reply to t_i: $rp_j^{t_i}$.
- Third, they can quote t_i: $q_j^{t_i}$.
- Fourth, they can favorite (or "like") t_i: $f_j^{t_i}$

All four of these engagements will be broadcast to j's followers and friends in their notifications. The likelihood that each of these engagements appears to the engager's follower's is not constant, and depends on the details of Twitter's proprietary algorithm.

Re-tweets and quote tweets are by far the most likely to be appear, as expected, but likes and replies sometimes appear as well.

Ideally, we would want to fully trace the way t_i is disseminated throughout the Twitter network, such that j's quote is re-tweeted by k which is replied to by l which is re-tweeted by n and so forth. Unfortunately, the way the Twitter API records data precludes this type of detailed tracing for re-tweets, the most common form of engagement in our data. Specifically, all re-tweets are assigned to the original tweet, regardless of whether j is re-tweeting the original tweet or is instead re-tweeting k's re-tweet. In other words, we can only observe $rt_j^{t_i}$, but not $rt_j^{rt_k^{t_i}}$.

This limitation only applies to re-tweets. We can observe $rt_j^{rp_k^{t_i}}$ and $rt_j^{q_k^{t_i}}$. Furthermore, the same logic doesn't apply to favorites of re-tweets, meaning that $f_j^{rt_k^{t_i}}$ is observable. Nevertheless, the inability to observe how re-tweets spread throughout the network is a non-trivial limitation for our ability to trace how research is disseminated, since 79% of engagements with original research take this form.

A.4 Network Graph

The full follower network is visualized in Figure 1, where users are colored by gender and tenure status, with directed links reflecting follows. Nodes are sized by the logged number of tweets associated with each account. We visualize the network using a radial axis layout in `Gephi` that groups nodes using degree centrality and organizes the groups radiating out from a central circle.

The full network contains many scholars – roughly 11% of the Twitterverse – who do not follow other Political Scientists. We drop these nodes for visual clarity but emphasize that these are not dormant accounts. Rather, these scholars are online and active but not connected to our #polisci Twitterverse. We treat the follower network as the antecedent for many of the behaviors we measure. Choosing who to follow defines what tweets one is most likely to see, and circumscribes the tweets that might then be engaged with.

B Quality of Life Online

Our manuscript focuses on statistics of the #polisci twitterverse suggesting that scholars on Twitter engage in behavior that both elevates voices of under-represented groups in academia, but also exhibits homophilous tendencies common in social networks. But all of the manuscript's results are primarily based on countable phenomena. Follows, re-tweets, and mentions can be summed, network centrality scores can be estimated, and communities based on these links can be detected. These different measures help us understand the shape of #polisci Twitter, neatly summarizing both behaviors and experiences in terms of these numeric quantities. But these measures are less useful in describing the subjective experience of using the platform.

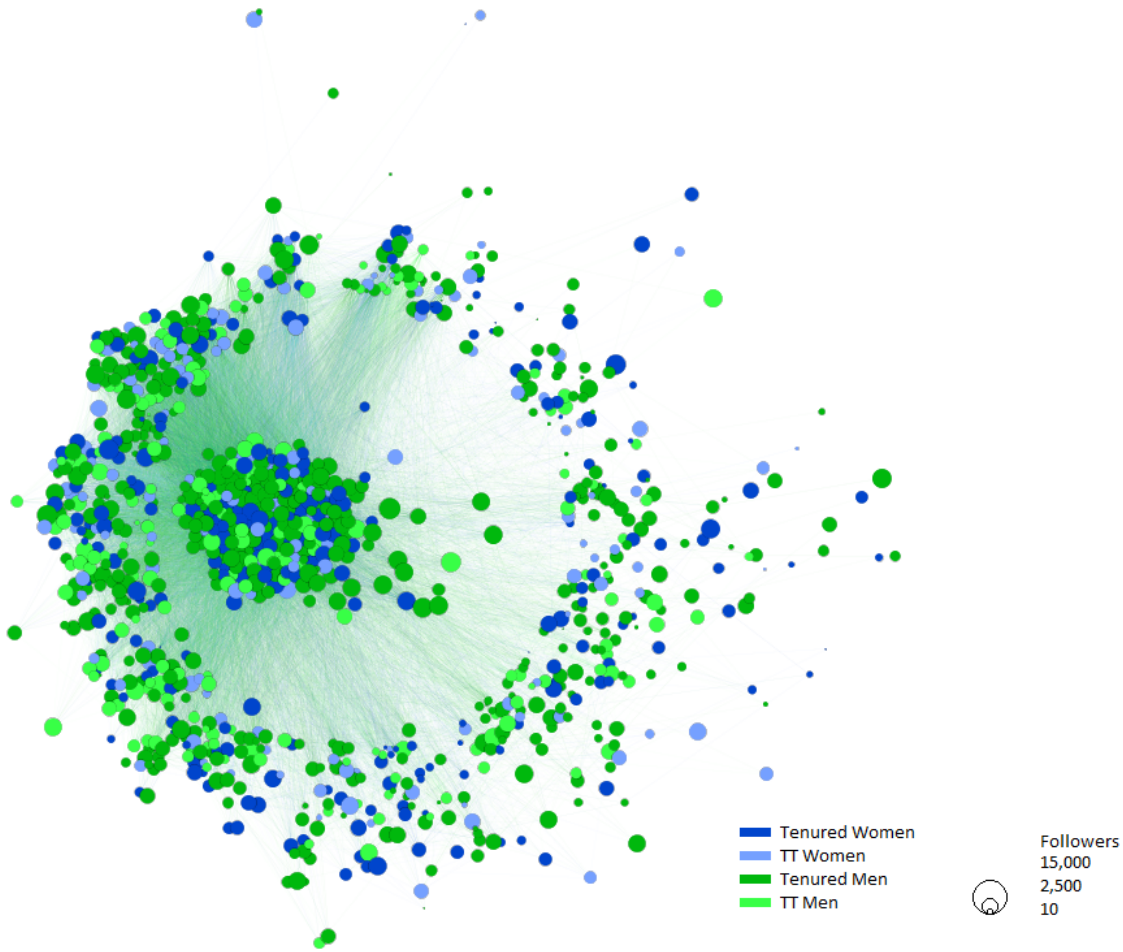

Figure 1: Full #polisci Twitterverse, as of April 2019. Nodes are size according to logged tweets as of April 2019. Nodes are colored to reflect male and female scholars (green and blue respectively) in tenure-track and tenured positions (light and dark shades respectively).

There are several reasons to be interested in the subjective experience of being a scholar in the #polisci twitterverse. For example, a well-cited framework for understanding female underrepresentation use the metaphor of a "leaky pipeline" in which women are more likely than men to leave academia for one reason or another (Bos and Schneider, 2012; Breuning et al., 2018). While these reasons vary, a common explanation is simply that the experience of being a female scholar is more burdensome, due to additional stresses such as micro-aggressions in the workplace, greater work required to achieve the same goals, and disproportionate shares of academic "service", to name a few. In the following section, we use natural language processing to examine whether the experience of being a female academic online constitutes an additional source of stress that can further develop our understanding of the leaky pipeline.

Specifically, we use the `peRspective` package for `R` which evaluates text for toxic content via a deep neural network, trained on hundreds of thousands of human-coded examples. Toxicity is divided into a variety of dimensions, ranging from profanity to attacks on identity, insults, threats, to sexually explicit content. These dimensions are based on human-coded content that is used to train the AI which takes as an input a small set of sentences (typically a comment, although it is easily amenable to tweets) and returns the probability that a human would consider this content objectionable for one or more of the provided dimensions. We apply this algorithm to each tweet issued by a scholar on #polisci Twitter that mentions another scholar on #polisci Twitter.

Overall, #polisci Twitter is a fairly polite environment. Figure 2 plots the distribution of the probability that a tweet is toxic in #polisci Twitter alongside a similar distribution of randomly sampled tweets by the entire Twitter population, recorded during the 2016 election (theoretically a time in which the average Twitter user was more likely to be discussing politics). On average, #polisci Twitter is significantly less toxic than the general public on October 19th, 2016. In addition, the general public have a much more pronounced bimodal distribution, with the majority of tweets being very benign but a handful being extremely toxic. Although the low end of the #polisci twitterverse is slightly more toxic than the low end of the general public, there is not nearly as much highly toxic tweeting among political scientists.

To describe the experience of being on #polisci Twitter, we subset the data by gender, position, and ideology and compare the distributions of toxicity among subgroup. We further divide these subgroups by who is doing the mentioning, allowing us to assess both the overall experience of being in a particular group and comparing how much of the toxicity is coming from members within the group versus outside it. We plot the average probability of being mentioned in a toxic tweet in gray bars in Figure 3 for men and women, tenured and tenure-track faculty, and liberals and conservatives. Below these overall averages, we disaggregate by the source of the mention. For example, we see that tweets that mention women are approximately 1% more likely to be toxic than tweets that mention men. But we also see that tweets mentioning women written by women are equally likely to be toxic as tweets written by men. Conversely, tweets mentioning men written by men are significantly more likely to be toxic than tweets written by women, but both probabilities are lower than tweets mentioning women.

Taken together, these results suggest that exposure to online toxicity is higher among women, tenured scholars, and liberals. Furthermore, the results indicate that these same groups are more responsible for the toxic content that their peers experience, although this is likely confounded by the homophily results discussed above. But it is important to remember that these results do *not* indicate that subsets of the #polisci twitterverse are exposed to horrendous toxicity on a daily basis. Rather, these results indicate that all groups enjoy a very low level of toxicity online, with the differences between groups being substantively small, albeit statistically significant.

Yet when it comes to toxicity, perhaps the overall average is not the measure of interest. It might be that one particularly toxic tweet can sour a scholar's experience online. An alternative measure is to extract most toxic tweets faced by everyone in our data and compare these distributions in Figure 4. Here we find much larger substantive differences, with males being mentioned in tweets written by men whose maximum

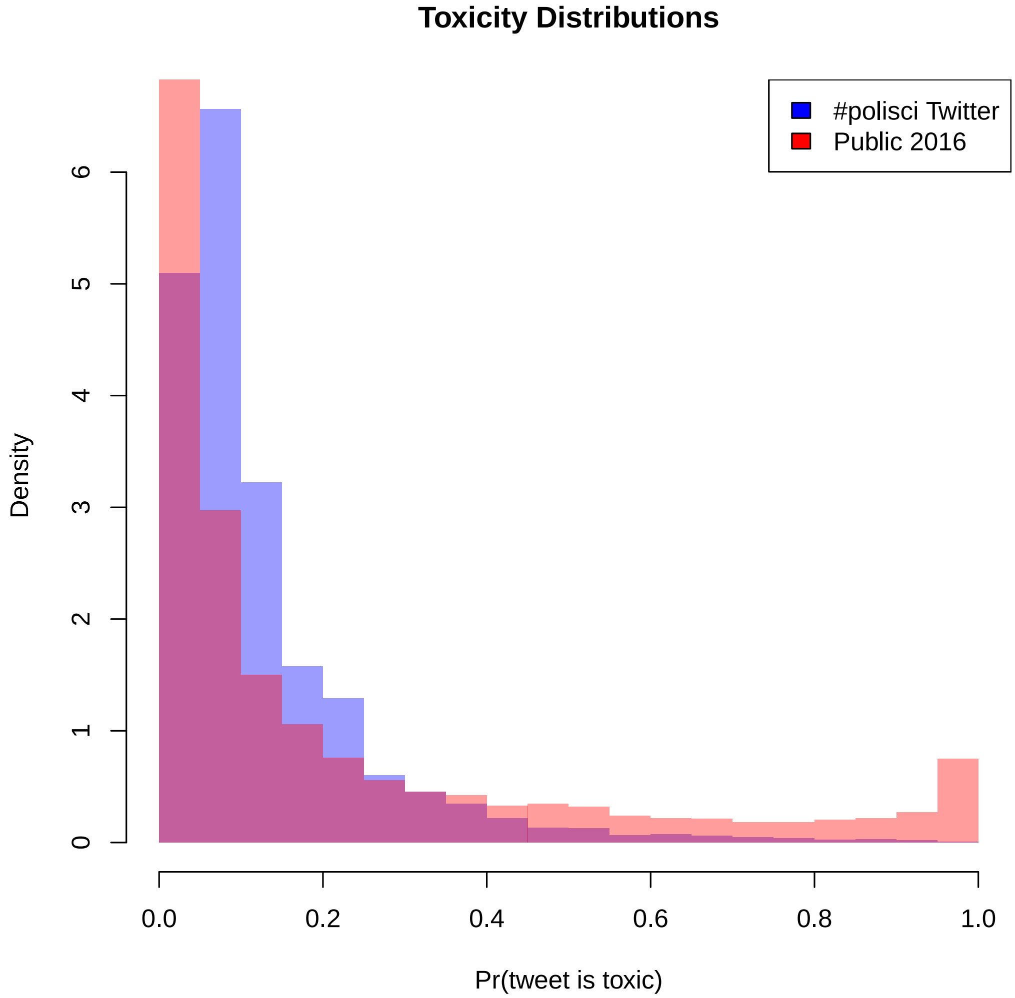

Figure 2: Distribution of toxicity measures across #policsi Twitter and the general public on October 19th, 2016.

probability is almost 10 percentage points more likely to be toxic than tweets written by women. But while this difference is statistically significant, very few of the other differences are with the exception of tweets written by liberals versus conservatives.

Note also that the maximum probability of toxicity results summarized in Figure 4 again point to the relatively non-toxic discourse that characterizes #polisci Twitter. Taking 0.5 as the threshold probability of identifying toxic tweets (i.e., tweets that are more likely than not of being toxic), less than 3% of political scientists ever faced a tweet that was more than 50% likely to be toxic. Furthermore, inspecting the most toxic tweets in the data reveals that some of them are not directed at the scholar that they mention. For example, the most likely to be toxic tweet in our data (98% chance of being toxic) was a reply to a tweet by @blumrm who was describing her worst

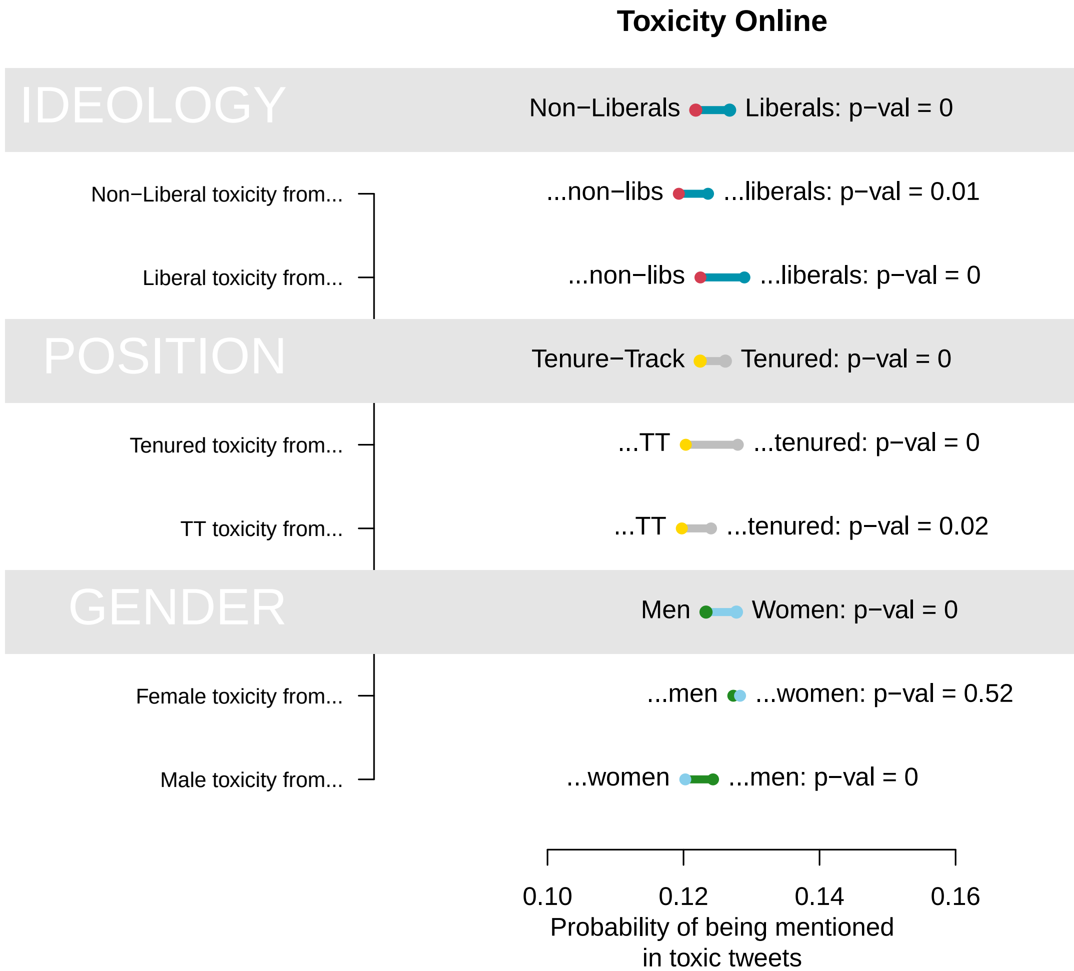

Figure 3: Average probability of being mentioned in a toxic tweet by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text. Probabilities calculated as average across all tweets.

experience in academia. The reply – "@blumrm Makes me want to punch someone in his fat f***ing face" – is clearly intended as a sign of support and solidarity, not an attach on @blumrm herself. Insofar as stronger language provides clearer signals of solidarity, it may not be accurate to assume that this measure of toxicity is perceived as such by those who are mentioned in tweets.

The `peRspective` API provides access to other categories beyond toxicity. Specifically, we can examine the probability that a tweet mentioning an author is considered sexually explicit, an attack on the author, profane, or an insult. We plot these results in the figures below, examining both the average probability and the maximum probability of being mentioned in a type of tweet. Across measures, we note that, while

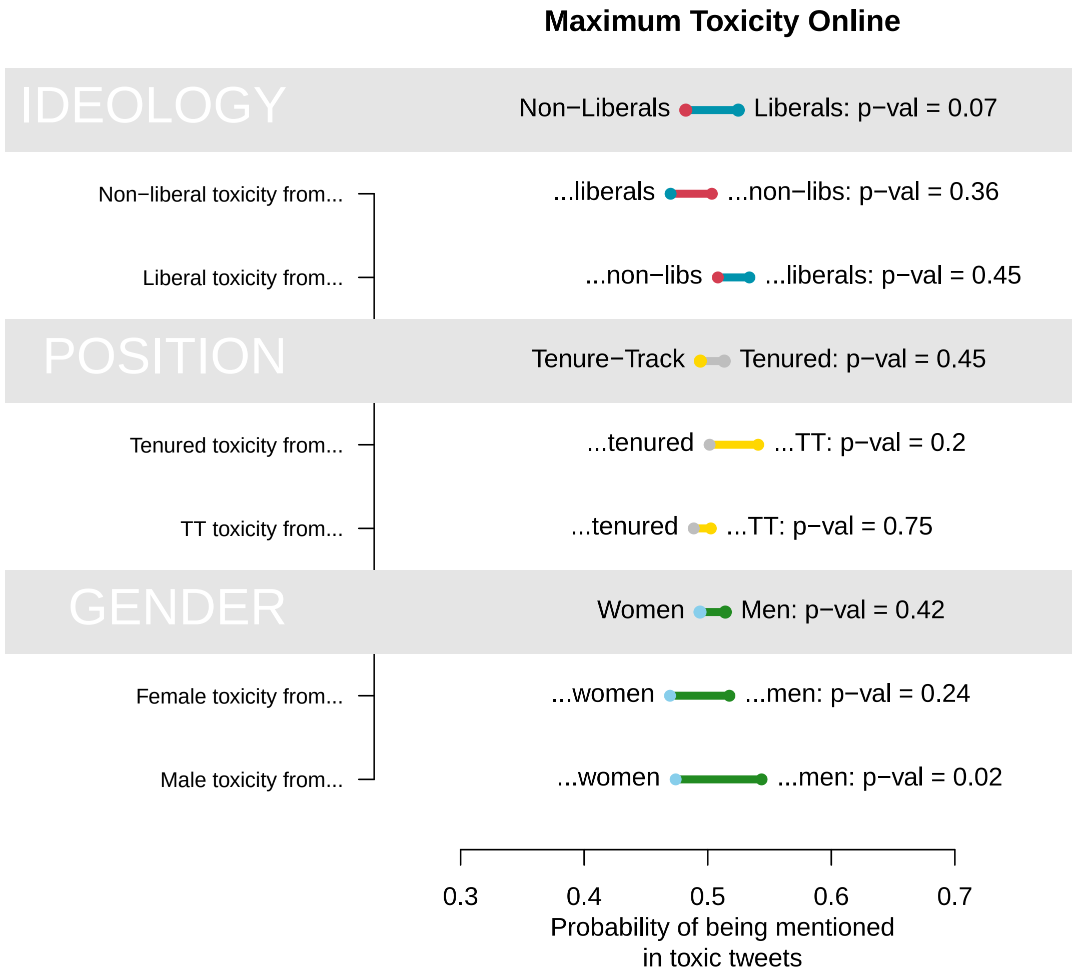

Figure 4: Average probability of being mentioned in a toxic tweet by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text. Probabilities calculated as average of most toxic tweets experienced by scholar.

women are on average more likely to be mentioned in these types of tweets, men are more often mentioned in the tweets most likely to be sexually explicit, an attack on the author, profane, or insulting.

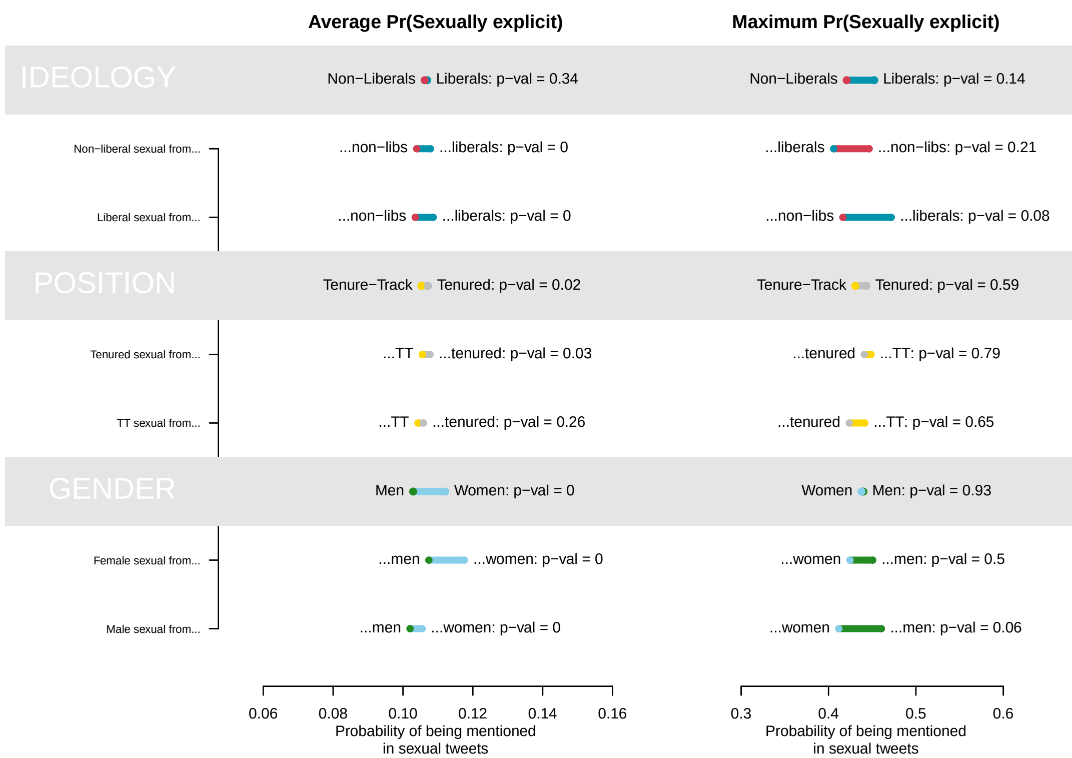

Figure 5: Average (left panel) and maximum (right panel) probability of being mentioned in a sexually explicit tweet by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text.

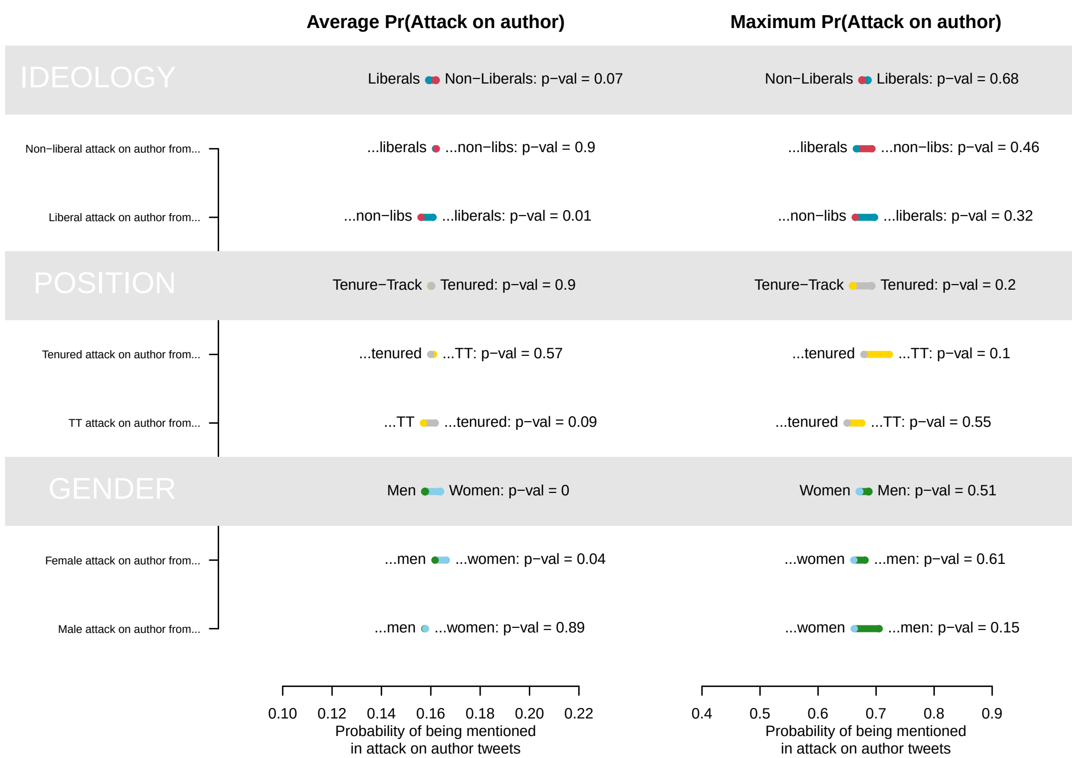

Figure 6: Average (left panel) and maximum (right panel) probability of being mentioned in a tweet that attacks the author by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text.

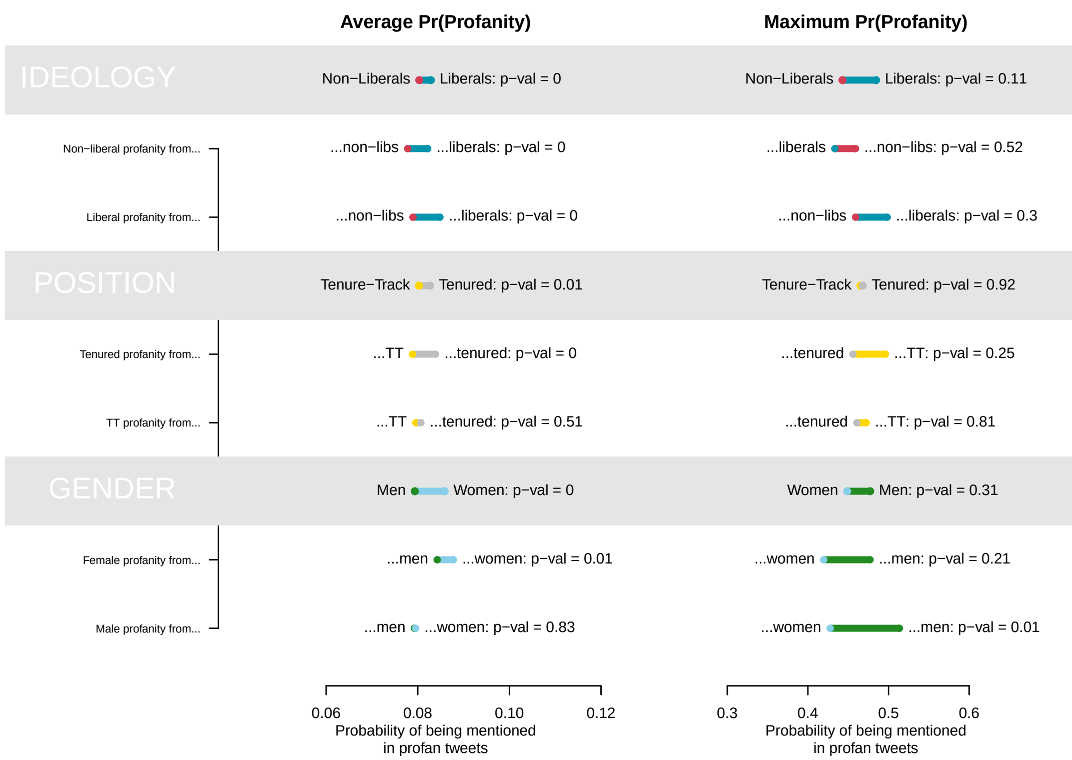

Figure 7: Average (left panel) and maximum (right panel) probability of being mentioned in a profane tweet by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text.

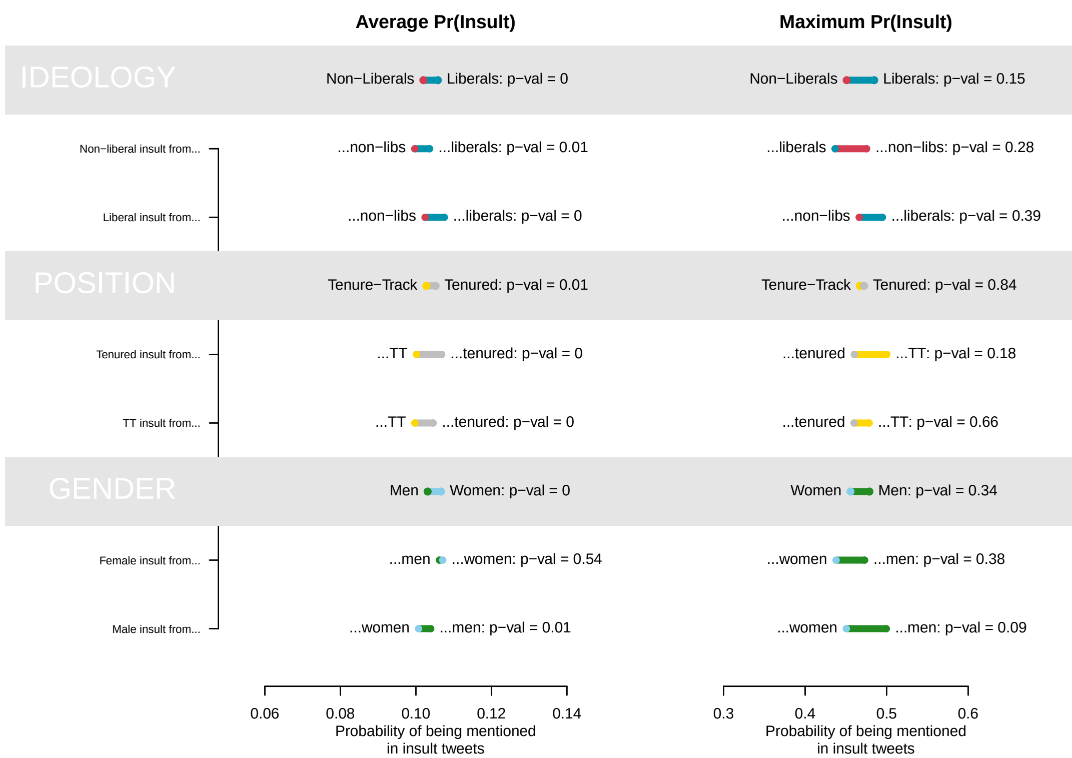

Figure 8: Average (left panel) and maximum (right panel) probability of being mentioned in an insulting tweet by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text.

C Heterogeneity across "also know" engagement

Our main results group all scholars together and examine correlations in relevant behavior across observed characteristics. However, we are able to formulate some coarse priors about which among these scholars would be more likely to actively work against homophilies based on researcher identity. Specifically, we look for all mentions of the two most prominent hashtags associated with amplifying underrepresented scholars in the discipline: #womenalsoknowstuff and #pocalsoknowstuff.

Do the patterns of homophily we document in the full sample persist when we subset our analysis to scholars who use these hashtags? We start by describing the difference between this subset at the sample writ large in Figures 9, 10, and 11. As illustrated, the gender imbalance is more apparent, particularly among tenured faculty. In addition, there is clear evidence of an ideological bias in the direction we would expect, with scholars who engage with the AlsoKnow hashtags being more liberal than those who don't. Finally, the tenor of the conversation among the #AlsoKnow scholars is more consistent, as illustrated by the distributions of toxicity among both subsets.

We re-analyze all of our main results, finding much weaker evidence of homophily among the scholars who mention either "also know" group one or more times. Or more accurately, while in-group homophily persists (i.e., men are more likely to be mentioned by other men), the cross-group results are attenuated (i.e., women are no longer less likely to be mentioned by men). The figures and table below reproduce those found in the manuscript, subsetting to these scholars.

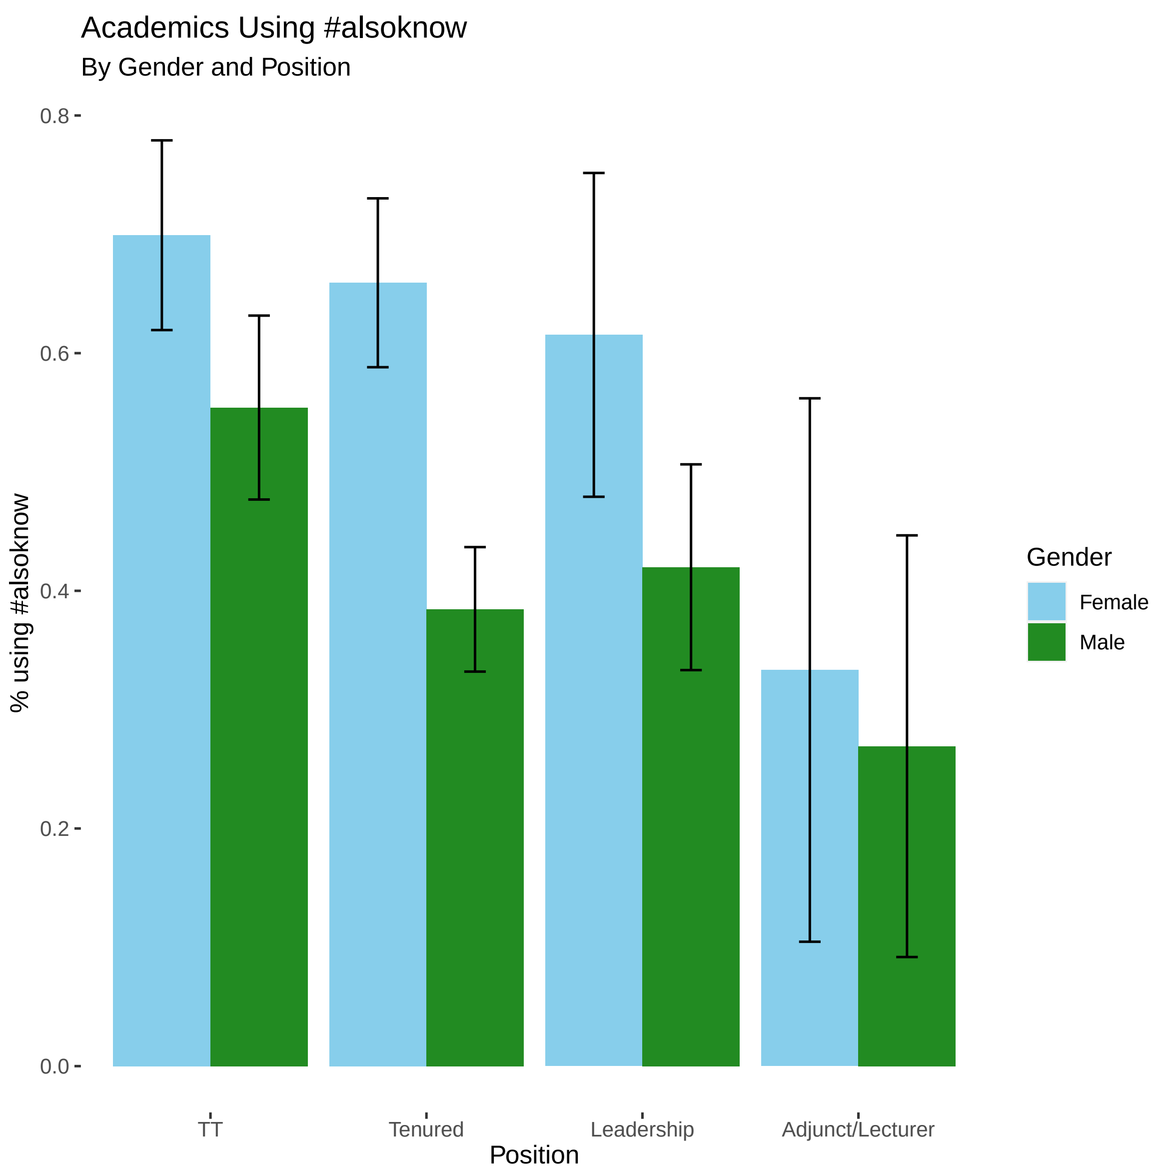

Figure 9: Who engages with the #AlsoKnow categories by gender (bars) and position (x-axis).

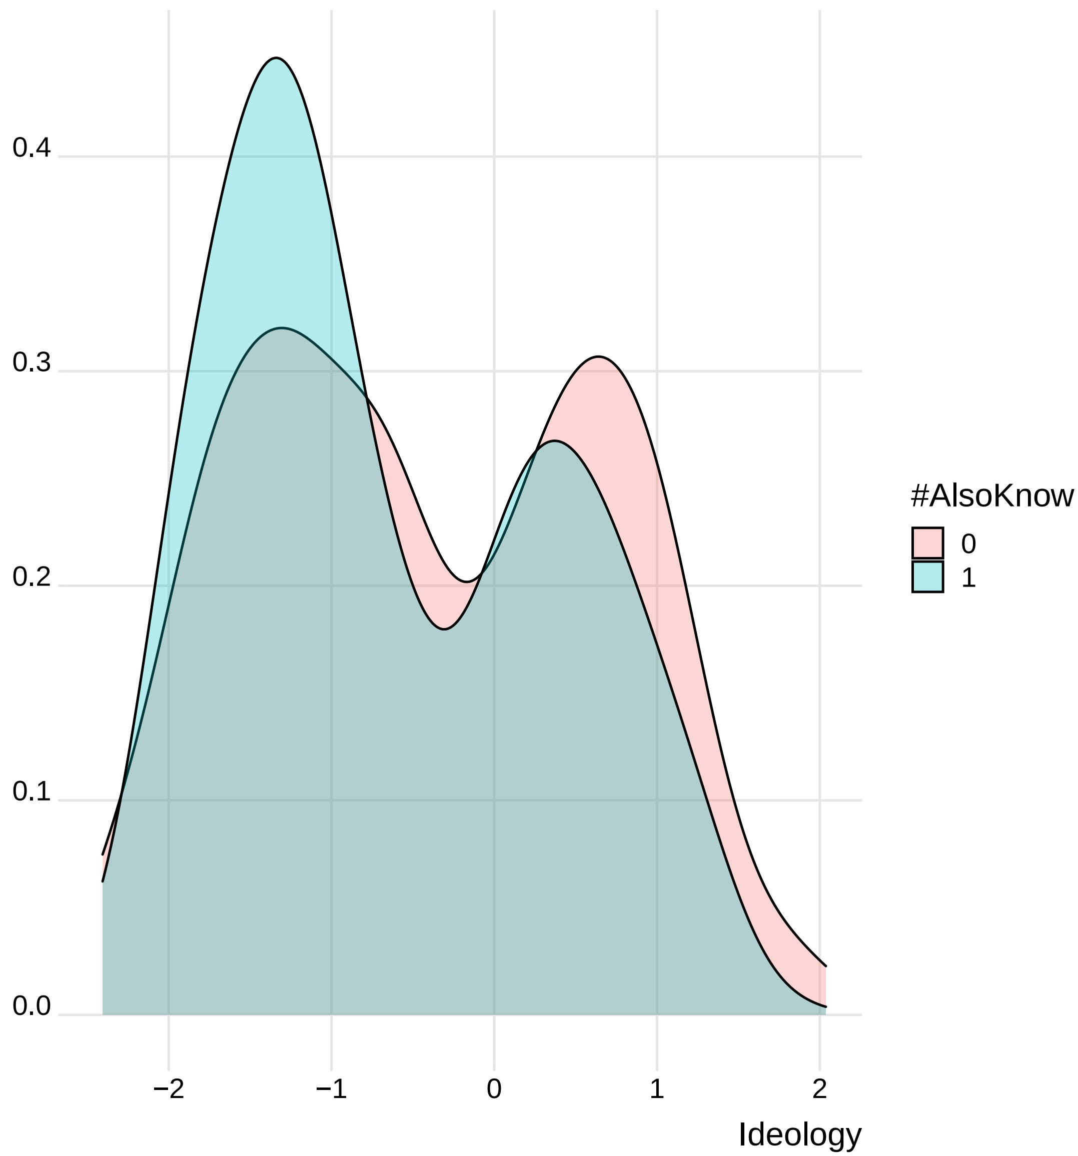

Figure 10: Ideology by #AlsoKnow subset.

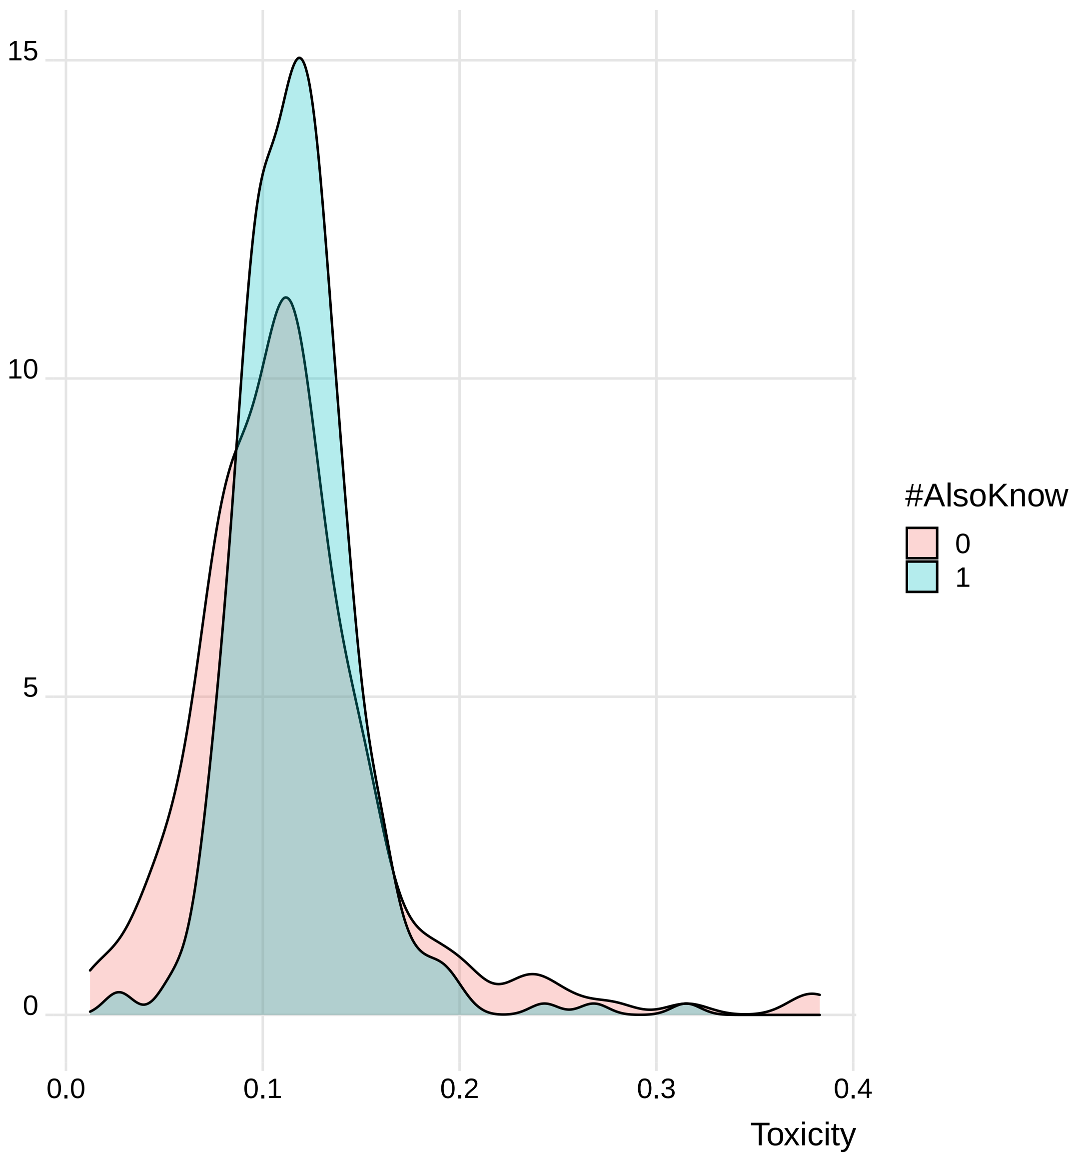

Figure 11: Toxicity by #AlsoKnow subset.

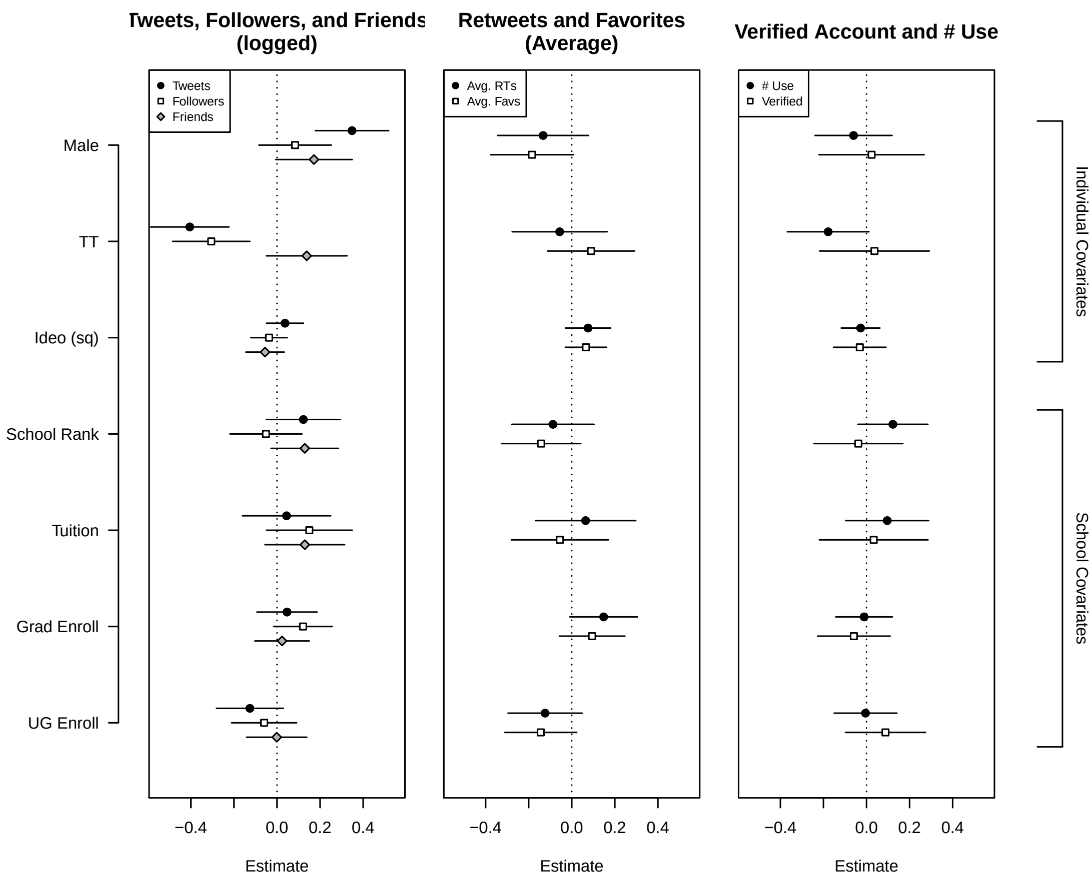

Figure 12: Correlates of behavior online among the subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

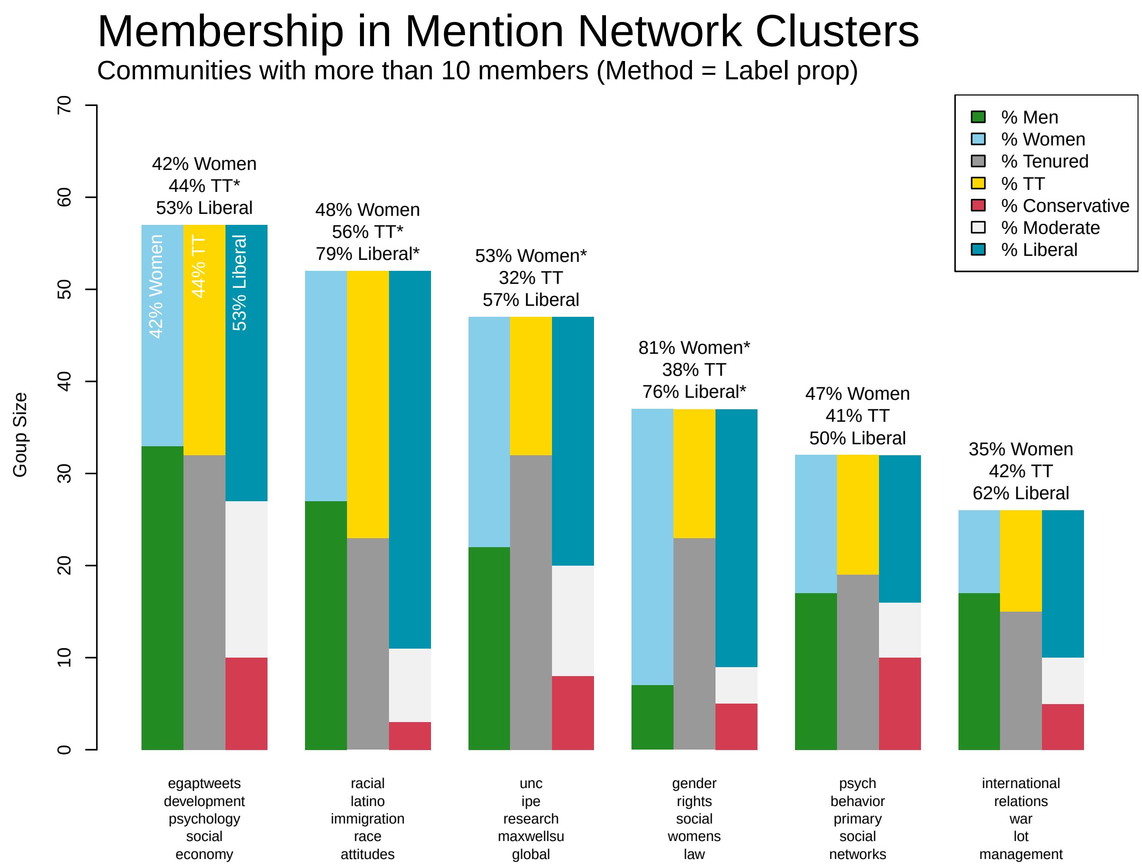

Figure 13: Communities from Network Analysis. Bars indicate the number of accounts assigned to each community where links are defined by mentions using the majority label proportion algorithm. Bars are colored according to the share of each community that is female (blue) and male (green), tenure-track (gold) and tenured (gray), and by ideology. Communities are labeled by the top five discriminating words (TF-IDF) based on text analysis of the concatenated Twitter bios of all members. Data subset to only those scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

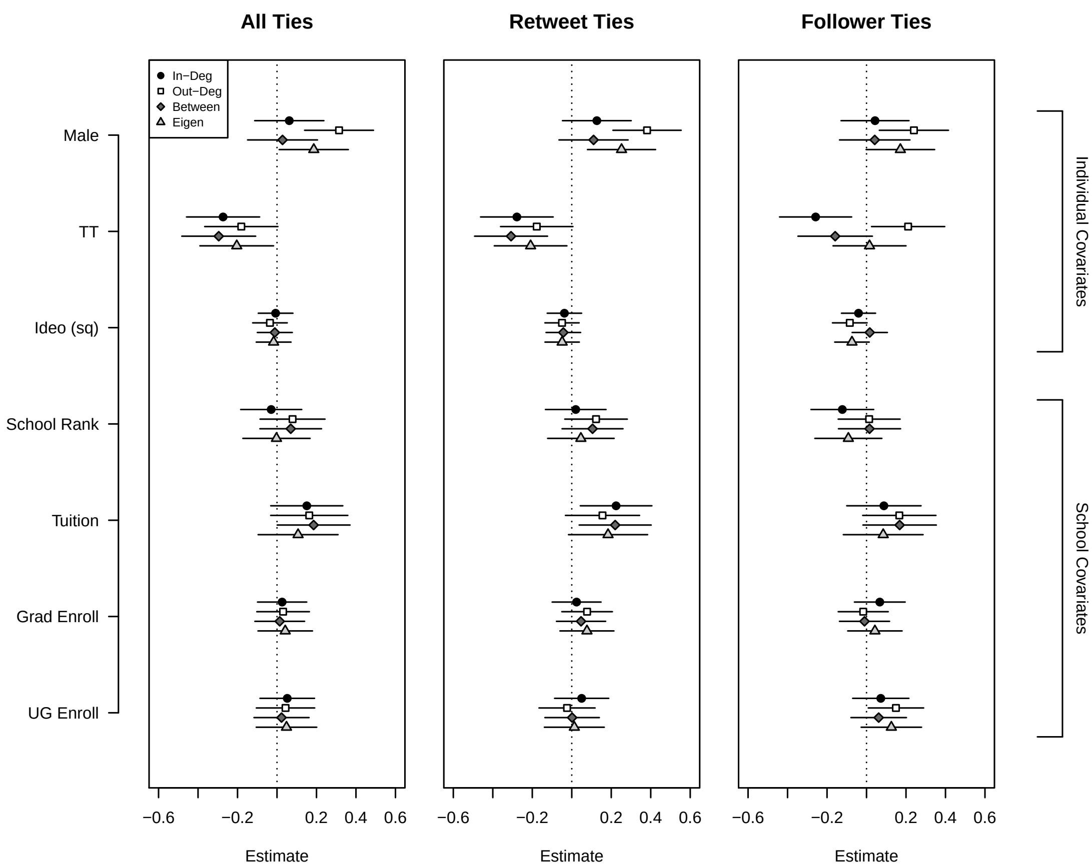

Figure 14: Correlates of node centrality among the subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

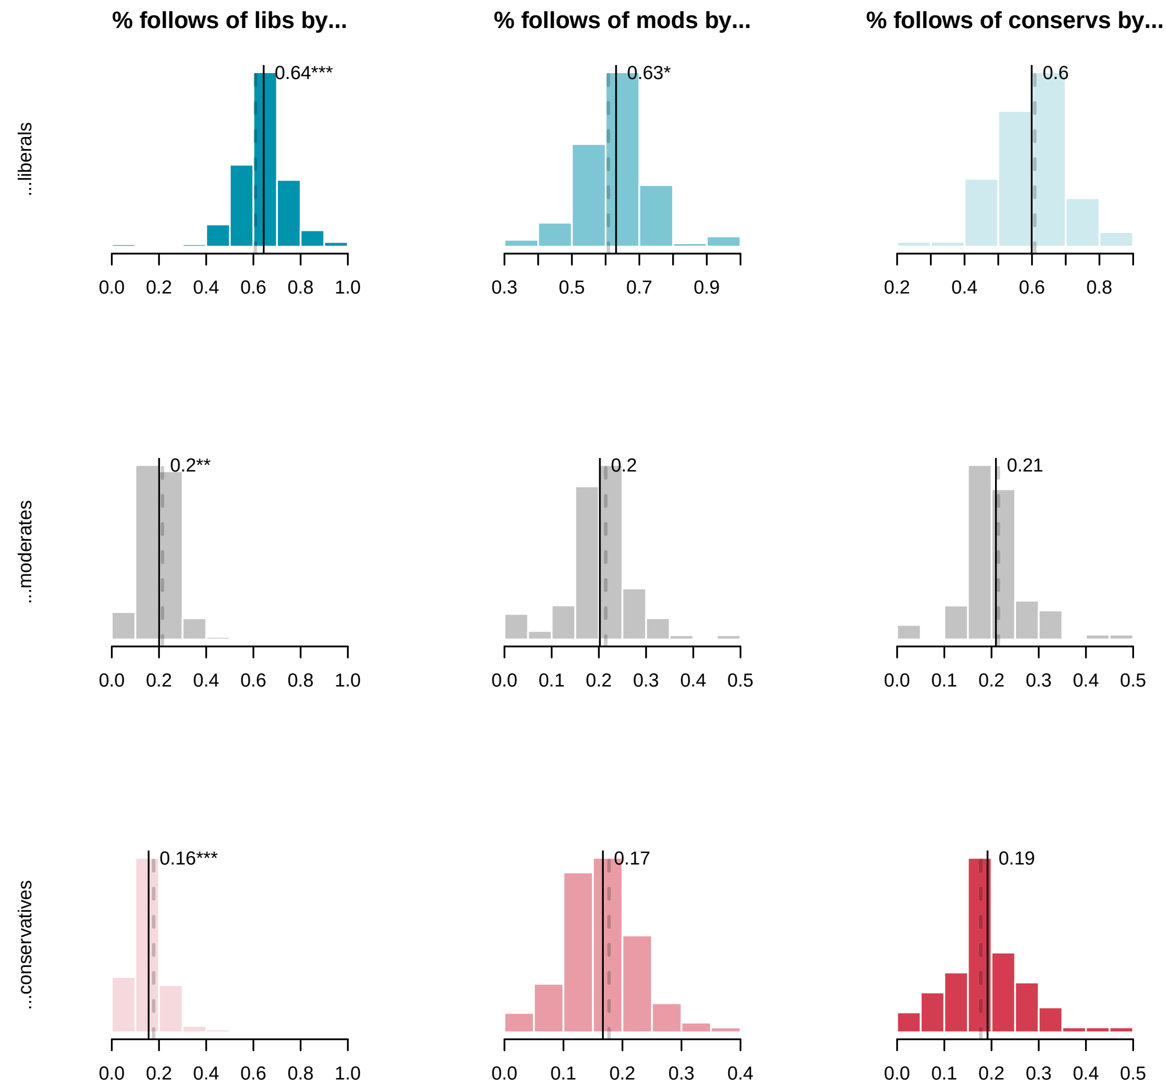

Figure 15: Mentions by ideology. Histograms reflect distribution of share of a user's mentions that come from other users of a given ideology. Columns indicate the ideology of the mentioned users. Rows indicate the ideology of the mentioning user. Vertical solid lines capture the mean of the distribution, given by the black text. Vertical dashed lines indicate the share of the group in the overall data. Significance stars indicate whether a simple t-test difference between the distributions mean and the share of the group in the overall data is significant at conventional levels. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

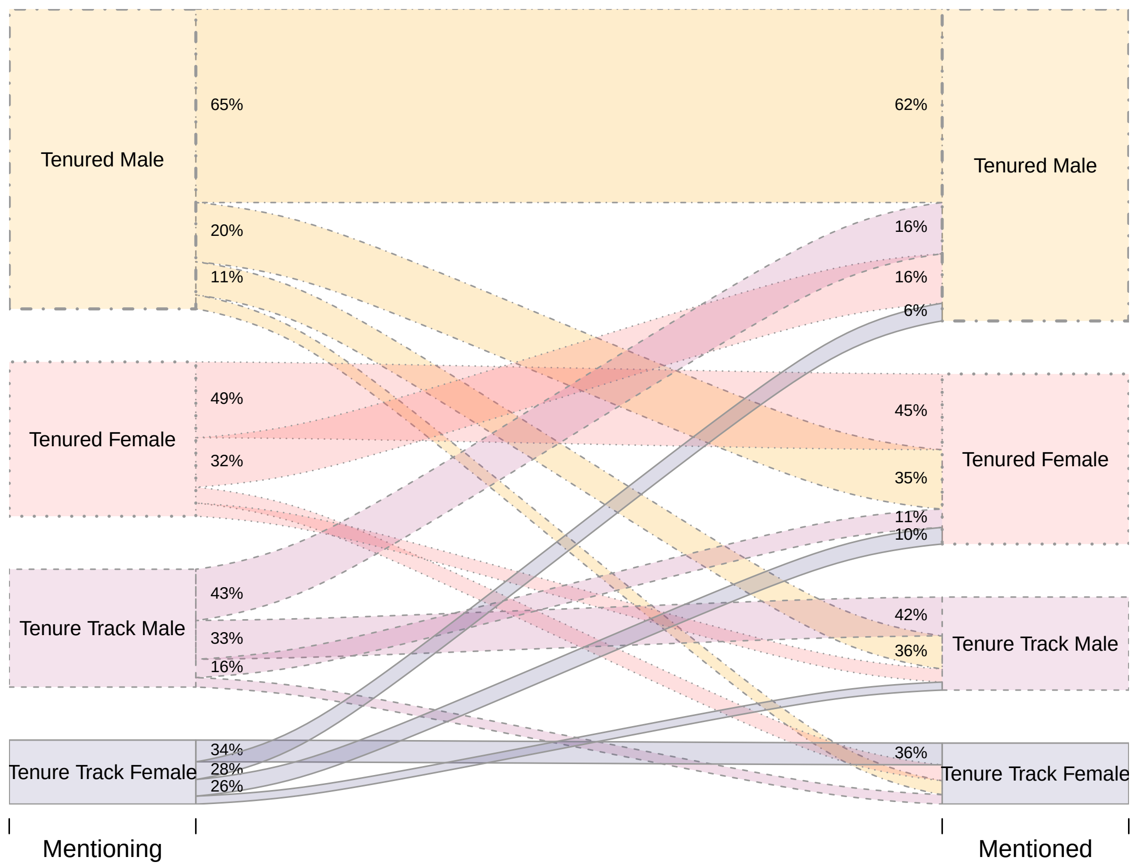

Figure 16: Mentions by position and gender. The left column describes the mentioning behavior by different groups, with the flows charting who they mention. Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

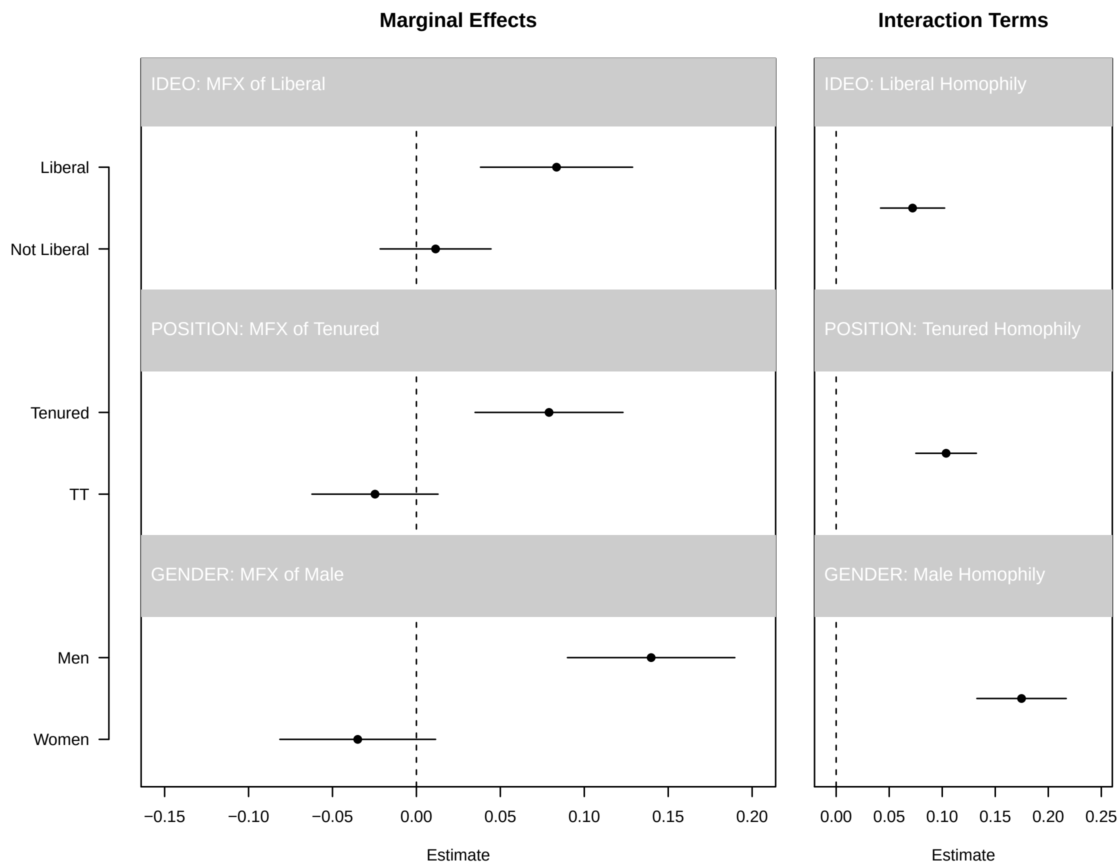

Figure 17: Interaction marginal effects (left panel) and interaction terms (right panel) estimated using dyadic data. Bars indicate two dyad-robust standard errors calculated via multiway decomposition (Aronow, Samii and Assenova, 2015). Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

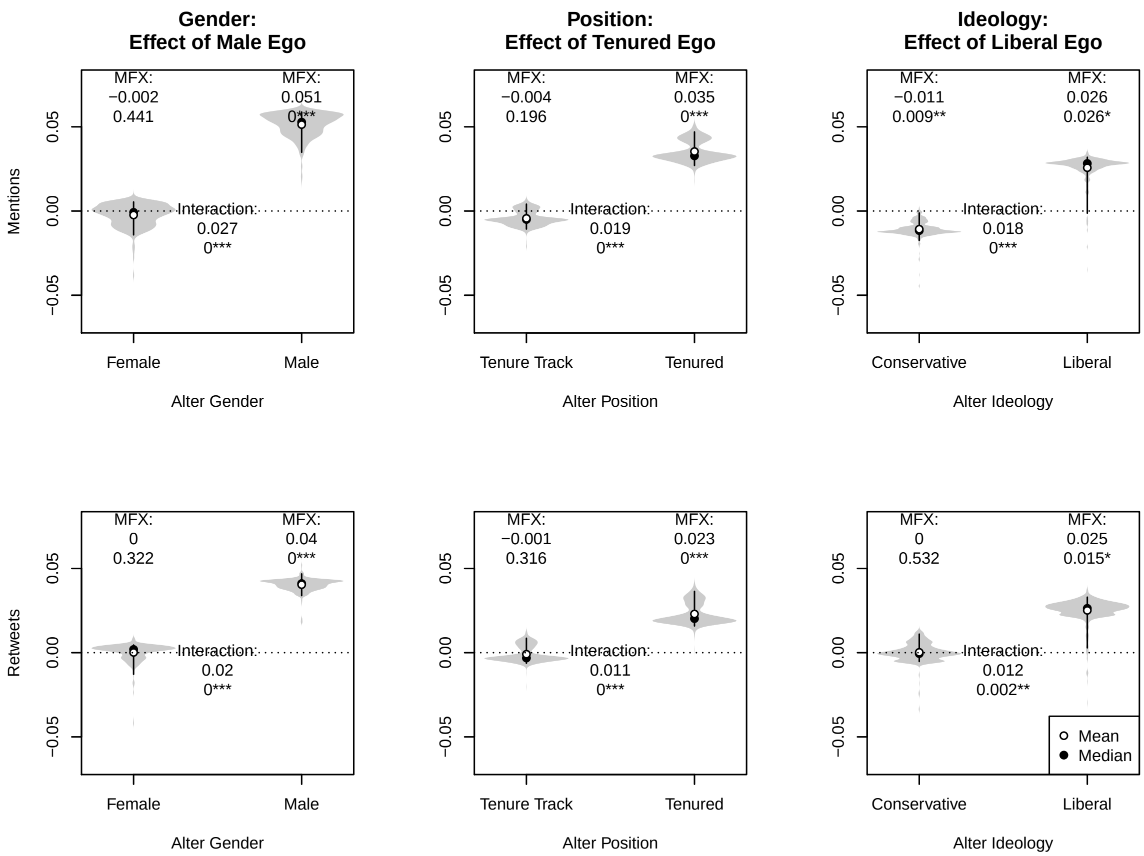

Figure 18: Dyadic Analyses Controlling for Followers. Bootstrapped marginal effects displayed as gray densities with average coefficient and p-value indicated in text. Mean and median bootstrapped estimated indicated by white and black circles, respectively. Bars represent the 95% confidence interval between 2.5% and 97.5%. Mean interaction coefficient and p-value also indicated in text in center of each plot. * p < 0.05, ** p < 0.01, *** p < 0.001. Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

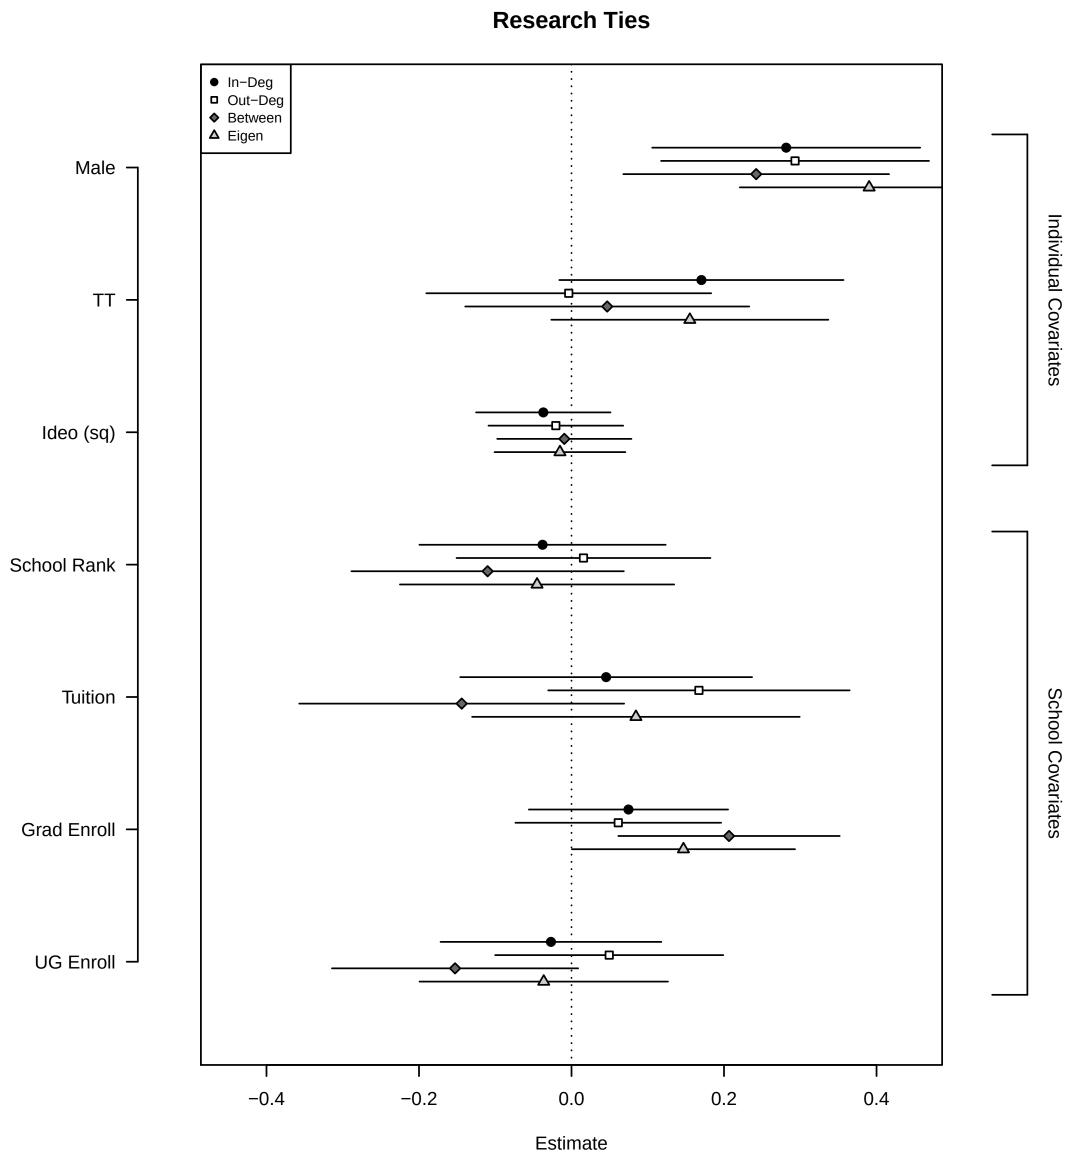

Figure 19: Centrality measures for links defined as research shares. Coefficients (points) and two standard errors (bars) based on multilevel regression of centrality measure on individual-level and school-level covariates (y-axis). Centrality measures given in the legend. Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

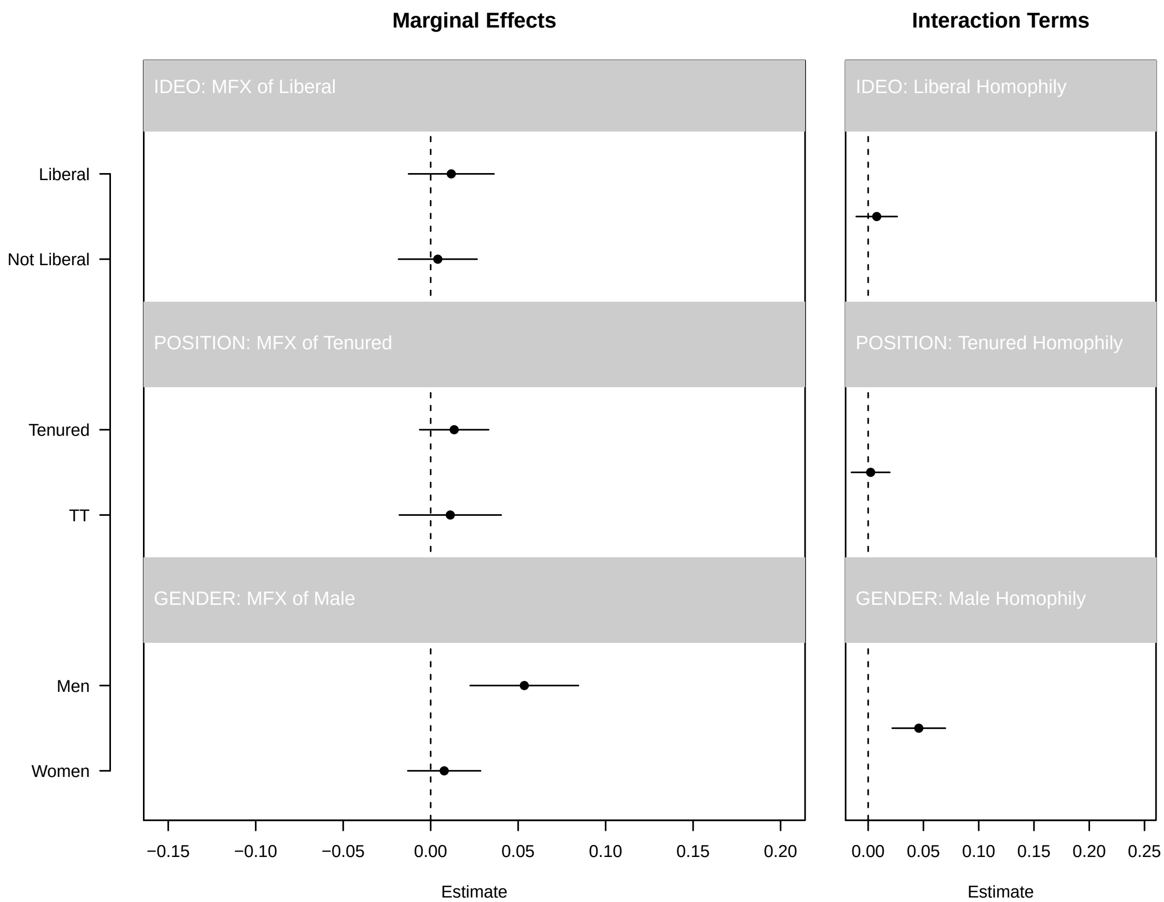

Figure 20: Interaction marginal effects (left panel) and interaction terms (right panel) estimated using dyadic data. Bars indicate two dyad-robust standard errors calculated via multiway decomposition (Aronow, Samii and Assenova, 2015). Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

Table 1: Research Dissemination

	Shares	Dissemination of Research			
	% Research (1)	# RTs (2)	# Favs (3)	Engage (4)	Follows (5)
Male	0.102**	0.108	0.164	0.506*	1.666**
	(0.039)	(0.184)	(0.171)	(0.227)	(0.518)
Tenure Track	0.063	0.233	0.289	0.195	0.301
	(0.040)	(0.191)	(0.178)	(0.236)	(0.539)
Moderate	0.029	0.070	−0.018	−0.015	0.024
	(0.048)	(0.225)	(0.209)	(0.278)	(0.633)
Conservative	0.041	−0.361	−0.438	−0.095	−0.095
	(0.052)	(0.243)	(0.227)	(0.300)	(0.685)
Years Online	0.022	−0.093	−0.136	−0.184	−0.237
	(0.020)	(0.095)	(0.089)	(0.118)	(0.269)
School Rank	0.008	0.261	0.134	0.003	0.066
	(0.030)	(0.138)	(0.127)	(0.169)	(0.385)
R1 Inst.	−0.084	−0.030	0.028	−0.221	−0.247
	(0.075)	(0.351)	(0.325)	(0.430)	(0.982)
State School	−0.139	−0.270	0.235	−0.226	−0.888
	(0.071)	(0.333)	(0.308)	(0.408)	(0.930)
Grad Enroll	0.009	0.218	0.214	0.030	0.066
	(0.027)	(0.126)	(0.117)	(0.155)	(0.353)
Observations	523	442	442	442	442
Log Likelihood	−324.109	−905.727	−875.587	−996.507	−1,350.298

Notes: Patterns of sharing research on Twitter. First column regresses share of tweets that contain link to research. Ensuing columns regress measures of how popular these tweets are overall (columns 2 and 3) and how much of an impact they make on Political Science Twitter (columns 4 and 5). * p $\leq$ .05, ** p $\leq$ .01, *** p $\leq$ .001. Subset of scholars who engage with either the #womenalsoknowstuff or the #pocalsoknowstuff hashtags.

D Citations and Network Centrality

To what degree are our findings correlated with a more traditional metric of scholarly success – citations? Put more simply, are the scholars who are more centrally placed in the Twitter network the same as those with more citations? Do citations and followers predict each other?

To answer these questions, we scraped the Google scholar profiles for all scholars in our dataset. We caution that this scraping procedure is imperfect, although we believe that whatever errors are made in the course of scraping are orthogonal to our quantities of interest. Furthermore, we emphasize that these results, as with those presented in the paper, are purely correlational. It may be that citations cause more engagement on Twitter, or it may be that more engagement on Twitter causes more citations. More plausibly, engagement on Twitter and citations are both caused by unobservable researcher qualities.

Nevertheless, we regress a scholar's logged total citations and her "recent" (since 2016) citations on her position in the #polisci Twitter network. The correlation coefficients of these analyses are presented in Figure 21. As illustrated, if anything there is a *negative* relationship between citations and network centrality. We believe this reflects the characteristics of the most active participants on #polisci Twitter, who are younger and therefore have fewer citations.

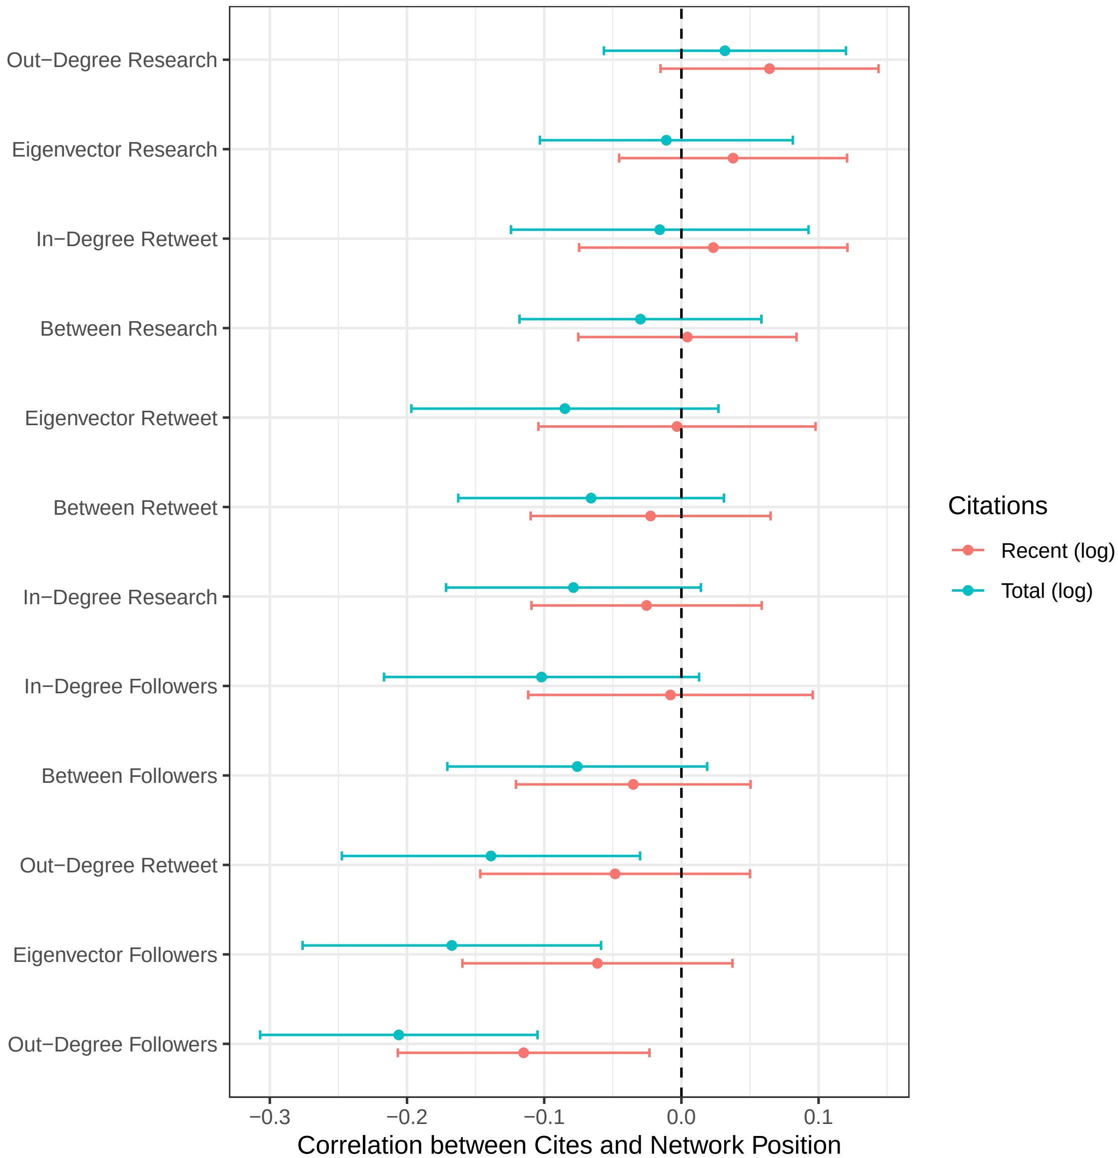

Figure 21: Correlation coefficients relating logged citations (type given by color) with network centrality (y-axis). Bars indicate two standard errors. Model controls for scholar and school characteristics.

E Robustness to Ivies and Productivity

The main results how that one of the few significant positive correlations is between the size of the graduate school program and the number of followers a scholar has on Twitter. We theorized that this may reflect greater participation of graduate students on #polisci Twitter, who follow their professors. However, a plausible alternative explanation is that the scholars at these larger programs are located at more prestigious universities and / or are more productive. While our main results control for school rank, we test an alternative measure of "prestige" by creating a dummy indicator for whether the institution is among the ivies. In addition, we proxy for scholar productivity with their number of recent citations (since 2016). As illustrated in Figure 22, there is compelling evidence to suggest that productivity and the logged number of followers are positively related, although we note that the positive coefficient on the size of the graduate student body persists. There is no systematic relationship between whether the scholar is at one of the Ivy league institutions and any of these metrics.

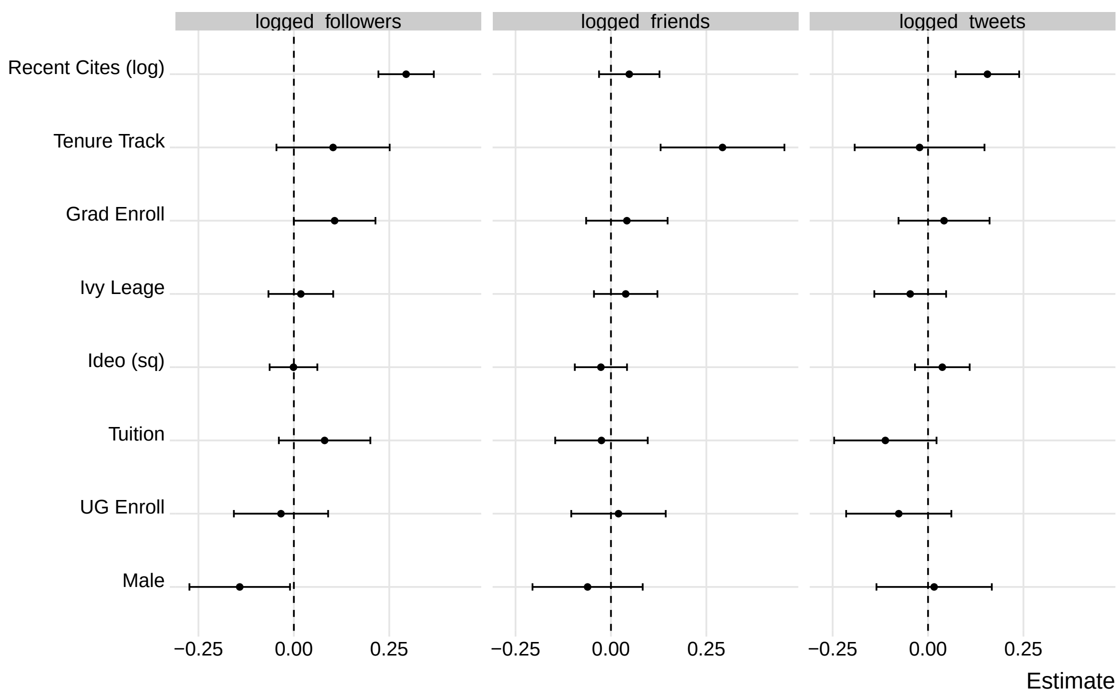

Figure 22: Correlation coefficients relating logged followers (left plot), logged friends (center plot), and logged tweets (right plot) with politician and school characteristics (y-axes). Bars indicate two standard errors. Coefficients generated by multilevel model nesting respondents within schools. School rank replaced with ivy league dummy. Recent citations (logged) used as a proxy for productivity.

F Lecturers and Adjuncts

The main findings group tenured scholars with directors, chairs, provosts, deans, and others in leadership positions, and drop adjuncts / lecturers. These groups are based on the observation that scholars within each group exhibit conceptual and empirical similarities to each other, but differ across groups. In other words, tenure-track scholars are similar to adjuncts, and both differ from tenured scholars and those with leadership positions. Nevertheless, we note that adoption of Twitter differs most meaningfully between tenure-track scholars and adjuncts / lecturers. As such, we replicate our main findings in Table 2 by disaggregating to the four categories of career, finding that our main results are not meaningfully changed. The hold-out category is the leadership position.

Table 2: Career Categories

	Disaggregated Career Categories				
	Tweets (ln) (1)	Followers (ln) (2)	Friends (ln) (3)	# Use (4)	Verified (5)
Male	0.069 (0.064)	0.015 (0.062)	0.013 (0.064)	−0.102 (0.065)	0.080 (0.065)
Adjunct/Lect	0.114 (0.168)	−0.158 (0.163)	0.110 (0.166)	0.492^{**} (0.168)	0.147 (0.168)
TT	−0.183 (0.094)	−0.133 (0.091)	0.327^{***} (0.093)	−0.158 (0.094)	−0.049 (0.094)
Tenured	−0.062 (0.085)	0.041 (0.082)	0.060 (0.084)	−0.094 (0.086)	0.026 (0.085)
Ideo (sq)	0.032 (0.031)	−0.003 (0.030)	−0.007 (0.031)	−0.039 (0.031)	−0.003 (0.031)
School Rank	0.096 (0.063)	−0.021 (0.068)	0.122^{*} (0.056)	0.091 (0.062)	−0.041 (0.056)
Tuition	−0.035 (0.075)	0.045 (0.081)	0.095 (0.066)	0.021 (0.073)	0.048 (0.065)
Grad Enroll	0.093 (0.049)	0.175^{***} (0.053)	0.058 (0.043)	0.065 (0.048)	−0.039 (0.043)
UG Enroll	−0.061 (0.057)	−0.082 (0.061)	0.040 (0.050)	−0.071 (0.056)	−0.005 (0.050)
Constant	0.022 (0.088)	−0.004 (0.087)	−0.124 (0.086)	0.135 (0.088)	−0.055 (0.087)
Observations	1,052	1,052	1,052	1,039	1,052

Notes: * p ≤ .05, ** p ≤ .01, *** p ≤ .001.

G Community Behavior and Cross-Pollination

Our manuscript used a label propagation method to identify communities among scholars online based on who they mention. But we can also examine the degree to which different communities interact with each other by measuring the share of users from one community who are mentioned by another. Figure 23 visualizes these results using a circular flow diagram. Communities are indicated by color and their top 3 TF-IDF terms. Moving clockwise, each community is divided into output ("Retweeted Statuses" or statuses that are retweeted by others) and input ("Statuses Retweeted" or statuses that members of the community retweet). The former percentages are summarized in dark text, the latter summarized in light text. Using the largest community comprised of scholars working on "economy, european, trade", we note that there are 1,422 statuses originating from scholars in this community that are retweeted, and 1,445 statuses that scholars in this community retweet. Of the statuses that originate in this community, the majority (52%) are retweeted by others in the same community.

Looking across all communities, we note a clear within-community bias. With the exception of the "gender, girl, mom" and "public, policy, social" communities, the majority of statuses are retweeted by members of the same community. There is also clear evidence of reciprocity between the "economy, european, trade" community and the others, as illustrated by the fact scholars from this community are in the top two communities for retweeting others' statuses, and are in the top two communities of who these other scholars retweet.

From a purely descriptive standpoint, we know that research moves across retweet communities online, as illustrated in Figure 24. This plot replicates Figure 23 but focuses on tweets about research. Again, we note both the strong in-group preference as well as similar reciprocal behaviors between several communities.

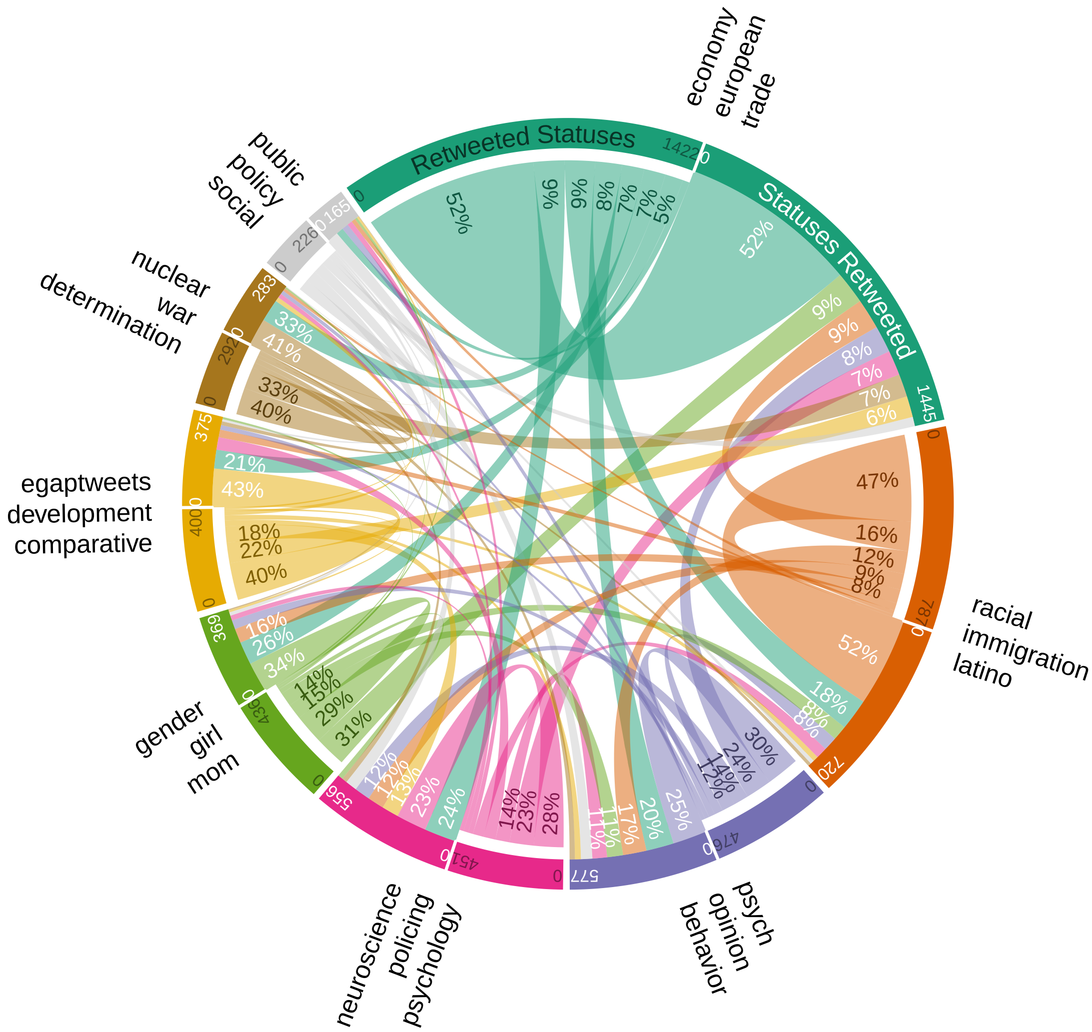

Figure 23: Cross-pollination among Ideological Communities. Links indicate number of members of each community (first bar moving clockwise) that are mentioned by members of other communities (second bar moving clockwise).

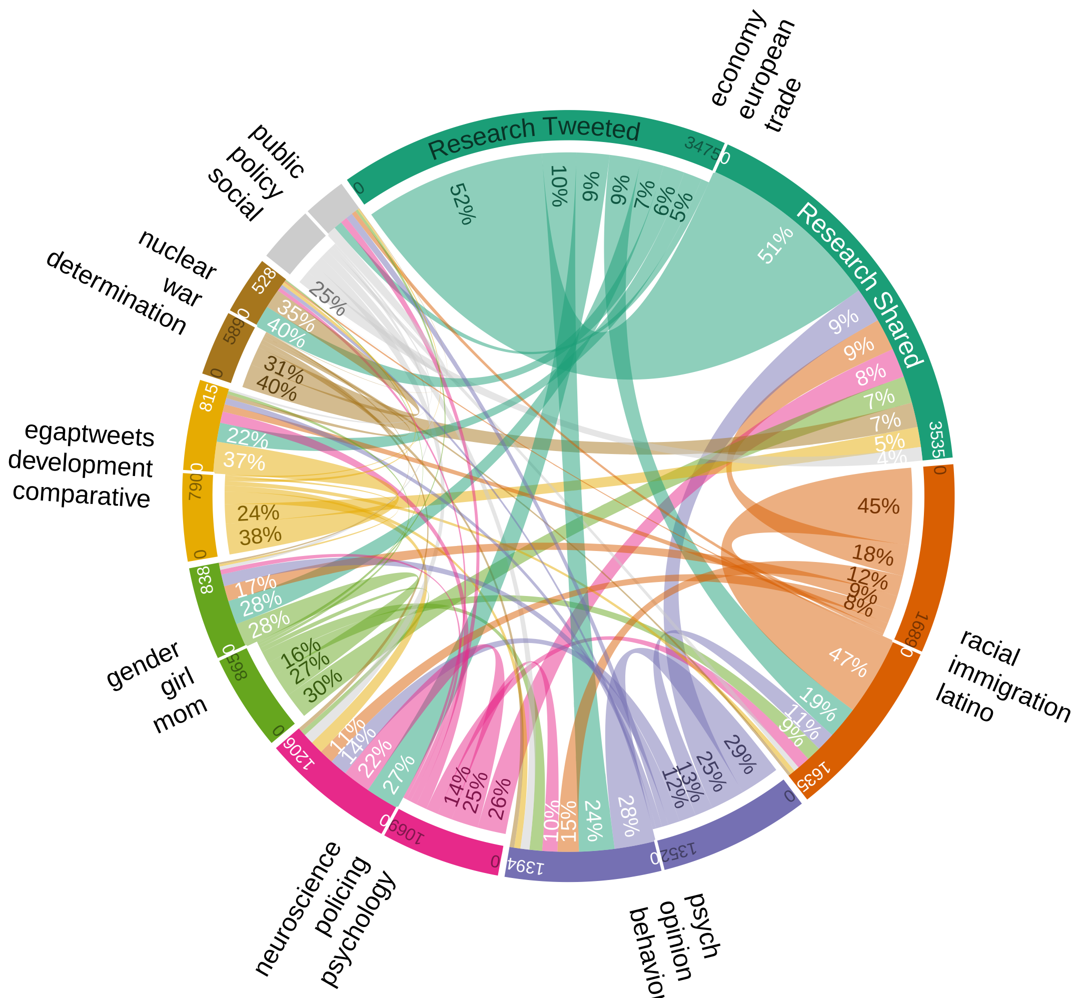

Figure 24: Research dissemination across Ideological Communities. Links indicate volume of tweets linking to research written by members of each community (first bar moving clockwise) that are shared by members of other communities (second bar moving clockwise).

H Homophily within Communities

The main results looked at dimensions of homophily for the sample writ large. In this section, we subset the data by the 8 largest communities returned by the label propagation method. Within each subset, we re-estimate the magnitude of the gender-based homophily and plot the marginal effect coefficients on being a male ego for female and male alters. Positive coefficients suggest that male scholars within these communities are more likely than female scholars to mention male or female alters. As illustrated in Figure 25, there is persistent evidence that gender-based homophily persists across the majority of political science communities on #polisci Twitter, although perhaps reassuringly, none of these marginal effects approach the magnitude of those found in the full data.

The most striking evidence of gender-based homophily is among scholars associated with the community characterized by scholars with the differentiating terms "gender", "girl", "mom", "judicial", and "loves" in their profile bios. Here we see no evidence that male scholars are more likely than female scholars to mention other male scholars. However, we see striking evidence that male scholars are far less likely than female scholars to mention female scholars. This may reflect a bias against female scholars among male scholars in this community, although we suspect a more plausible explanation is that female scholars in this group are significantly more proactive in amplifying other female scholars. These patterns are almost totally attenuated when restricting attention to engagement with research, as illustrated in Figure 26.

Similar patterns obtain when we compare ideological homophily by these notional sub-fields of research interests. Future work that measures the true subfield of the scholars and investigates the field-specific types of homophily that obtain would be valuable.

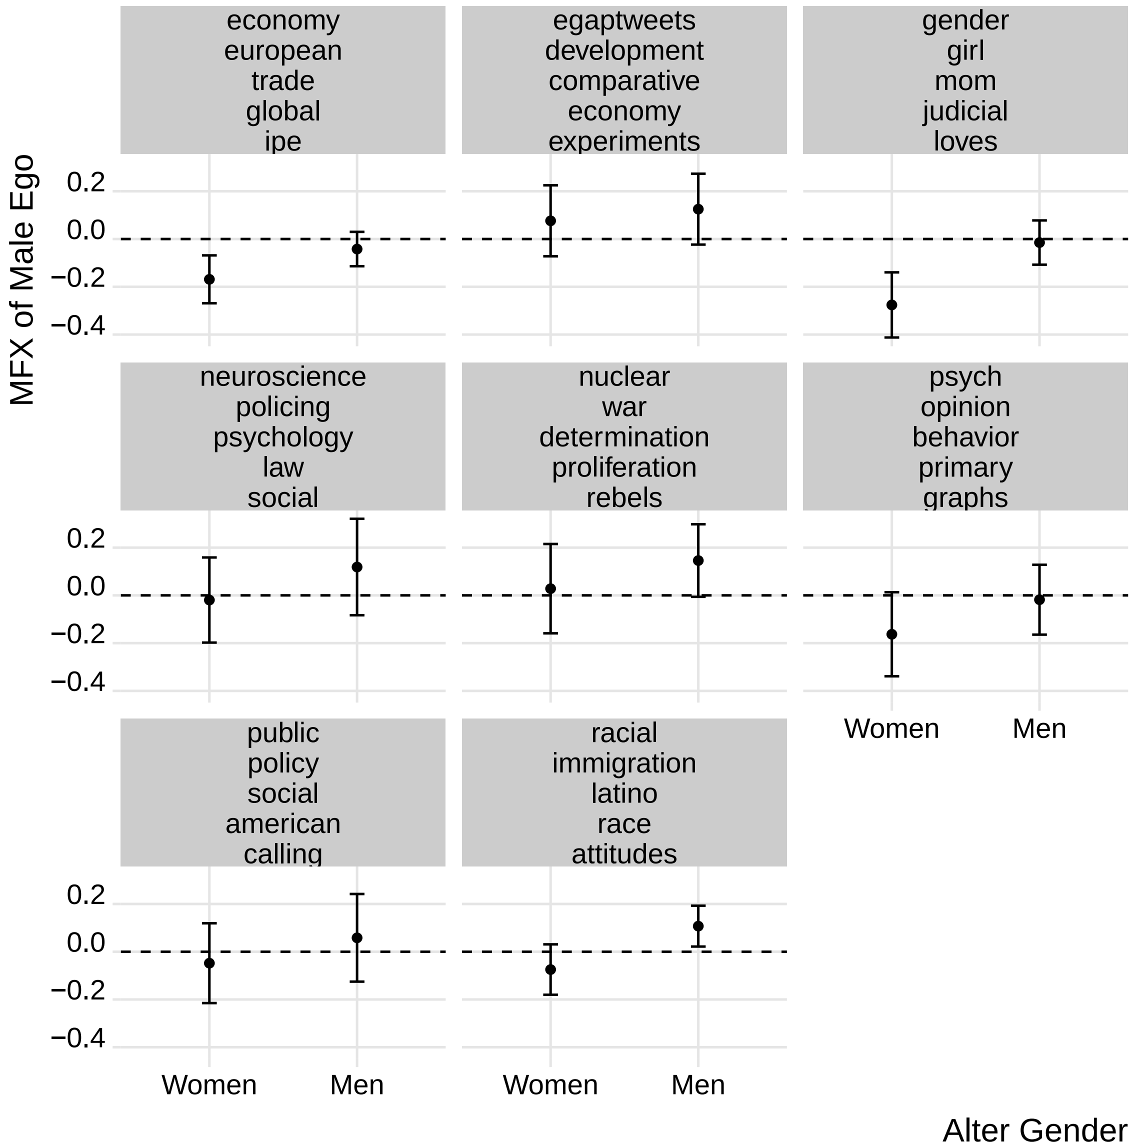

Figure 25: Marginal effect coefficients on male dummy for the ego on the gender of the alter across the 8 largest communities returned by the label propagation method. Behavior is probability of the ego mentioning the alter one or more times. Plot titles depict the top five TF-IDF words appearing in group member Twitter bios.

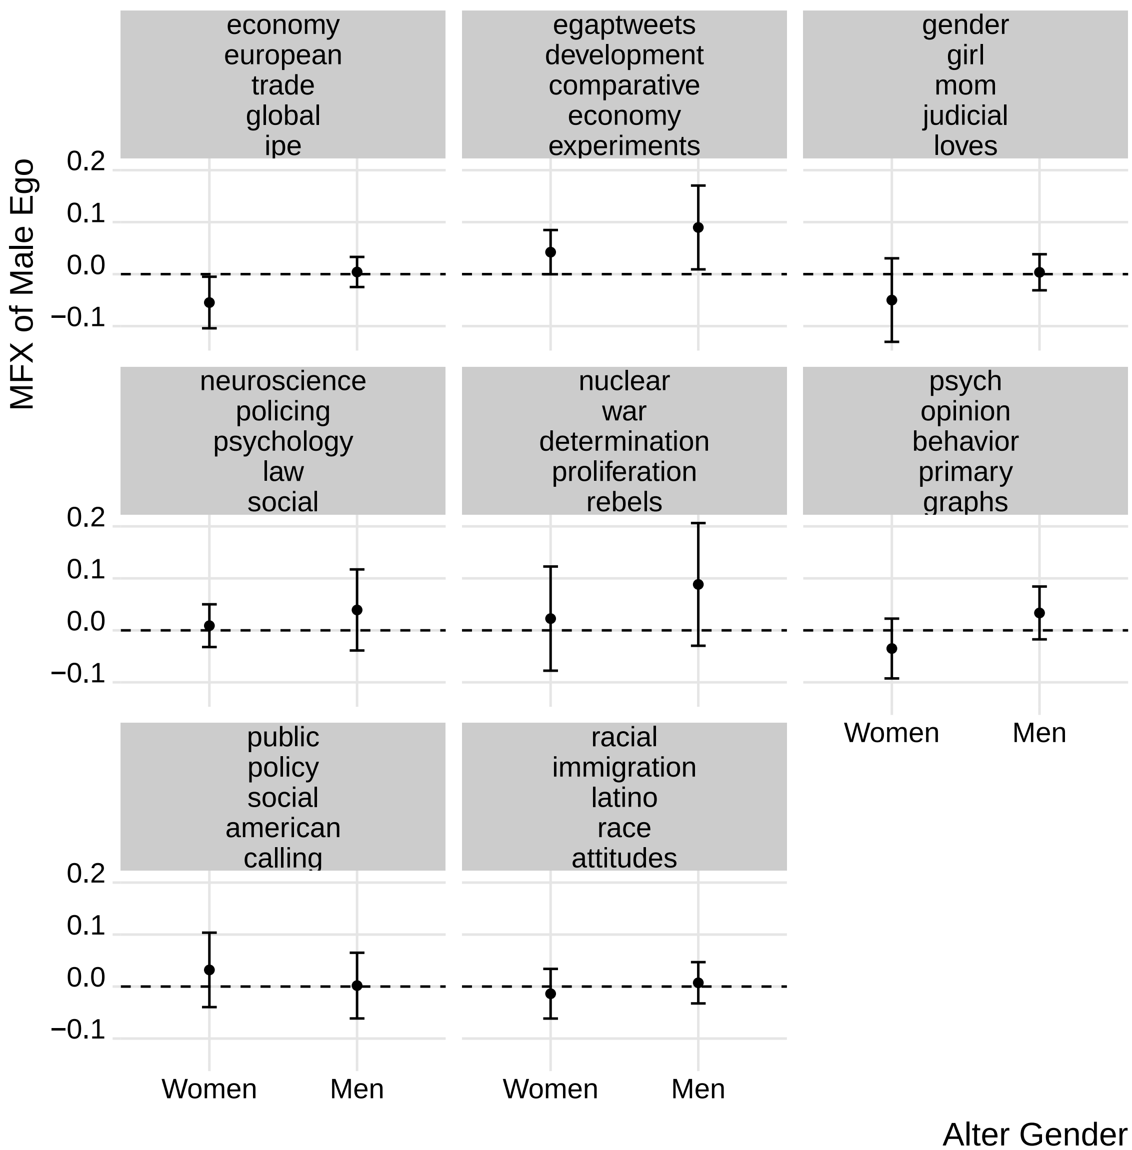

Figure 26: Marginal effect coefficients on male dummy for the ego on the gender of the alter across the 8 largest communities returned by the label propagation method. Behavior is probability of the ego engaging with research shared by the alter one or more times. Plot titles depict the top five TF-IDF words appearing in group member Twitter bios.

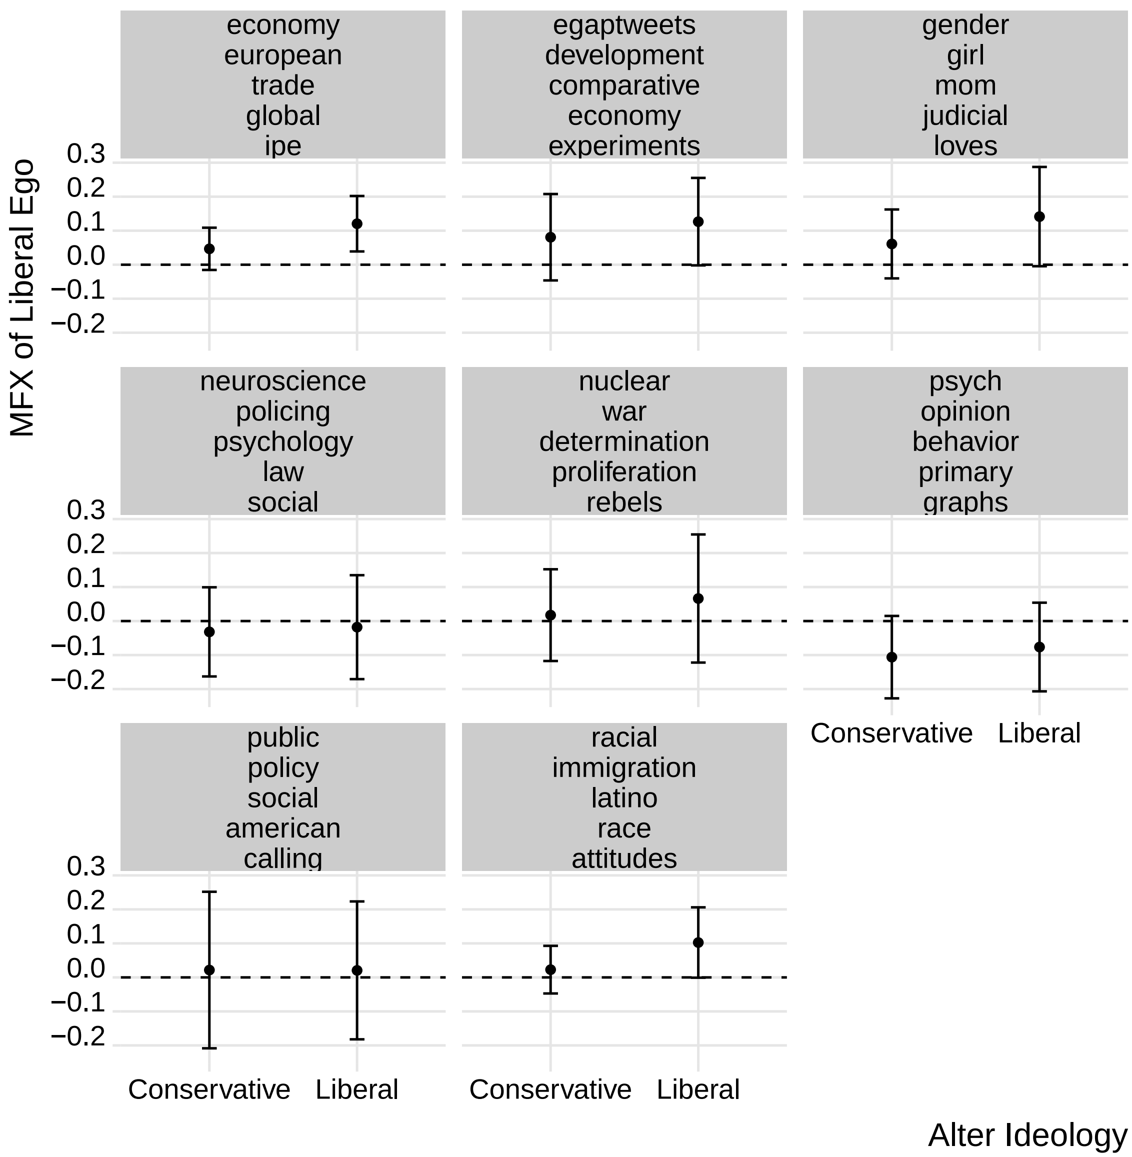

Figure 27: Marginal effect coefficients on male dummy for the ego on the gender of the alter across the 8 largest communities returned by the label propagation method. Behavior is probability of the ego engaging with research shared by the alter one or more times. Plot titles depict the top five TF-IDF words appearing in group member Twitter bios.

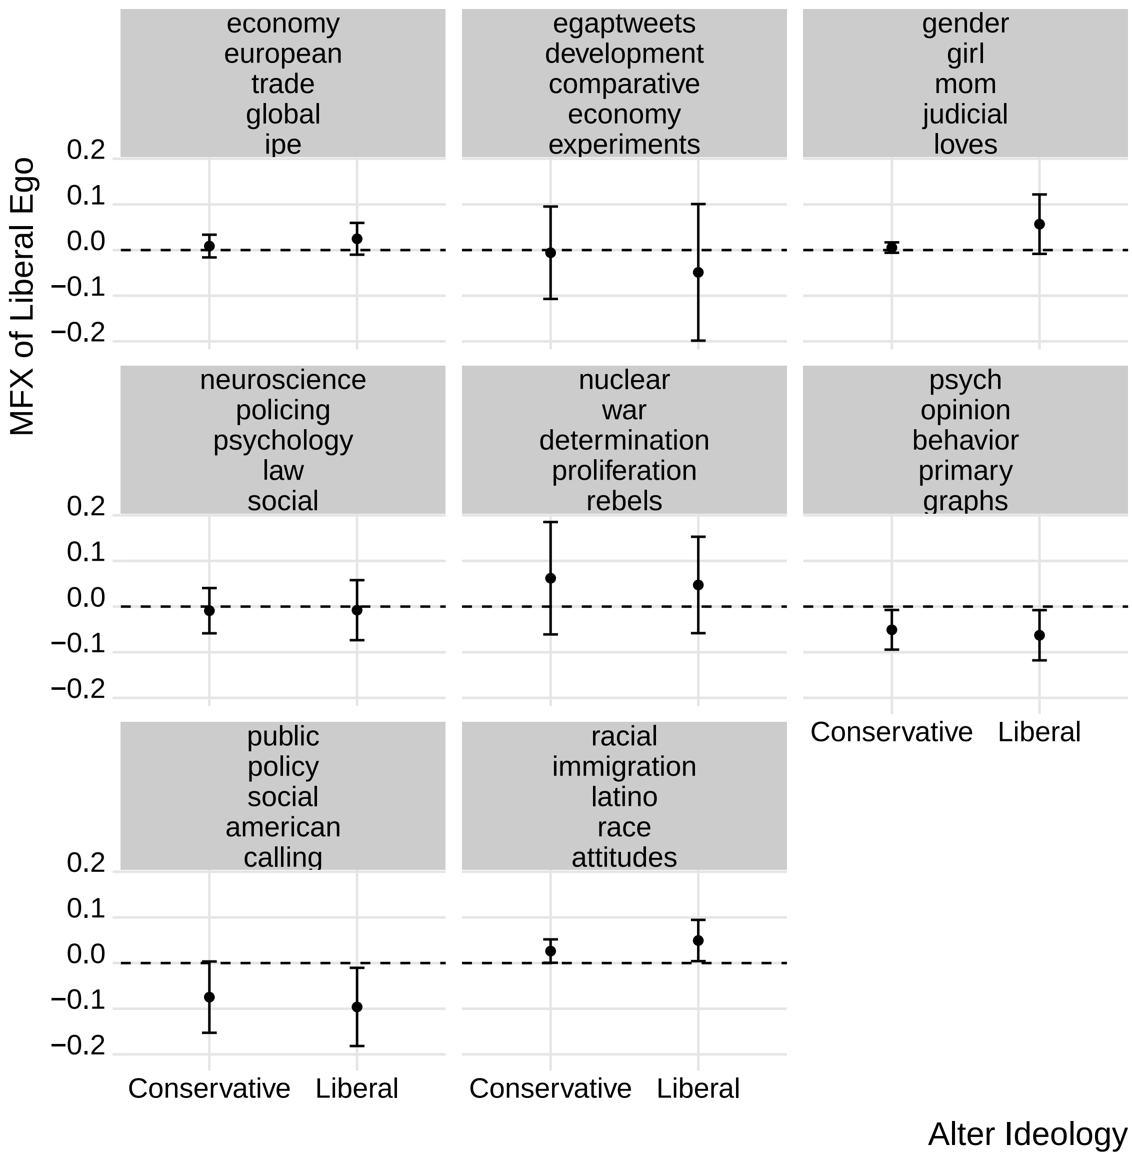

Figure 28: Marginal effect coefficients on male dummy for the ego on the gender of the alter across the 8 largest communities returned by the label propagation method. Behavior is probability of the ego engaging with research shared by the alter one or more times. Plot titles depict the top five TF-IDF words appearing in group member Twitter bios.

I ERGMs: Homophily Robustness

Our manuscript characterizes homophily using a variety of analytical tools, ranging from descriptive analyses to regression models. However, there are additional methods that are specific to network data which are designed specifically to evaluate homophily. We run an exponential random graph model (ERGM) on different networks to test whether covariates such as gender, position, and ideology predict a link in the network.

We plot the results of this test in Figure 29 which divides the results into link type (rows) and whether we look at a simple match (left column) or detailed matches (right column). In the simple match, we note that the volume of tweets is most prognostic of a tie existing in terms of followers (top-left plot) or mentions (bottom-left plot). Substantively, this means that two scholars are most likely to follow or mention one another the more each of them tweets. But when we disaggregate the node matches by specific groups (right column), we see that position and gender are now the most prognostic of a tie existing. Specifically, two scholars are most likely to follow or mention one another if they are both female or both tenure-track.

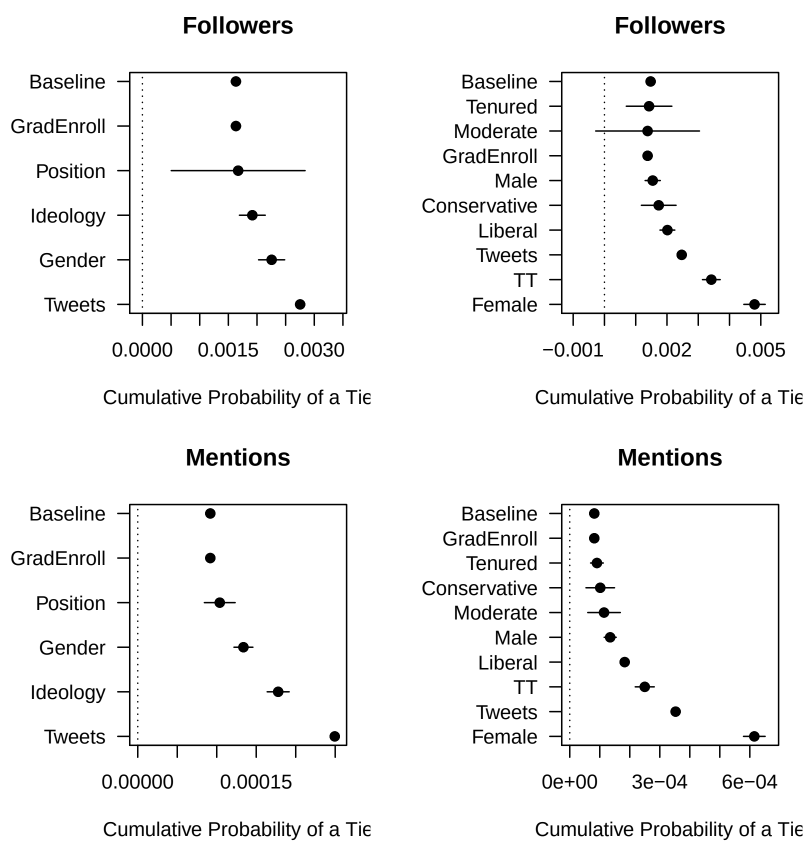

Figure 29: ERGM results predicting follows by individual-level covariates only (left panel) and controlling for mutual follows (right panel).

J Ideology

Our main analyses estimated user ideology using Barberá (2013) which applies Bayesian ideal point estimation to Twitter users' following choices. The underlying assumption for this method is that social networks are homophilic. Under this assumption, Barberá (2013) argues that an individual's choice about which political actors to follow reveals their latent ideology using standard Bayesian ideal point estimation techniques. This distribution is visualized in Figure 30, which suggests a bi-modal distribution that skews liberal.

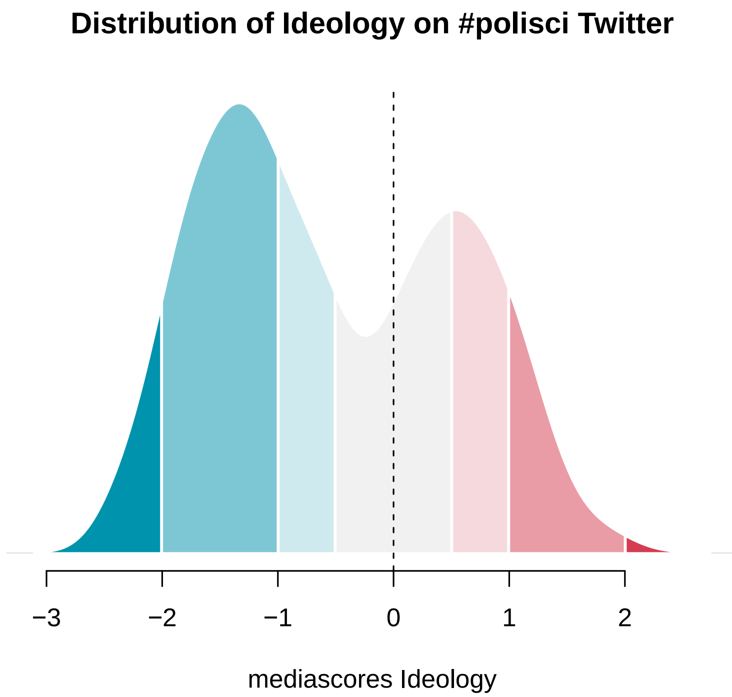

Figure 30: Density of ideology scores among 1,236 political scientists with active Twitter accounts. Ideology scores estimated according to Barberá (2015). Color shades represent substantive bins defined by the authors, ranging from very liberal (dark blue, ideology less than -2) to very conservative (dark red, ideology more than +2) with the range between -0.5 and +0.5 defined as moderate.

However, we recognize that this motivating assumption might not hold when looking at political scientists. Specifically, it is plausible that a political scientist's decision about who to follow is driven not only by their preference for ideologically similar content, but also is driven by their substantive area of academic research. For example, American politics scholars might follow President Donald Trump not because they feel ideologically proximate to him but rather because he generates news that is relevant to the scholars' research. Similarly, scholars of conflict, security, or terrorism might follow ideologically extreme accounts, despite not holding similarly extreme personal beliefs.

NYU's Social Media and Political Participation Lab (the SMaPP Lab) has an alternative method for estimating ideology which is determined not by who users follow but by the news media stories they share. The method – developed for `R` as the `mediascores` package – also relies on the assumption of homophily. Unlike Barberá (2013)'s approach (available in the `tweetscores` package for `R`), the `mediascores` package looks at the decision to share or retweet content from either liberal or conservative media outlets, such as the New York Times or Fox News. Unlike the choice over who to follow, decisions about which news stories to highlight are made on a daily basis and are more likely to reflect a mix of personal and professional interests. However, for scholars who are less active on Twitter, this method yields a lot more noise. We are unable to estimate ideology for 154 scholars due to the lack of sufficient data, motivating our decision to relegate this measure to the appendix.

We re-run our main analyses using this alternative measure of ideology. The overall distribution of ideology is shifted left, with the majority of our population consisting of liberals (61%) and moderates (36%), with only 3% of scholars being scored as conservative. We also note that the correlation between these two measures is very poor (-0.16), as illustrated in the scatterplot in Figure 32. Given these issues with estimating ideology for political scientists on Twitter, we caution against concrete conclusions using either measure and our main results using both.

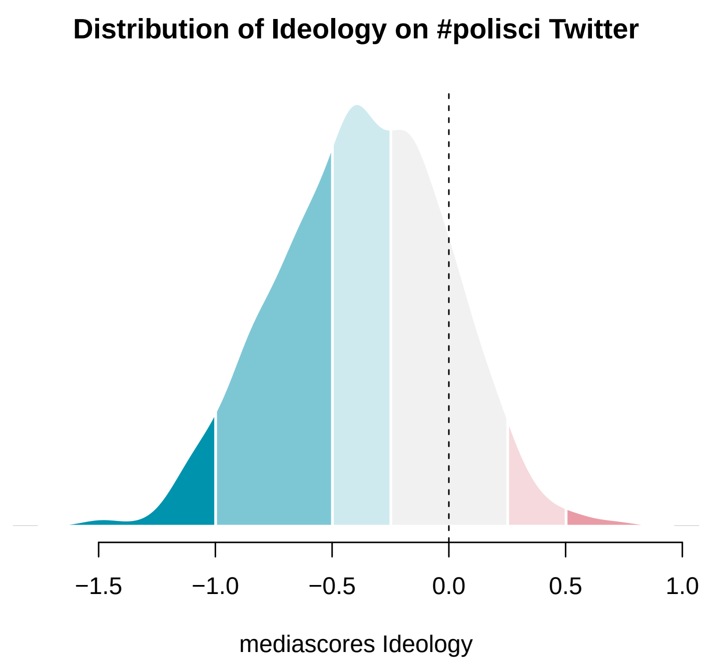

Figure 31: Distribution of ideologies generated by the `mediascores` package.

Replacing the tweet scores with the media scores yields much weaker face validity in our communities, as illustrated in Figure 33. Unlike in our manuscript, which showed

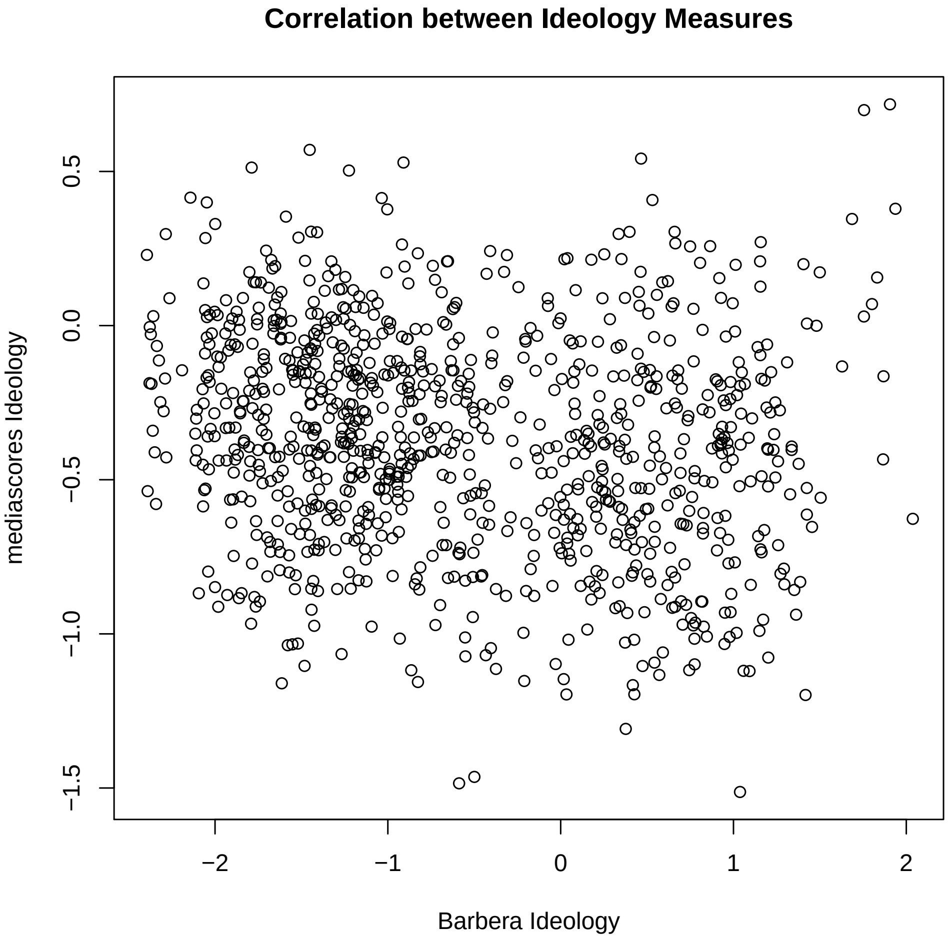

Figure 32: Scatter plot comparing Barberá (2013) ideology scores with those generated by the `mediascores` package from NYU's SMaPP Lab.

the greatest shares of liberals being members in the race and immigration community (#2) and the gender and judicial community (#4), we now see that liberals are significantly *less* likely to be part of these groups when using the `mediascores` package. Conversely, liberals constitute almost the entirety of the nuclear community (#7).

The biggest substantive impact the choice of ideology makes on our results is with our findings on homophily. With only a small fraction of our dataset comprised of conservatives, we restrict our attention to only the difference between liberals and non-liberals. But here we find much stronger evidence of homophily in terms of mentions (Figure 34), and even some evidence of a significant relationship with research dissemination (Figure 35).

We also note that these results persist even when we restrict our attention to comparisons of scholars who follow approximately the same accounts. As illustrated in Figure 36, liberals are far more likely to engage with other liberals on Twitter even

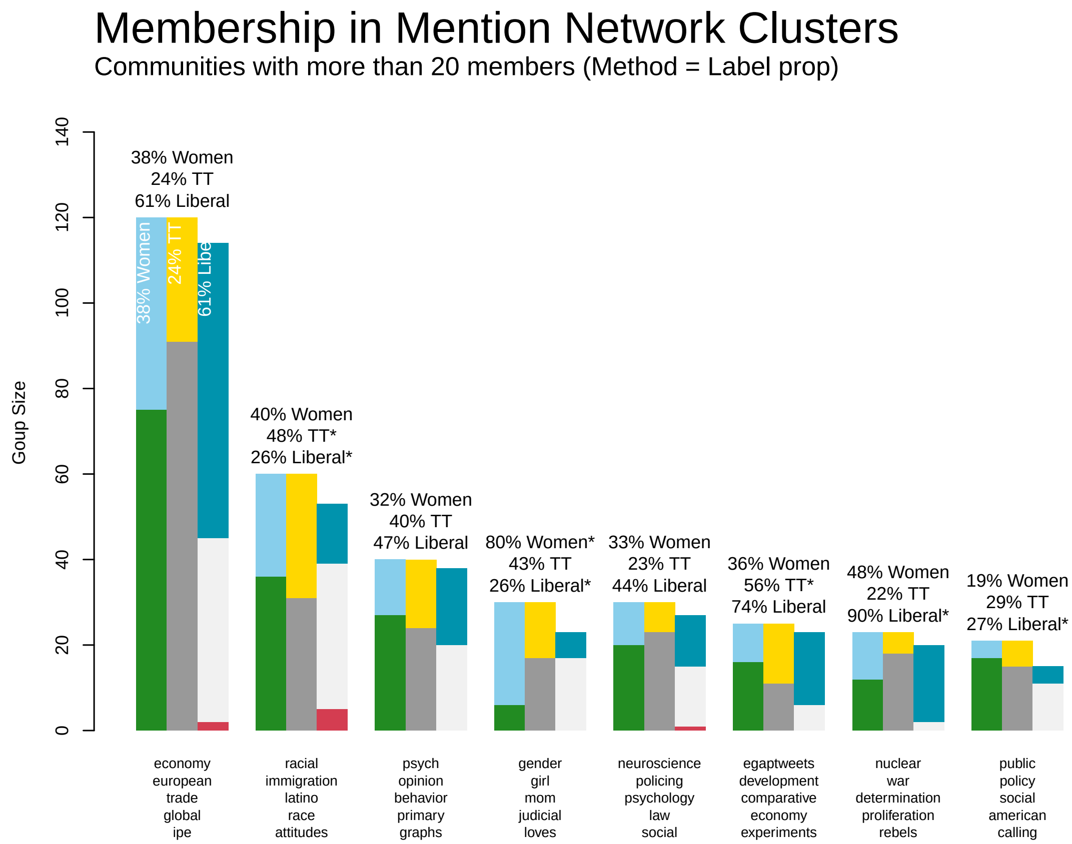

Figure 33: Communities by group. Ideology bars do not add up to the total due to 154 scholars for whom the `mediascores` package couldn't estimate ideologies.

when we compare them to non-liberals who follow similar accounts.

Finally, replicating the toxicity analysis suggests that non-liberals are mentioned in tweets that are more likely to be toxic on average (left panel of Figure 37) and in the extreme (right panel of Figure 37).

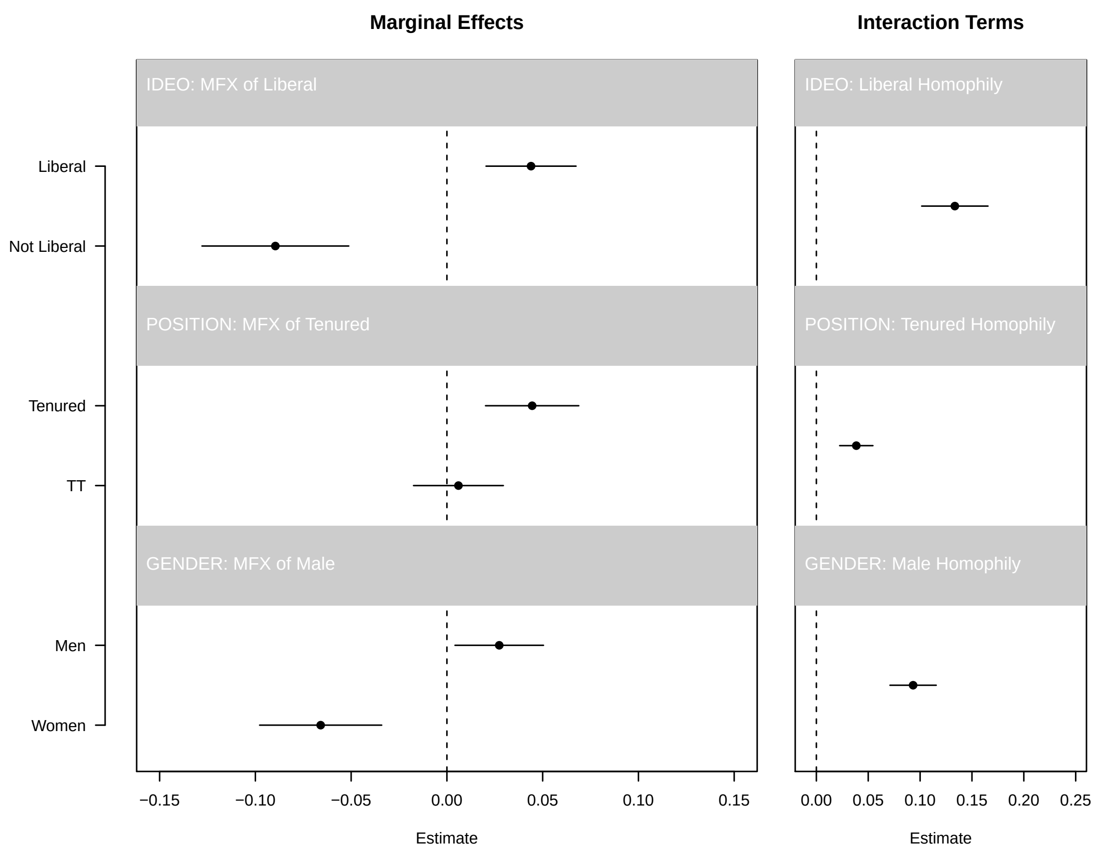

Figure 34: Dyadic results interacting ego and alter groups for links defined by mentions.

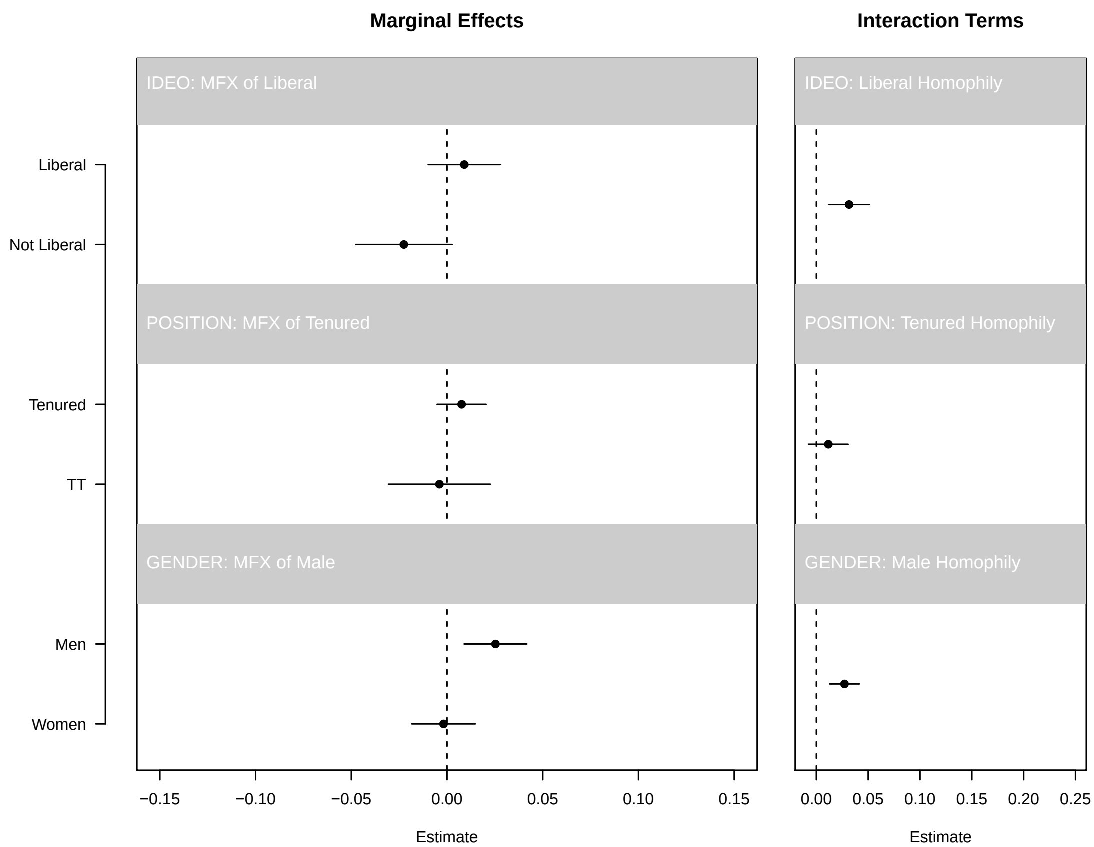

Figure 35: Dyadic results interacting ego and alter groups for links defined by sharing each others' research.

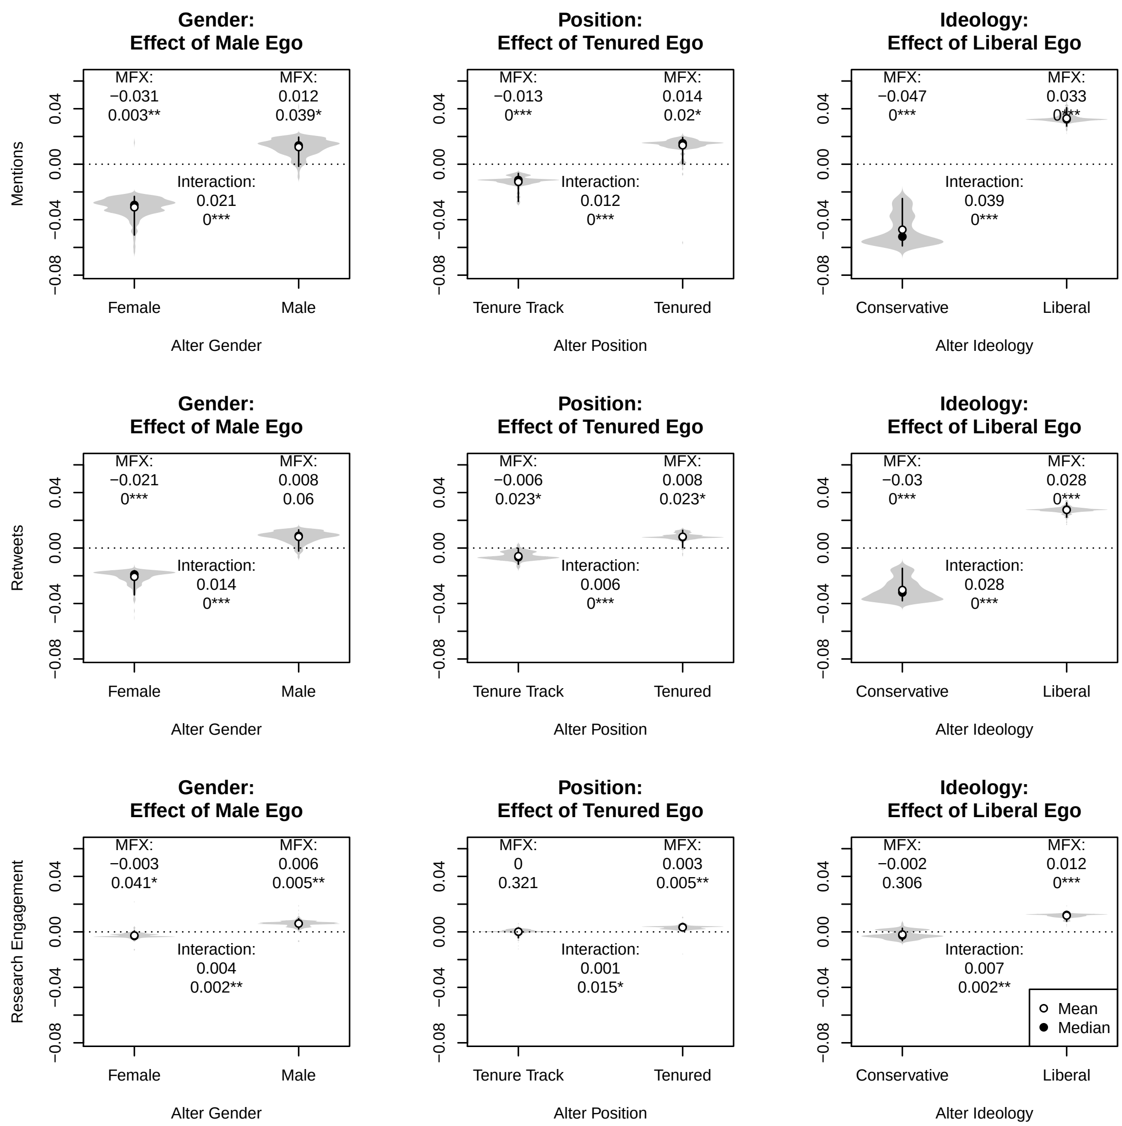

Figure 36: Confidence distributions for marginal effects (densities) and interaction coefficients (text) from bootstrap sampling the data and subsetting to all other scholars who follow similar accounts, where "similarity" is defined as the Euclidean distance between an ego's vector of friends and all other scholars. Top row of plots evaluates homophily in mentions, center row evaluates homophily in retweets, and bottom row evaluates homophily in research engagement.

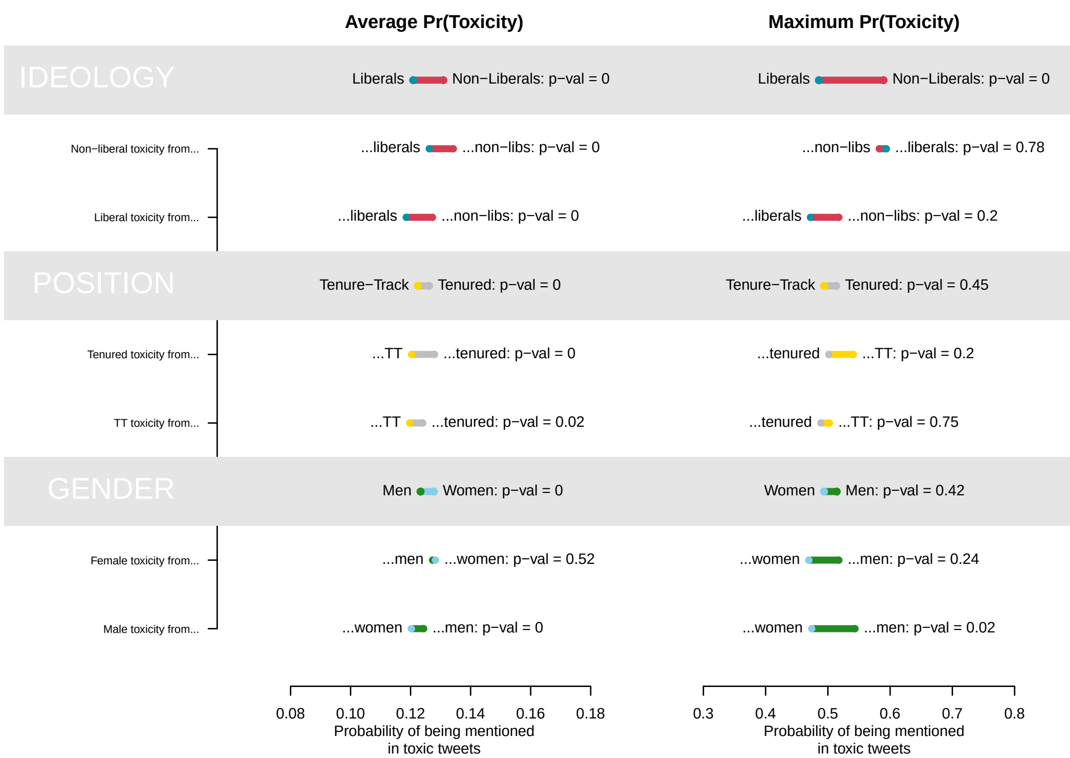

Figure 37: Average (left panel) and maximum (right panel) probability of being mentioned in a toxic tweet by group (gray headers) and disaggregated by who sends the tweet. P-values from t-test of difference of means given in text.

K Community Detection Robustness

Our manuscript used a label propagation method to identify mention clusters in our network. But there are several competing methods for identifying communities in social network data. We re-estimated communities using four of the fastest methods that are included in the `igraph` package for `R`, including:

- Leading Eigenvectors
- Infomap
- Fast Greedy
- Walktrap

The communities generated by these methods are given in the figures below. We highlight that, while the communities that are detected do change, there is consistent evidence of many of the same groups we find in our manuscript. Specifically, we highlight the persistence of a community involving gender and women, a community involving race and ethnic politics, and the broad categories of IPE on the one hand, and American politics on the other. The relative size of each of these communities varies depending on whether subsidiary communities are lumped together or disaggregated.

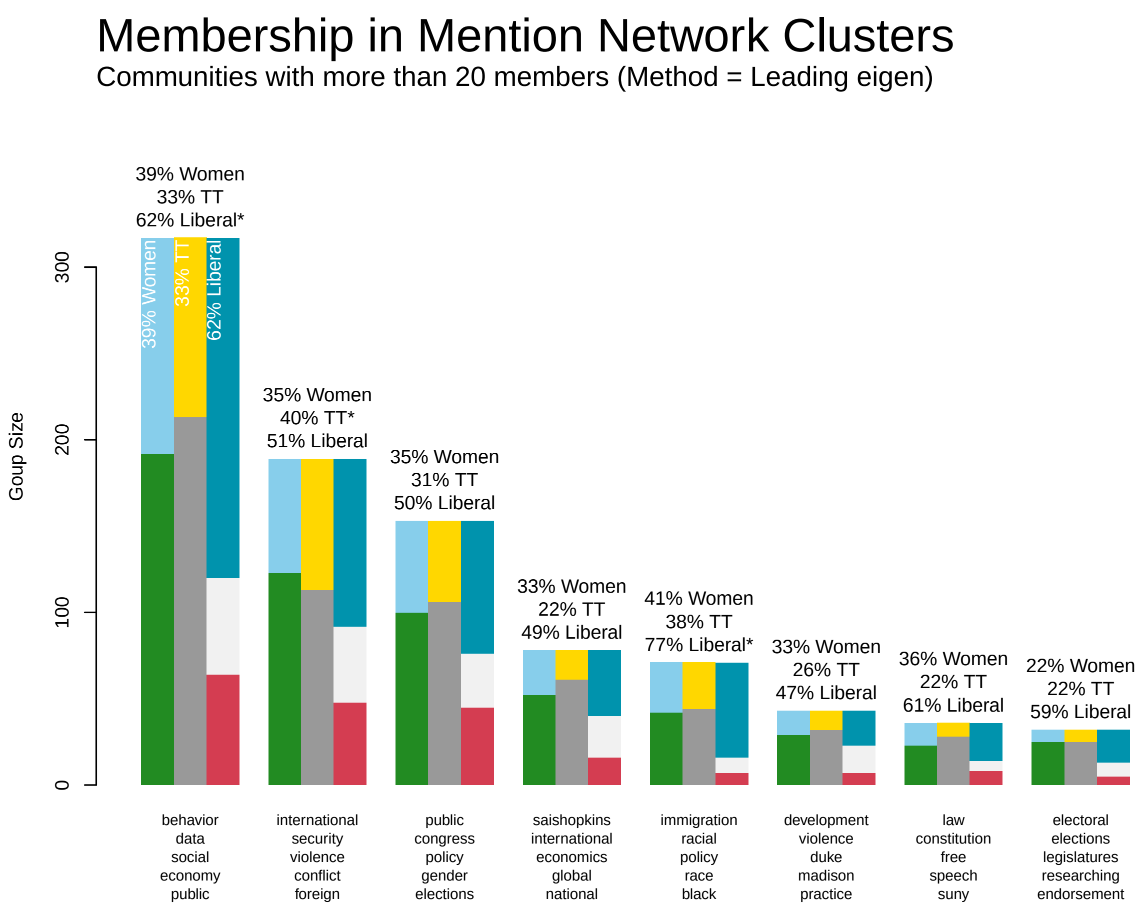

Figure 38: Communities by membership, calculated using the `cluster_leading_eigen` algorithm.

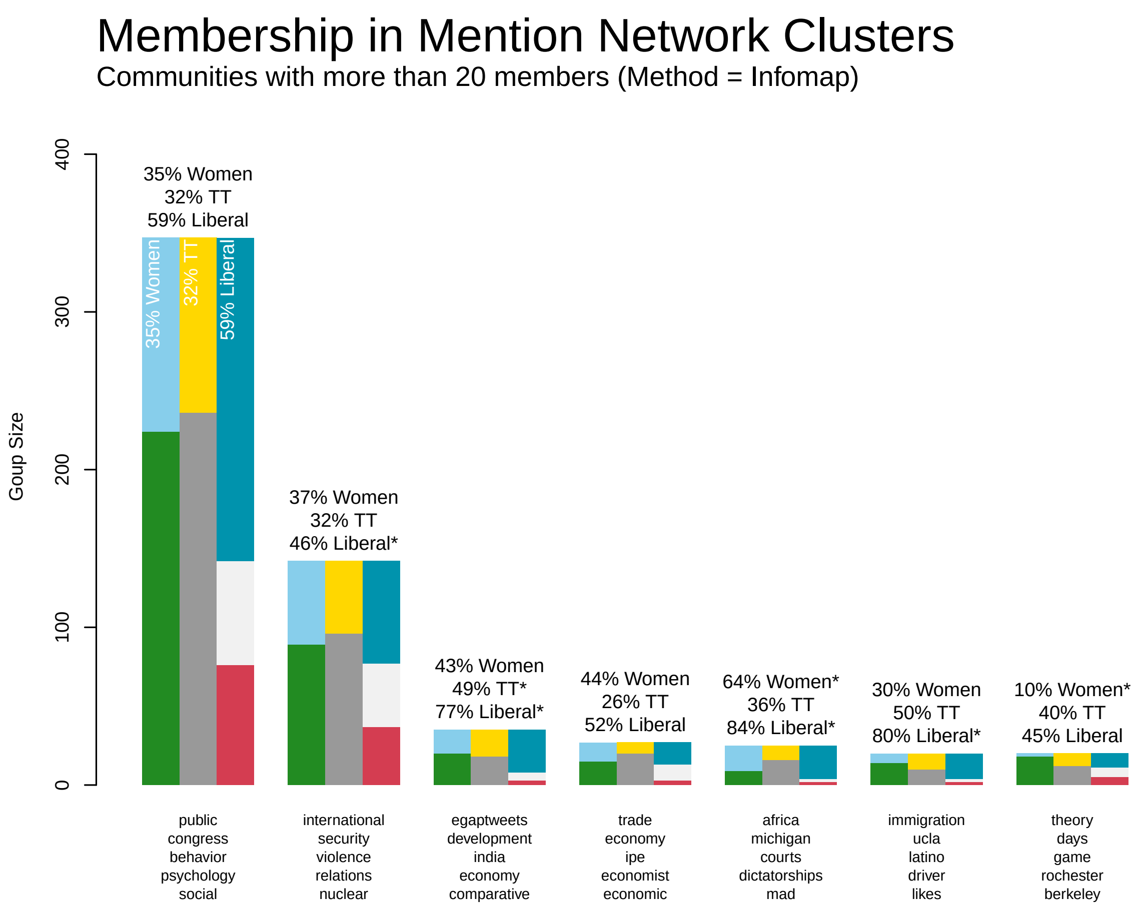

Figure 39: Communities by membership, calculated using the `cluster_infomap` algorithm.

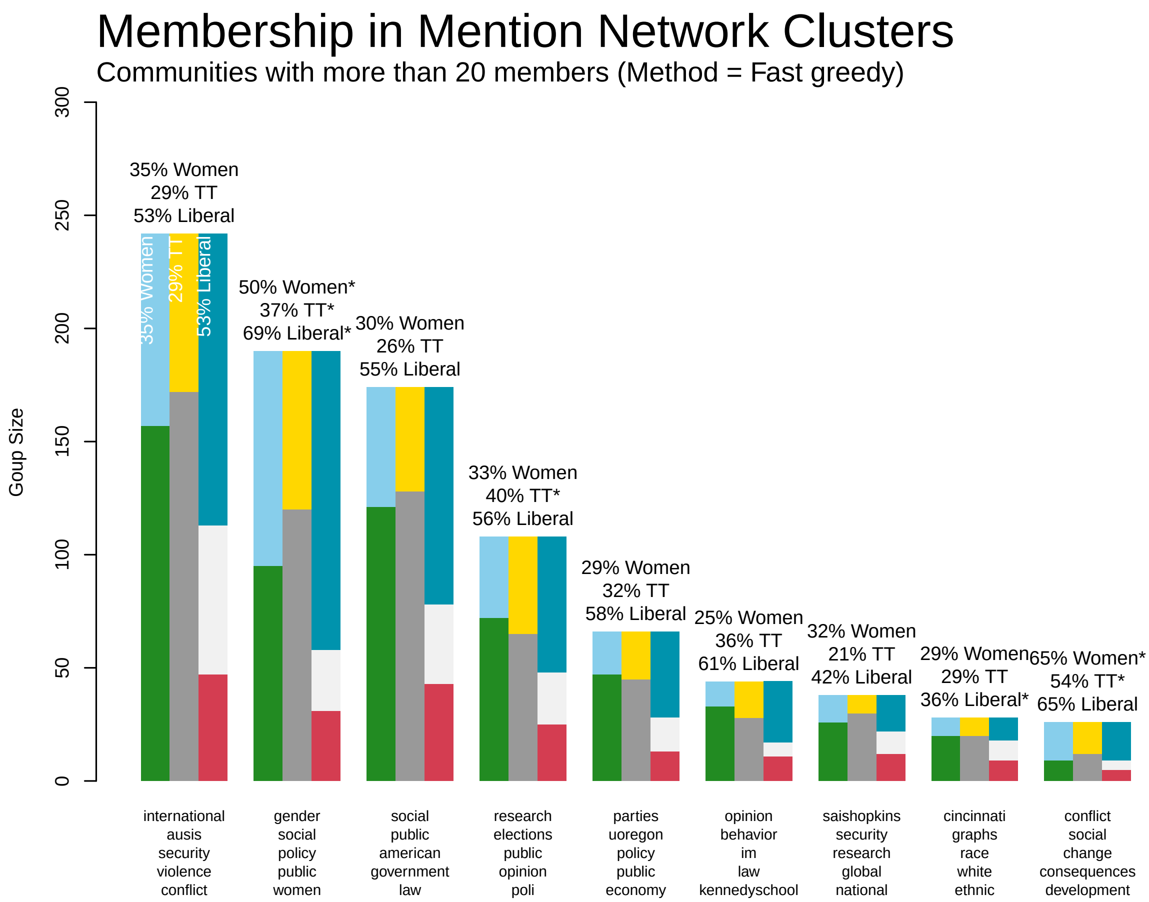

Figure 40: Communities by membership, calculated using the `cluster_fast_greedy` algorithm.

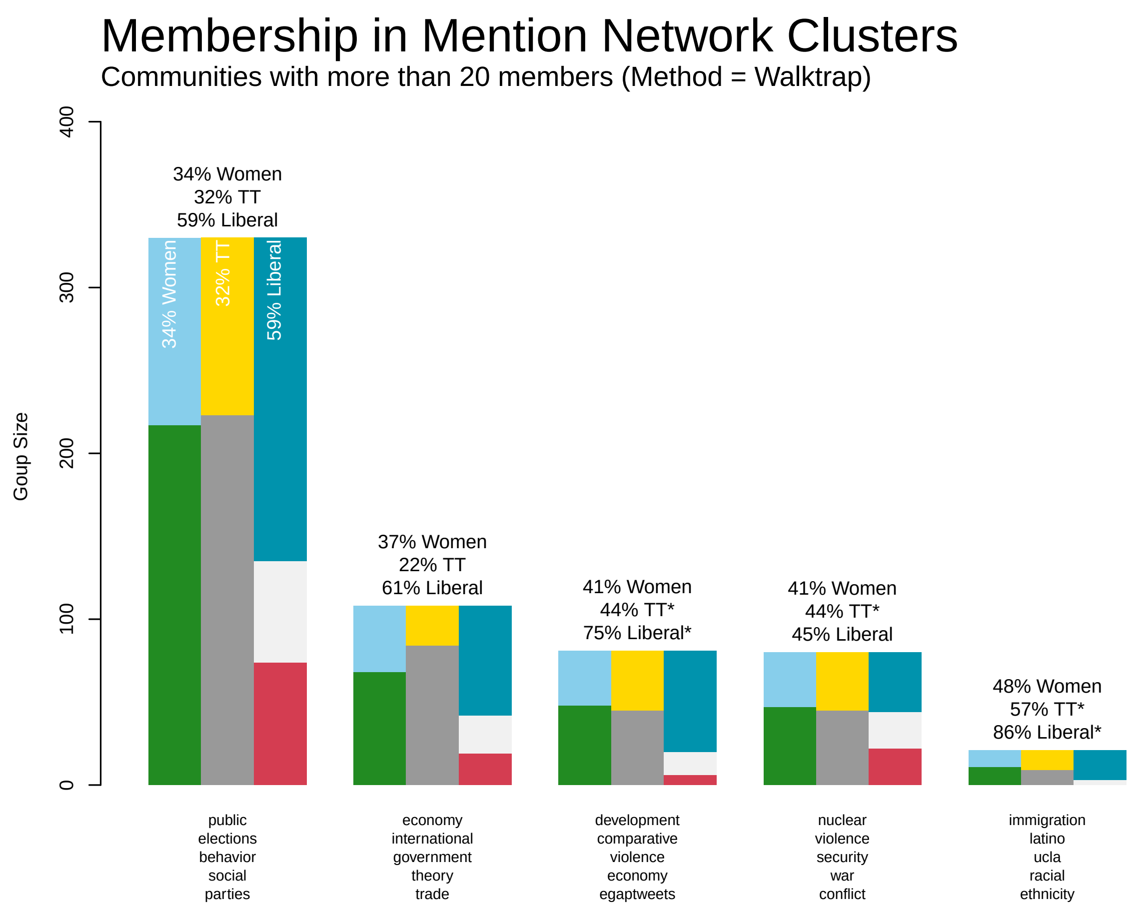

Figure 41: Communities by membership, calculated using the `cluster_walktrap` algorithm.

References

Aronow, Peter M, Cyrus Samii and Valentina A Assenova. 2015. "Cluster–robust variance estimation for dyadic data." *Political Analysis* 23(4):564–577.

Barberá, Pablo. 2013. "Birds of the Same Feather Tweet Together. Bayesian Ideal Point Estimation Using Twitter Data." *Proceedings of the Social Media and Political Participation, Florence, Italy* pp. 10–11.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together. Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23(1):76–91.

Bos, Angela L and Monica C Schneider. 2012. "New Research on Gender in Political Psychology: Mentoring to Fix the Leaky Pipeline." *PS: Political Science & Politics* 45(2):223–231.

Breuning, Marijke, Benjamin Isaak Gross, Ayal Feinberg, Melissa Martinez, Ramesh Sharma and John Ishiyama. 2018. "Clearing the Pipeline? Gender and the Review Process at the American Political Science Review." *PS: Political Science & Politics* 51(3):629–634.