Northwestern | INSTITUTE FOR POLICY RESEARCH

IPR Working Paper Series

WP-23-08

Liberals Engage With More Diverse Policy Topics and Toxic Content Than Do Conservatives on Social Media

Ho-Chun Herbert Chang
University of Southern California

James Druckman
Northwestern University and IPR

Emilio Ferrara
University of Southern California

Robb Willer
Stanford University

Version: February 14, 2023

DRAFT

Please do not quote or distribute without permission.

Abstract

The rise of social media provides citizens with direct access to information shared by political elites. Yet, more than ever before, citizens play a critical role in diffusing elite-generated content. What kinds of content spread on social media? Do conservative and liberal citizens differ in the elite content with which they engage? These questions relate to long-standing academic and popular debates about whether political behavior is symmetric or asymmetric with respect to political ideology. The researchers analyze more than 13 million users' retweets of messages by U.S. Members of Congress on Twitter from 2009 to 2019, leveraging estimates of users' political ideology constructed from over 3.5 billion prior retweets. They find limited evidence for ideological symmetry where the strength of ideology predicts diffusion choices similarly on the left and right. In contrast, the authors find robust support for ideological asymmetry. Consistent with theories of ideological asymmetry, liberals retweeted a more diverse set of policy topics than conservatives by 19.4%. They also engaged more with toxic content from in-group elites by 56%. Given the tendency for people to follow like-minded others on social media, these diffusion patterns imply that liberals are exposed to more politically diverse and toxic elite-generated content on social media, while conservatives receive more politically homogeneous and less toxic content. The demand and supply dynamics of political information suggest the existence of polarized information bubbles such that liberals and conservatives reside in distinct information ecosystems.

Significance Statement: Most Americans now receive at least some political information through social media platforms. However, social media users not only consume information, but they also play a critical role in determining its — potentially polarizing — diffusion. While existing research considers how messages spread, scant work explores the audience's perspective. Analyzing ten years of Twitter data with the labeling of users' ideologies (derived from the content of their prior social media activity), the researchers explore how liberals and conservatives engage with political information from political elites. They find evidence that liberals engage (measured via re-tweeting) with more diverse perspectives (i.e., a wider range of policy topics) than conservatives. They also exhibit greater receptivity to toxic content. Because people tend to follow other users with similar views to their own, these findings suggest that liberals may find themselves in more diverse, but also more toxic, information environments than their conservative counterparts on social media.

Social media plays a crucial role in shaping contemporary political information environments, both in the U.S. and globally. Audience members have an enormous amount of influence over what political information does or does not become prevalent via their decisions to diffuse social media information to others (Rathje et al. 2021). Do liberals and conservatives make these decisions in similar ways, or are there ideological asymmetries? Despite sizable literatures on the sharing of misinformation (e.g., (Chen et al. 2021; Osmundsen et al. 2021; Guay et al. 2022; Pennycook and Rand 2021) and message diffusion based on content (e.g., (Brady et al. 2017; Rathje et al. 2021; Frimer et al. 2022), scant work explores differential engagement (i.e., retweeting) with elites' messages based on the audience's ideology. This is a surprising gap given long-standing debates about whether liberals and conservatives behave in fundamentally different ways. Some scholars argue for pervasive asymmetries (e.g., (Barberá et al. 2015; Grossmann and Hopkins 2016; Jost 2017; Jost 2021; Baron and Jost 2019), while others assert symmetric behaviors by liberals and conservatives ((Ditto et al. 2019; Guay and Johnston 2022; Enders et al. 2022). We test these perspectives in the context of individuals receiving information from a highly influential category of elites in American politics — members of Congress (i.e., MCs)—and then deciding whether to diffuse it to other citizens.

Isolating whether asymmetries exist when it comes to engagement with elites' messages has crucial implications for understanding if ideologues experience distinct downstream information ecologies. Liberals generally obtain social media information from other liberals and conservatives from other conservatives (Boutyline and Willer 2017); thus, if ideologues who consume elite content pass along distinct types of messages (e.g., with different policy foci), it could alter the information bubbles in which those with different ideologies find themselves. Significantly different information bubbles would suggest a polarized situation. That is, insofar as polarization refers to "a large or increasing gap between two or more groups" (Jost et al. 2022), asymmetry in elite content engagement could lead to distinctive liberal and conservative information ecologies. This, in turn, could make cross-ideological exchanges and perspective-taking difficult as the sides have different information realities.

Turning to potential sources of asymmetries, we have reason to expect ideologically divergent responses with regard to two information attributes. The first concerns engagement with political issues or policies. Ideologues and partisans tend to care more about some issues compared to others. The theory of issue agendas and, relatedly, issue ownership suggests that parties and candidates prioritize (i.e., pay attention to) some issue bundles more than others (Baumgartner et al. 2006, 2019), often those on which they are advantaged (Petrocik 1996; Druckman and Jacobs 2015). For instance, Democrats prioritize the environment, education, and civil rights, while Republicans put more weight on defense, the economy, and international affairs.

We expect ideologues to react differently to partisan policy agendas. To see why, consider that conservatives tend to prefer situations of reduced uncertainty, threat, and discord; this is exemplified by their lower openness to new experiences (Jost 2021). This leads them to be averse to engaging in new ideas. For example, Rogers and Jost (2022) show that conservatives are less likely than liberals to participate in various cultural activities (e.g., consuming music, adopting new hobbies, attending live performances) due partially to lower openness. It follows that conservatives will be less likely to engage with (i.e., retweet) policies that come from the other party's (i.e., liberals' / Democrats') agenda. Consequently, relative to liberals, when conservatives encounter issues that are not part of their existing agenda in a tweet from an MC, they will be less likely to retweet it (Hypothesis 1, also see Barberá et al. 2015: 1537). Conservatives have less of an interest in engaging with issues with which they are not already familiar.

From the perspective of MCs, Hypothesis 1 suggests that Democratic MCs elicit a less diverse audience than Republican MCs. To see why, consider that Democratic MCs will focus on liberal issues; in so doing, they will not attract a conservative audience, which would be averse to their agenda (and, thus, their audience will tend toward mostly liberals). In contrast, Republican MCs who mostly emphasize conservative issues will draw both a conservative audience and a liberal audience that is more open to variation in agendas (and, thus, their audience will be more diverse since it will include both conservatives and liberals). Thus, another way to consider Hypothesis 1 is that Democratic MCs will have less diverse audiences than Republican MCs. We will test Hypothesis 1 in this latter manner.

The second type of information to which we expect liberals and conservatives to respond differently is toxic or uncivil content. Conservatives have a relatively greater need for order and place greater value on tradition than do liberals (Jost 2017). Toxicity violates social norms (Bormann et al. 2022) and thus counters convention, order, and tradition. Along these lines, Mutz (2015: 106) states, "Republicans and conservatives may be particularly sensitive to norm-violating threats to the social structure." Mutz shows that negative emotional arousal is especially high among Republicans (and Independents) when exposed to toxic incivility (also see Druckman et al. 2019). This suggests they may be less inclined to share toxic communications given the negative emotions associated and the desire to not violate norms (and maintain order). It follows that, relative to liberals, conservatives will be less likely to retweet issues that contain toxicity (Hypothesis 2, also see (Theocharis et al. 2020; Frimer et al. 2022). Prior work makes clear that toxicity (or animosity) is invoked when discussing the other party (and this triggers virality) (Rathje et al. 2021). Consequently, conservatives will be less likely to retweet Republican MCs' tweets about Democrats than liberals to retweet Democratic MCs' tweets about Republicans (since liberals are not as bothered by toxicity common in tweets about the other party).[1] The result is less cross-cutting exposure and a distinct information environment. We thus expect that ideologues will show different responses to the political information they receive from elites.

A further implication is that we have no reason to expect variation based on ideological extremity per se, as it is ideological predilection rather than placement that we expect to matter. This is the crux of the asymmetry versus symmetry accounts – the former, as explained, predict differential response based on policy topics and toxicity while a version of the latter would suggest similar ideological responses but with extremists on both sides generating more variation, potentially being more toxic (e.g., (Schraufnagel et al. 2021; Ballard et al. 2022). Understanding whether there is ideological asymmetry in retweeting behaviors has substantial

[1] Rathje et al. (2021) report no difference by the ideology of members of Congress or media outlets; however, they focus on differences in sharing out-group messages (or in-group messages) against messages that do not invoke either group at the supply-level. Our focus is on audience response (demand-level) analysis for understanding engagement with the other side.

implications for isolating the information environments to which those with different ideologies are exposed. Given the general tendency for ideologues to obtain information from like-minded ideologues on social media (Boutyline and Willer 2017), ideological asymmetry – if correct – would suggest a potentially widespread disparity in information environments. Liberals would be exposed to more policy topics and more toxicity than would conservatives and, thus, would have distinct bases on which to make decisions and interact.

Data

We use Frimer et al. (2022) as the seed dataset (N=1,293,753), which contains **toxicity** (or incivility) scores based on the *Google Perspective API*. We additionally add the following variables: **policy topic** (machine-learned labels with an accuracy of 88% based on the Comparative Agenda Project) and **user ideology** based on user timeline data (weighted score based on media domains tweeted). The dataset contains a total of 10,451,080 retweet events, 1,435,132 unique tweeters, and 4.04 billion timeline tweets. Our dataset differs from Rathje et al. (2021) primarily in timespan, including MC tweets from between 2009 and 2019. Notably, this then includes early online discussion of the 2016 primaries (which does not seem to be as present in Rathje et al. (2021), whose data are from 2016 to 2020), and contains the subsequent presidential election and COVID-19 discourse. As mentioned, we label audience members as liberals (Democrats) or conservatives (Republicans) based on their media diet over time. While we suggest this captures ideology (partisanship), an alternative (and perhaps less presumptive) interpretation is to think of it as a political diet.

Using Twitter: There are many social media platforms, each with its own structure. Which platform a researcher uses for a given study depends on a variety of considerations, including the function of the data, the availability of the data, and the hypotheses under study (Guess 2021). Our interest lies in how individuals with varying ideologies react to data from elite sources. This requires that we identify the political predilections of both elites (i.e., members of Congress) and Twitter users. We rely on partisanship for the former since it is easily identifiable. Users' partisan affiliations on Twitter, in contrast, are not explicitly declared but can be accurately inferred by relying on patterns of Twitter use over time to categorize users as either liberal or conservative (Ferrara et al. 2020; Chang et al. 2021; Chen et al. 2021). In this regard, Twitter is unique among widely used social media platforms in that it allows public and retroactive analysis detailed at the user level. Our scope is 10 years, which falls outside the range of Facebook's data retention. More generally, Facebook user timelines are much more privacy-restrictive, which poses challenges to user-level ideological modeling (Sharma and Verma 2018). As Steinert-Threlkeld (2018: 6) states, "researchers interested in diffusion processes and emergent behavior find Twitter a natural resource." We recognize that, like other studies of a single platform, there are invariably questions about generalizability. Those on Twitter are far from a representative sample of Americans. We view our study in a Popperian sense – if the data cohere with our predictions, that means the continued accumulation of evidence for theories of partisan asymmetry.

Consequently, we employ partisanship and ideology as proxies for one another, assuming that Democratic (Republican) MCs share an outlook and identity with liberal (conservative) users. We will use the terms interchangeably, noting the high degree of ideological sorting among partisans (e.g., Levendusky 2009).

Results

We begin with a crucial descriptive figure. **Figure 1** shows the ideology of those who retweet MCs' tweets. The x-axis provides the political position of the retweeter with lower (higher) scores being more liberal (conservative). That is, negative values (between -2 and 0) indicate a liberal/Democrat-leaning user, while positive values (between 0 and 2) indicate a conservative/Republican-leaning user. The y-axis provides the probability density estimate, with the solid blue line indicating Democratic MCs and the dashed red line Republican MCs. The figure reveals that for Democratic MCs, the retweeters are much more likely to be homogeneously liberal; in contrast, for Republican MCs, the retweeters are much more likely to be widely distributed, including a non-trivial proportion of liberals (i.e., only a slight skew to the right). This implies liberals are more likely to tune in to Republican MCs than conservatives to Democratic MCs (also see (Barberá et al. 2015; Wojcieszak et al. 2022): an asymmetry in engagement with other-party MCs. We next turn to uncover whether our predictions capture (can explain) the sources of the asymmetry.

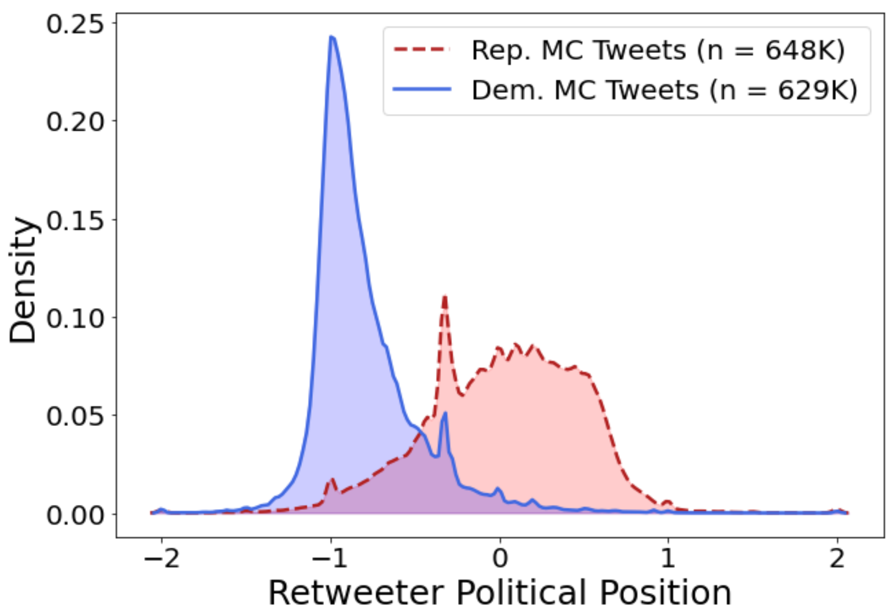

Figure 1: Density estimate of audience response by Republican MCs (red) and Democratic MCs (blue). Republican MCs garner a more diverse audience – more liberal users tune in.

Policy Engagement

Our first hypothesis suggests that conservatives will be less likely to retweet Democratic MCs than liberals to retweet Republican MCs. This occurs because conservatives avoid engaging with the liberal policies that Democratic MCs likely discuss more so than liberals avoid conservative topics. To identify issues, we trained a supervised machine learning classifier using labeled tweets from Russell (2018), with the assumption that the linguistic characteristics of tweets between Representatives and Senators are similar. The dataset includes a total of 68,398 tweets, 45,402 of which labeled with codes from the Comparative Agenda Project's (CAP's) codebook (Bevan 2017) and 22,996 labeled as non-related tweets. We implemented a variant of the large BERTweet architecture, a pre-trained model from Nguyen et al. (2020), and trained two classifiers – one for multiclass classification of 20 topics and one for binary classification to see if a tweet relates to policy or not, distinguishing if policies are mentioned. Our first (multiclass) model produces an accuracy of 86.4% for classification across all 20 policy topics (from CAP). Our second (binary) model produces a score of 88.5%. This beats the current best model at 79% (Hemphill and Schöpke-Gonzalez 2020). Methodological details are further discussed in the Methods section.

We then classified issues as liberal/Democratic or conservative/Republican based on their relative supply, with the theoretical rationale being that partisans emphasize issues on which they are advantaged (e.g., Druckman and Jacobs 2015). For **Figure 2**, the x-axis shows the difference in tweet supply for each policy (Republicans minus Democrats). This operationalizes right and left agenda items. Positive values indicate Republican MCs tweet more; negative values indicate Democratic MCs tweet more. The y-axis shows the political diversity via the standard deviation of political position of the audience – that is, the political diversity for the given topic (and thus the results are *not* affected by the frequency with which a given topic is tweeted by MCs). Thus, the blue dots are aggregate tweets from Democratic MCs and the red dots from Republican MCs. The figure is consistent with work on issue ownership: Democratic MCs tweet Democratic topics, including civil rights, health care, environment, immigration, and education, while Republican MCs tweet Republican topics such as defense, macroeconomics, and international affairs (e.g., Petrocik et al. 2003, Fagan 2019, Craig and Cossette 2020). To further ensure the validity of the variance, we also plot the average political position (Figure S2 in the SI) and show that the direction of diversity does indeed come from the out-group. This is a useful confirmation since our definition of liberal and conservative issues is based on the supply difference (see Figure S1 in the SI); it also shows that MCs of a given party pursue the strategy we suggested by emphasizing their owned issues (the one exception issue is Democratic MCs' relative emphasis on law and crime despite it being a traditional Republican issue. This likely reflects Democrats using the terms "criminal" or "crime" to discuss Trump during the Russia probe, which would inflate the supply of this topic for Democrats).

The figure clearly supports Hypothesis 1; when Republican MCs tweet, the audience diversity is always very high (consistently around 45%, with a slight downslope). In contrast, when Democratic MCs tweet on issues usually owned by liberals (i.e., civil rights), diversity is low, meaning cross-cutting engagement is low. Moreover, when Democratic MCs tweet about issues usually owned by conservatives (i.e., defense), diversity is high, which suggests cross-cutting engagement is high. Liberals respond consistently to Republican MCs across policy topics, whereas conservatives are more selective and only engage with Democratic MCs for conservative agenda items. Conservatives do not appear to engage with Democratic MCs on liberal issues. However, liberals engage with Republican MCs on conservative issues.

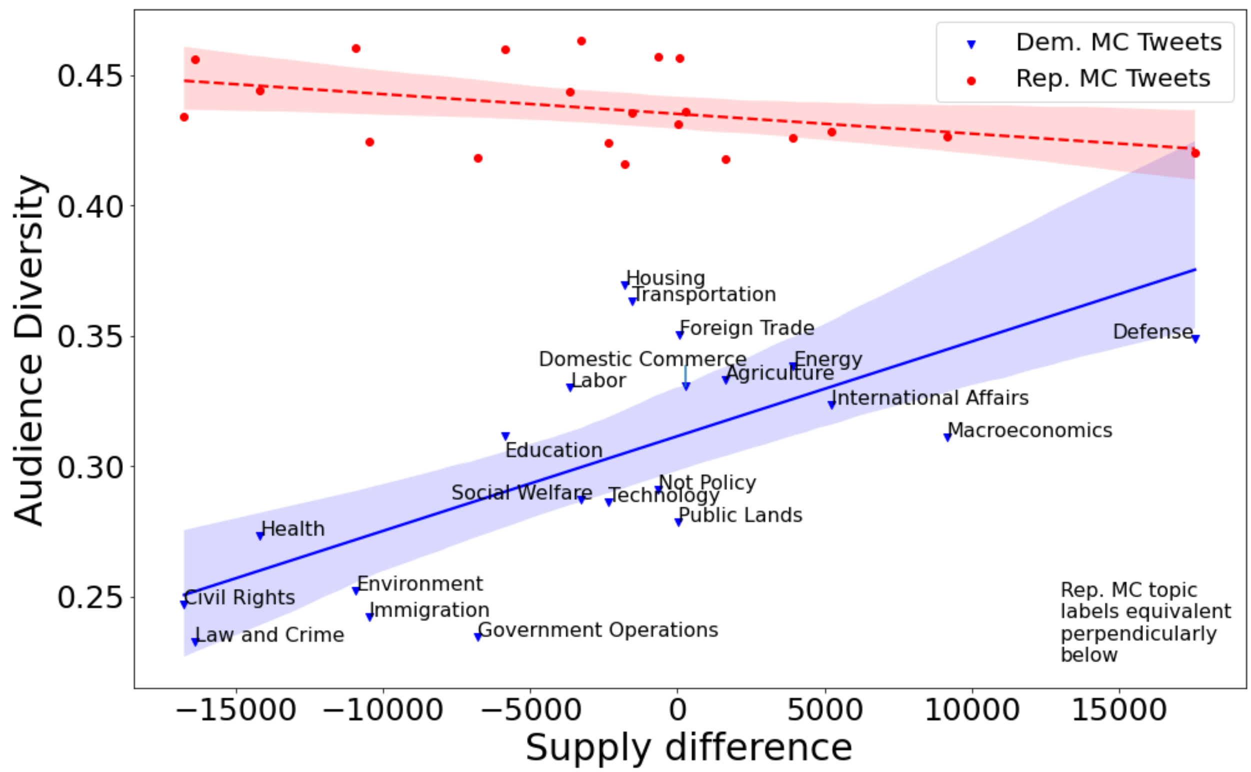

Figure 2: Left-/right-heaving agenda items versus audience diversity. Negative values indicate Democratic MCs tweet more and positive values Republican MCs tweet more. Republican topic labels are equivalent to Democrat labels perpendicularly below. Republican MCs elicit greater response from liberal users than do Democratic MCs from conservative users. When Democratic MCs tweet about more 'conservative' issues (e.g., defense, macroeconomics), diversity is high. When Democratic MCs tweet about more 'liberal' issues (e.g., civil rights, the environment), diversity is low. Republican MCs have high audience diversity across all topics, which indicates the presence of Democrat retweets. Thus, Democrat cross-cutting engagement is overall higher.

We further summarize these results in **Figure 3**. **Figure 3a** shows the user distribution of unique topics. The x-axis represents the number of unique topics with which a user has engaged, and the y-axis the proportion. The heavier tail of liberals show that, in general, liberals engage more diversely. This, however, does not take into account the skew within the topic categories: for instance, a user can engage with two topics, but fifty times with topic 1 and just once with topic 2. For every user, we then attach a score for topic diversity based on Shannon entropy:

$$H(x) = -\sum p(x) \log p(x)$$

Here, *p(x)* denotes the probability of an event (in our case, the frequency), and *log p(x)*, the information content. Thus, a higher Shannon entropy in this context means greater diversity. **Figure 3b** shows the cumulative distribution; the more leftward the S-shaped curve is, the more aggregate diversity. This shows, once again, that liberals engage more diversely. The extent of the partisan difference is sizable, captured by the yellow shading in **Figure 3a**: when aggregated, it is equivalent to a 19.4% divergence, where liberals engage more diversely than conservatives. Or, put another way, liberals retweeted 19.4% more diverse topics than did conservatives (a ratio of 1.19). We can further show that this scales with ideological extremism. **Figure 3c** shows that as we scale the absolute value of ideology. topic entropy decreases using LOESS smoothing, and this happens faster for conservatives. This affirms a prior hypothesis from Boutyline and Willer

[(2017)](#) that homogeneity increases with extremism, and conservatives exhibit this more acutely (suggesting that symmetry does not exist among extremists on each side). The compounding of homophilous ties and lower topic diversity for these intermediaries would have downstream effects for who generates virality, especially with toxicity. One further extension is the relationship between bots and media diversity, as bots may share a select few media sources by design. As they share in higher volumes, they may be over-represented at certain extremity bins, and due to their biases toward certain topics, skew topic diversity scores (Chang and Ferrara 2022).

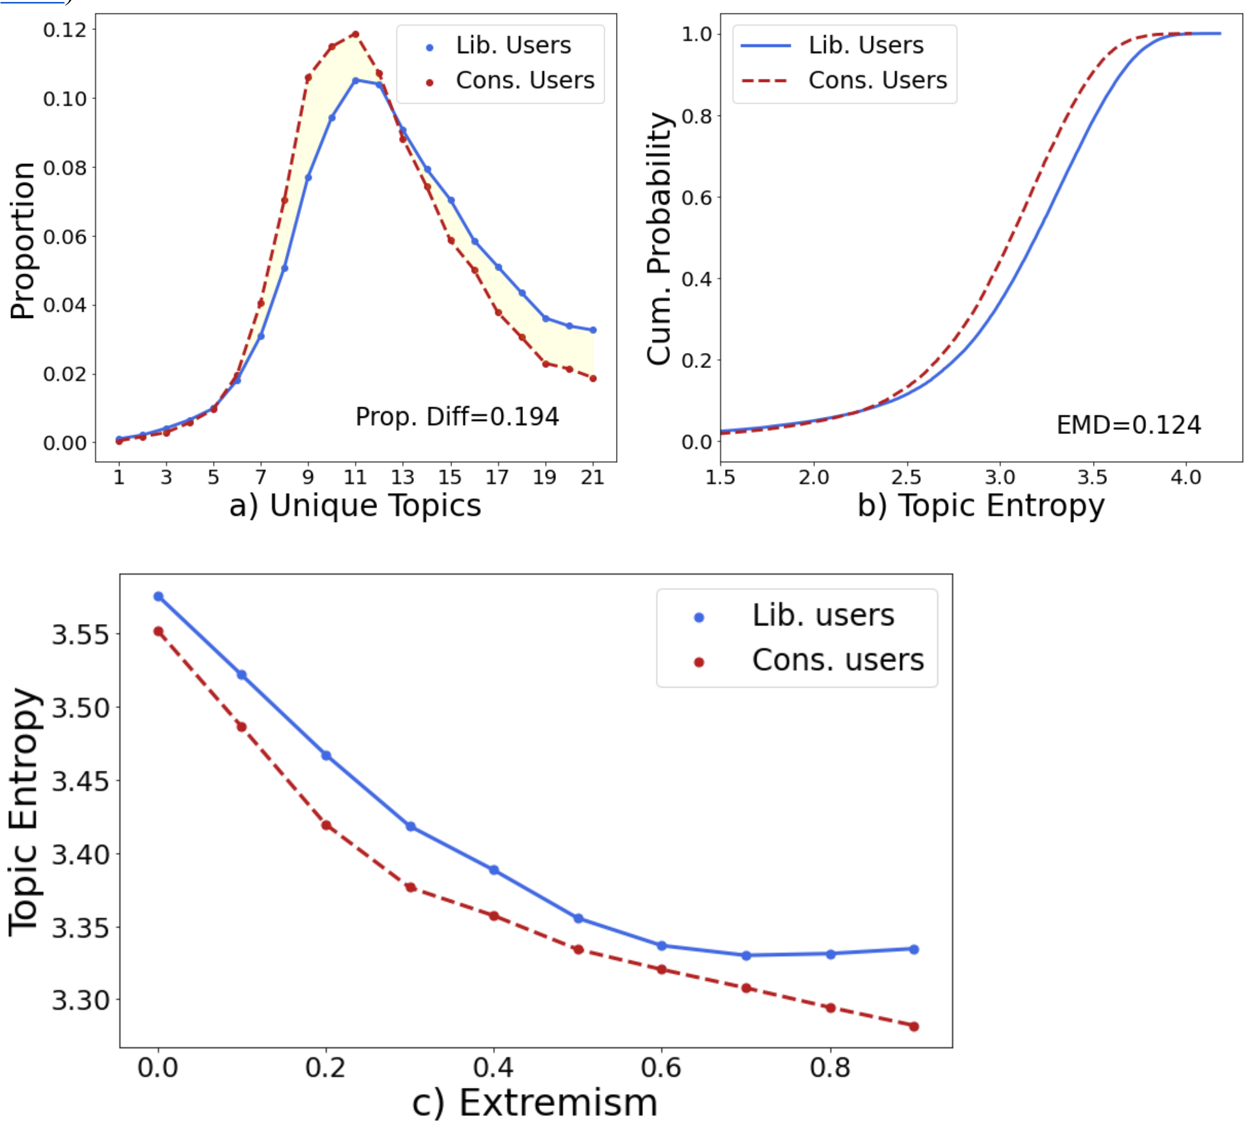

Figure 3: Liberals are exposed to and engage with more policy topics than do conservatives, shown through (a) unique policy topics retweeted and (b) aggregate measures of entropy (the more left the sigmoid curve, the more diversity); (c) scales topic entropy against ideological extremism with LOESS, and we observe that topic entropy decreases with extremism until a pullback at higher levels of extremism.

Toxicity

We next turn to tweets that invoke the parties regardless of issues. **Figure 4a** displays the relationship between the toxicity of an MC's tweet about their or the other party and the extent to which the tweet is retweeted. The figure shows, consistent with prior work, that in every case, as toxicity increases, so does retweeting (e.g., Brady et al. 2017). It also shows that Democratic MCs' out-group tweets that are particularly toxic are most likely to be retweeted. In contrast, Republican out-group MCs' tweets have a much lower level of retweeting. **Figure 4b** plots different audience responses to distinct tweets by MCs. It shows that liberals generally retweet much more toxic tweets than do conservatives. Indeed, conditioning upon the partisan source (e.g., same party or other party), the graph shows that the gap between users' reactions grows as toxicity increases (i.e., the slopes on the liberal user lines are much steeper than those on the conservative user lines). Substantively, liberals retweeted toxicity from their own party 1.56 times more often than did conservatives. Moreover, liberals retweeted toxicity from the other party -8.30 times more often than did conservatives, where the negative sign indicates different directions in response.

In summary, retweeting occurs much more when Democratic MCs are toxic against Republicans than vice versa, and this reflects that liberals are more likely to retweet toxic tweets. This is consistent with Hypothesis 2. Liberals will retweet toxic critiques of Republicans, but conservatives are less likely to retweet toxic tweets of Democrats, from out-group elites.

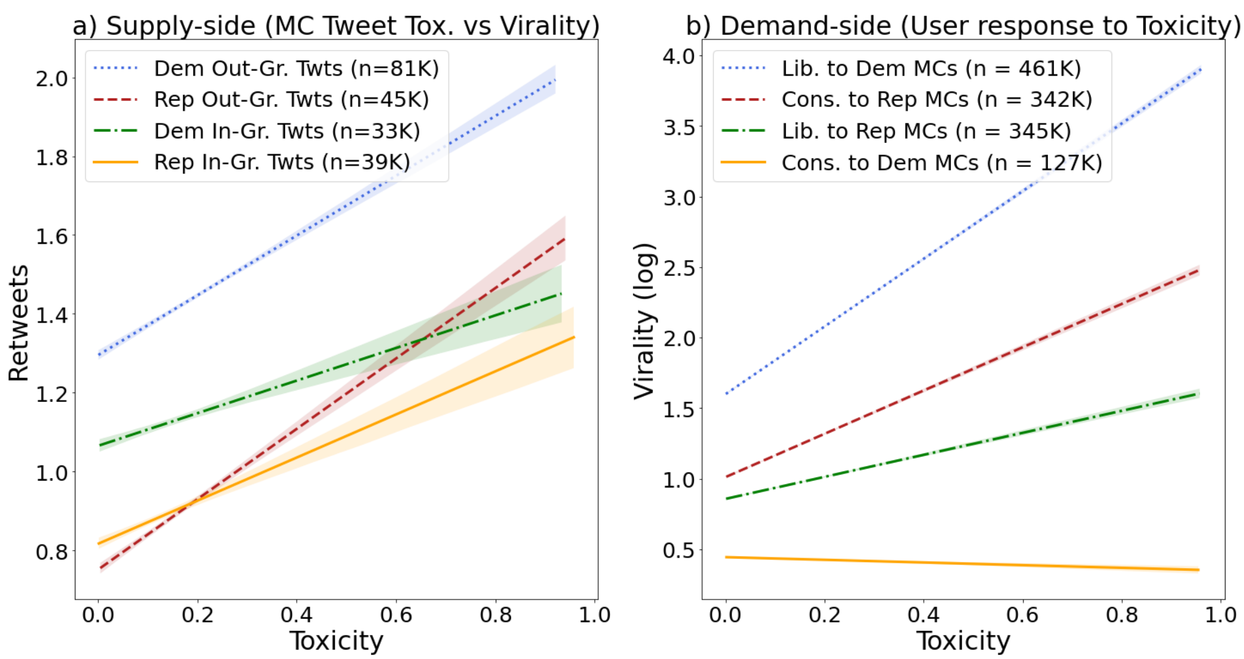

Figure 4: Virality versus toxicity, split by (a) partisan in-/out-group mentions and (b) user-level response. Out-group tweets are generally more viral and uncivil (a). Republican out-group tweets can be viral (red)—due to response from the left (a). Liberal users respond much more drastically to toxicity than do Conservatives (b).

In **Figure 5**, we provide more descriptions of the tweets. **Figure 5a** reports the average amount of toxicity in distinct types of tweets, showing that tweets about the out-group from either party contain more toxicity. This is particularly true among Democratic MCs (also see Frimer et al. 2022). **Figure 5b** and **c** describe the profile of retweeters in terms of ideology, whereas **Figure 5d** displays the aggregate amount of retweeting. **Figure 5b** shows in-group audience composition of retweets, revealing – consistent with the diversity dynamic – that Democratic MCs are overwhelmingly retweeted by liberals, while Republican MCs have a smaller in-group retweet percentage (although the clear majority are still in-party). That said, for both parties, out-group tweets draw more in-party users. This is further evidenced by **Figure 5c**, which shows Democratic MCs get retweeted by an extreme set of liberal users, while Republican MCs draw only a moderately conservative set of users. Looking specifically at out-group tweets, we see that the ideology of the audience is more extreme for out-group tweets from both parties. Out-group tweets are more toxic, homogenous, and extreme. As with **Figure 3,** we thus see some extremity impact, but, importantly, it is not the extremes on each side that evade topic diversity or endorse toxicity (as a symmetry hypothesis might suggest) – this stems from ideological predilection. The final panel (**Figure 5d**) reveals that Democratic MC outgroup tweets get the most retweets, reflecting that liberals are much more likely to retweet such communications (which tend to be toxic). Republican MC in-group tweets are more viral than their out-group tweets. This reflects conservatives' relative toxicity aversion (where toxicity is more present in out-party tweets) and the presence of Trump in the 2015 primaries and 2016 election driving more in-party retweeting. In sum, liberals engage in retweeting toxic critiques of Republicans, but conservatives are less likely to retweet toxic critiques of Democrats.

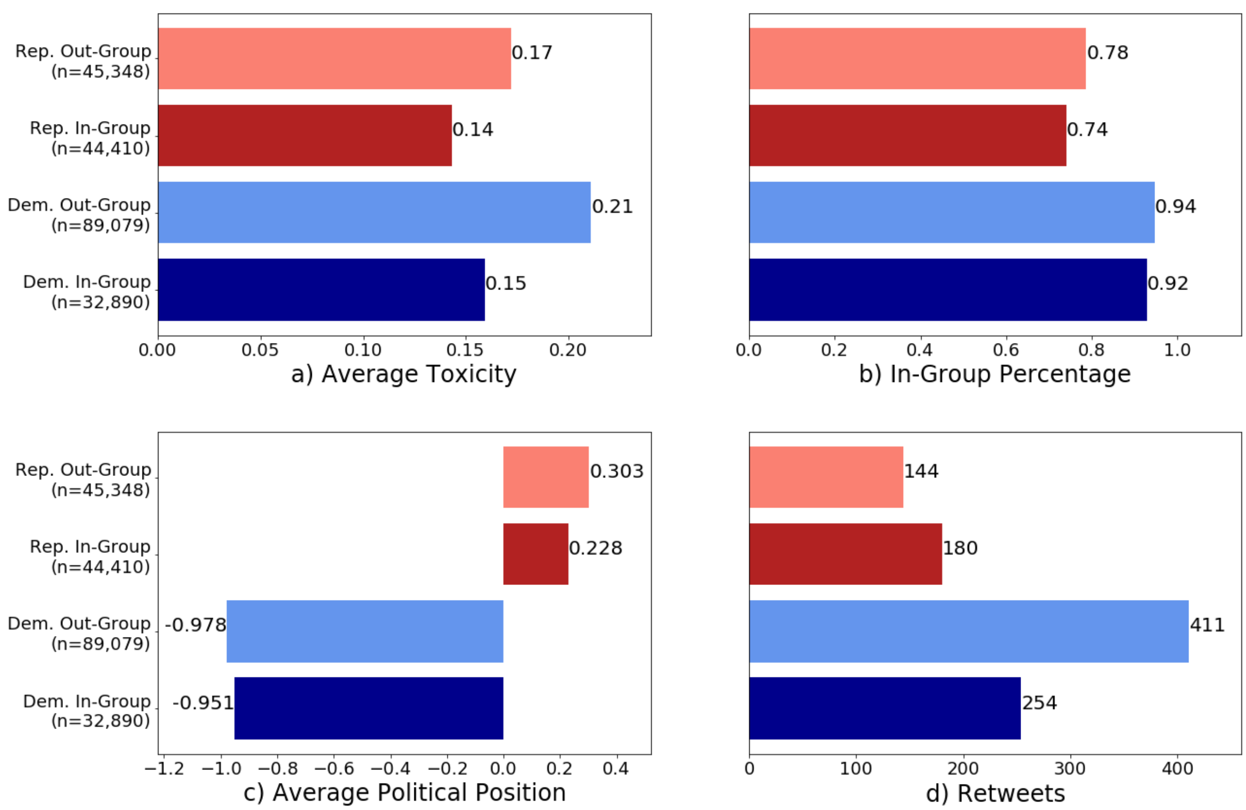

Figure 5: In-/out-group effects on (a) toxicity, (b) in-group composition, (c) political position, and (d) virality.

In the previous section, we showed that liberals are much more likely to retweet Republican issues than conservatives are to retweet Democratic issues. This section showed that liberals are much more likely to retweet discussions of Republicans, which tend to be toxic, than conservatives are to retweet discussions of Democrats (which also tend to be toxic). This aligns with each MC's incentives for maximizing virality. Therefore, liberals engage with (i.e., retweet), more diverse issue agendas and more discussion of the other side (and, thus, more toxicity). In short, our results show the audience diversity displayed in the audiences for MCs from different parties (Figure 1) does in fact reflect asymmetric engagement based on policy issues and toxicity.

Supply Side

It is useful to explore whether supply-side effects drive any of the asymmetry; that is, are the dynamics we display not only demand-driven (e.g., Frimer et al. 2022)? In **Table 1**, we present confusion matrices of MC tweets by non-toxic/toxic and not policy/policy content, split by party. The bottom row shows the ratio of toxic to non-toxic tweets per that policy category. We classify a tweet as toxic if it is greater than the mean plus one standard deviation (0.237). These results remain the same when using the mean (see Table S4, the additional standard deviation is more conservative.) The main finding can be captured by comparing the non-toxic/not policy category versus the toxic/policy category. For Democratic MCs, the respective percentages are 32% and 12%, whereas for Republicans they are 41% and 7%. In short, Democrats tweet much more about policy, and they do so with much more toxicity than do Republicans. The overall toxicity is 15% for Democratic MCs and 10% for Republican MCs. This suggests that the supply of content differs between partisan MCs.

Table 1: Policy versus non-toxic/toxic content, across Democratic and Republican MCs. Democrats tweet about policy with more toxicity than Republicans (4.4 versus 7.3 civility ratio).

(a) Democratic MCs	Not Policy	Policy	**N.Pol + Pol.**
Non-toxic	32.3%	52.6%	84.9%
Toxic	3.3%	11.8%	15.1%
Non-toxic / Toxic Ratio	9.739	4.447	–

(b) Republican MCs	Not Policy	Policy	**N.Pol + Pol.**
Non-toxic	41.4%	48.7%	90.2%
Toxic	3.1%	6.7%	9.8%
Non-toxic / Toxic Ratio	13.25	7.27	–

Of course, supply depends on demand, and, thus, we next look at the relationship of virality with MC choices. Using boosting and Shapley value analysis, we regress on virality (number of retweets) per MC party, using the following variables:

- Policy Difference: The supply difference of the 20 Comparative Agenda Project's Policy Topics.

- Policy: Whether it is a policy issue or not.
- Toxicity: Toxicity level of that tweet.
- Group status: Whether a tweet is in-group, out-group, or not related.

The results are in **Figure 6**. We find a clear asymmetry, with Democratic MC tweets becoming much more viral when they invoke the other group and when they are toxic. In contrast, Republican MC tweets become more viral based on policy, specifically when the policies are their policies presumably. This outlines the different incentives for MCs based on party if they are seeking to maximize their virality. Furthermore, we show that for MCs of both parties, Democrat-dominated topics generate more virality. This explains most clearly the toxicity asymmetry in supply: Democrats MCs tweet with more toxicity, as that draws them more of an audience. This makes sense given liberal users are relatively drawn to toxicity.

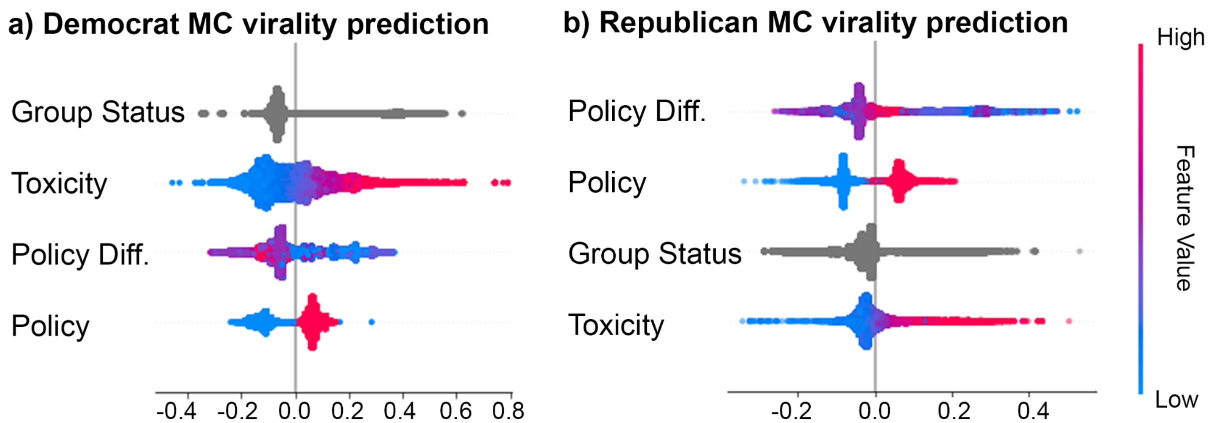

Figure 6: CatBoost regression on the four variables, then ranked top to bottom in terms of feature importance, separated by MC tweets (Democrats/left, Republicans/right). We find toxicity and group status more predictive for Democratic MC tweets and policy topic more for Republican MC tweets.

In **Figure 7**, we further look at the demand for toxicity. Here, we plot the effect of toxicity on retweeting on the y-axis (i.e., the regressed slope from virality on toxicity; tabular results shown in Table S3). This is a proxy for sensitivity (i.e., y = 0 indicates no change to virality based on toxicity). The x-axis shows the ideological dispersion on the given issue with each dot indicating an issue tweeted by a Republican MC (red) or a Democratic MC (blue). In line with what we have already shown, when toxicity is used in a tweet, liberals are substantially more likely to retweet it than are conservatives. The figure reveals a large effect. On Democratic issues, when there is toxicity, conservatives in fact are less likely to retweet it (lower beta/sensitivity). Even on Republican issues, they retweet toxic tweets at much lower rates than liberals do. There are *no* issues that conservatives are more likely to retweet when they are toxic, though this reflects somewhat that conservatives retweet less and also spend less time on Twitter in general. In sum, the asymmetry in retweets and, thus, information environments is a result of both supply and demand preferences. Republican MCs gain the attention of liberals for cross-cutting issues through toxicity (without compromising their in-group followers), but not vice versa.

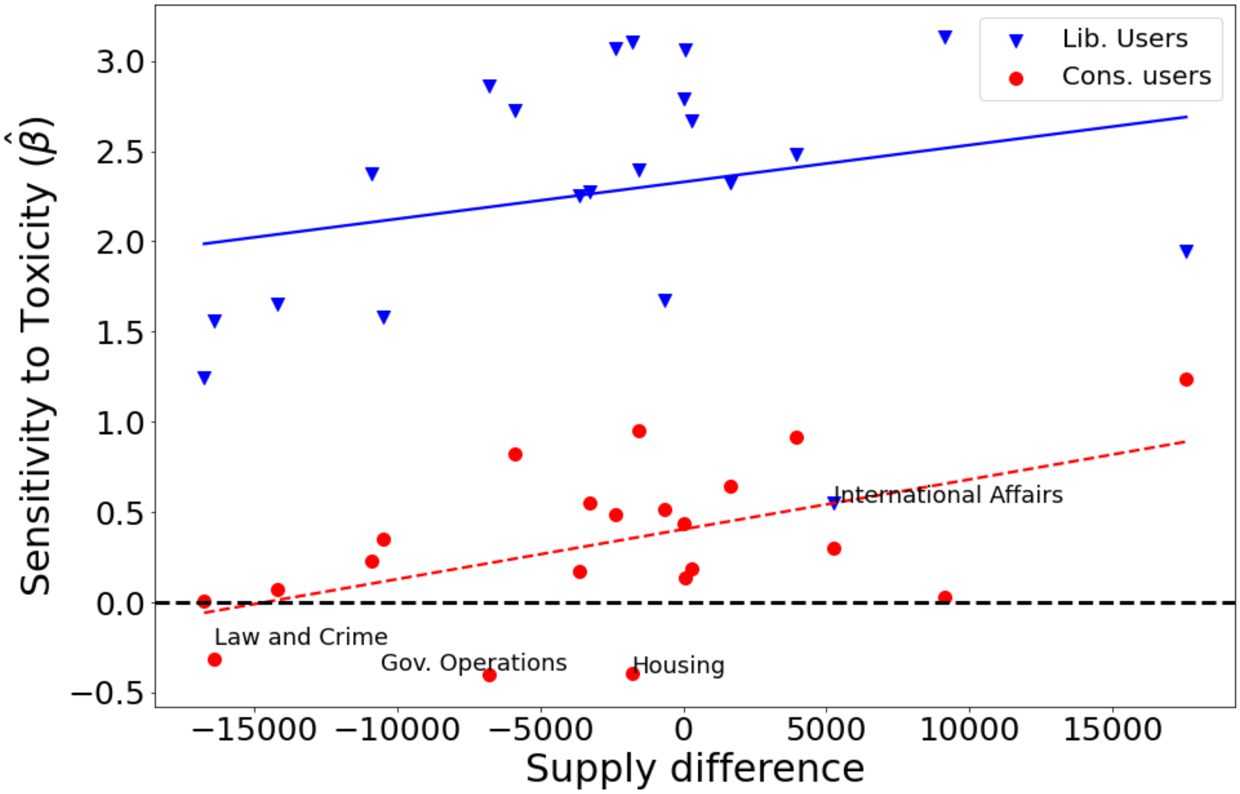

Figure 7: Partisan agenda items versus sensitivity to toxicity (slope of toxicity to virality), broken down by liberal and conservative users. Democrats seem to discuss international affairs with less toxicity. Labeled dots are explicitly discussed.

Discussion

We investigated whether there are asymmetric fractions to social media where liberals and conservatives spread information differently. We hypothesized ideological asymmetry due to cross-ideology policy and toxicity aversion among conservatives. We tested our predictions on Twitter, the only available social media data that provide the relevant measures.

We have three main findings. First, Republican members of Congress have a more diverse audience, implying that liberal users engage in more cross-cutting engagement, while conservatives pass along less diverse information. Second, conservative users are more selective with which policies they retweet, avoiding topics that are not on the conservative agenda. This may drive them to exist in an homogenous policy information bubble. This is consistent with work showing socially-driven (but not necessarily news-driven) policy bubbles (e.g., Barberá et al 2015, Cinelli et al. 2021, Mosleh et al. 2021). Ours is one of the only analyses, however, to document the asymmetry in the nature of those networks (although, see Boutyline and Willer 2017 for tie-level homophily). Relatedly, Eady et al. (2019) find that conservatives follow more politicians across the aisle. A key difference is they focus on a user subset based on follows, which is not required for retweeting. Our dataset also begins from the supply level, whereas they start from the demand level. Third, toxicity creates cross-cutting engagement asymmetrically: liberals respond to and retweet toxicity more than conservatives do.

Put together, our results show that if the forces that generate engagement for one party diminish engagement for the other, then we are pushed toward an asymmetric situation where one group engages with much less information. Specifically, liberals and conservatives react differently to information and, consequently, rebroadcast distinct types of information. Liberals, then, will view a host of policies and toxicity, while conservatives will mostly see issues on the conservative agenda and relatively less toxicity. These results add more evidence on behalf of the ideological asymmetry hypothesis from a new domain of testing. The findings also have potentially crucial downstream implications. Wojcieszak et al. (2022) find that most Twitter users do not follow partisan elites, which indicates a diffusion gap between partisan elites and the general public. It is plausible, if not probable, that most users receive more political information from other users via retweets. In that sense, the individuals we studied here serve as intermediaries as per the canonical theory of two-step information flow, where the general public receives political information not directly from elites but secondhand, from other citizens (e.g., (Katz et al. 1955; Weeks et al. ; Druckman et al. 2018; Carlson 2019).

While the composition of the intermediaries' audiences is not something we measured, there has been extensive work detailing the extent of homophily across liberal and conservative lines on Twitter. Boutyline and Willer (2017) find that conservative and ideologically extreme individuals exhibit greater levels of homophily and posit that this is caused by a stronger "preference for certainty." Putting this together with our finding that liberals retweet a more diverse array of topics suggests that this diversity is compounded with liberals' more heterophilous social ties. Vice versa for conservatives, sharing less diverse and homophilic ties likely creates an "information bubble." This bubble reflects liberals being both more likely to engage with heterogeneous content and with heterogeneous sources relative to conservatives. Moreover, the compounding effects of toxicity, tie homophily, and topic diversity mean the demand for out-group animosity may not necessarily arise from their expected constituents.

The consequence could be very distinct information environments for liberals and conservatives, with liberals being exposed to a wider array of policy topics but also more toxicity relative to conservatives. We earlier referred to these as polarized information bubbles since the information ecologies in which liberals and conservatives end up being very different from one another. This type of polarization (i.e., where there is a large gap between groups; Jost et al. 2022) is not the same as ideological polarization or affective polarization; here polarization refers to information content. That content, however, could contribute to other forms of polarization. For instance, since ideologues tend to have more extreme positions on policies that they own (i.e., conservative positions on conservative policies) (Egan 2013, Banda 2016), a homogenous policy bubble could increase ideological polarization. Alternatively, an information bubble that envelopes incivility from the other party (e.g., liberals engaging with conservative generate toxicity) could make them more affectively polarized (Druckman et al. 2019). More generally, these bubbles likely make cross-ideological or cross-partisan interactions and perspective-taking more difficult since those with distinct outlooks engage from very different places.[2] Finally, our findings clarify downstream implications for how social media should be

[2] Further, what appears to be a symmetric increase in virality to toxicity (Rathje et al. 2021) is actually driven by asymmetric reactions in the demand, where liberals react to Republican toxicity but not the other way around.

designed, if for the purpose of increasing exposure and engagement with outgroups. Symmetric approaches, such as blanket thresholds for toxicity or out-group topics, would have heterogeneous success based on the ideology of the user.

Methods

We extend Frimer et al. (2022)'s data source to include the following covariates (at the MC tweet level, i.e., the supply level): comparative agendas project (CAP) **topic**, in-/out-**group labeling**, MC **source party labeling**, and, upon scraping retweeter timelines, **user ideology** labeling. We use the **toxicity** scores provided by Frimer et al. (2022) (which are based on an independently validated operationalization of PerspectiveAPI's "toxicity").

Topic (Supervised): We trained a supervised machine learning classifier using labeled tweets from Russell (2018), with the assumption that the linguistic characteristics of tweets between the House and the Senate are similar. The dataset includes a total of 68,398 tweets, 45,402 tweets labeled with codes from the Comparative Agenda Project's (CAP's) codebook (Bevan 2019) and 22,996 labeled as non-related tweets.

The best models from Russell (2018) feature F1-scores of 79% and include augmentation via the Linguistic Inquiry and Word Count (LIWC) scores. LIWC is a gold standard package for computerized text analysis covering a range of psychological and topical categories and social, cognitive, and affective processes (Tausczik and Pennebaker 2010). Using these features, we implement deep learning, specifically a variant of BERT (Bidirectional Encoder Representations from Transformers), which is a variant of the transformer architecture fine-tuned for the English language. We implement the large BERTweet architecture (a specialized BERT model for Twitter and tweets), using a pre-trained model from Nguyen and colleagues (2020). We train two classifiers – one for granularly extracting topics and another for distinguishing if policies are mentioned. Our first (multiclass) model produces 86.4% accuracy for classification across all 20 policy topics. Our second (binary) model yields 88.5% accuracy, both of which out-perform those by Russell. Full classification results are given in Table S1 and Table S2 in the Supplementary Information.

Group-labeling: We implemented the same approach as Rathje et al. (2021), using a list of keywords that refer to Democrats (such as left-wing, liberal, etc.), Republicans (right-wing, conservatives, etc.), and the most popular politicians based on polling from YouGov.

Source Party Labeling: A total of 831 members of Congress had Twitter labels and were labeled. Several have changed accounts, and these were manually merged. For those who have switched parties, their party affiliation at the time of the tweet was used.

User ideology: Ideology scores were labeled based on a user's timeline history (up to 3,200 tweets per the limit of the API). This yielded a total of 3,522,734,792 tweets. URL domains were extracted from the timeline. Associated with each URL is a political score of [-2, -1, 0, 1, 2], corresponding to "left," "left-center," "center," "right-center, and "right," based on Media-bias/fact-check. A weighted political score is then calculated based on the proportion of tweets from each category. This approach has been adopted in many contexts related to Twitter ideological modeling (Ferrara et al. 2020; Huszár et al. 2022; Chang et al. 2021; Chen et al. 2021). Only tweets prior to a retweet are used to attribute ideology to allow for the possibility that a retweeter switch is small.

References

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(37), 9216–9221.

Ballard, A., DeTamble, R., Dorsey, S., Heseltine, M., & Johnson, M. (2022). Incivility in Congressional Tweets. *American Politics Research*, *50*(6), 769-780.

Barberá, P., Casas, A., Nagler, J., Egan, P. J., Bonneau, R., Jost, J. T., & Tucker, J. A. (2019). Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data. *The American Political Science Review*, *113*(4), 883–901.

Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science*, *26*(10), 1531–1542.

Baumgartner, F. R., Breunig, C., & Grossman, E. (2019). The Comparative Agendas Project. In *Comparative Policy Agendas* (pp. 3–16).

Baumgartner, F. R., Green-Pedersen, C., & Jones, B. D. (2006). Comparative studies of policy agendas. *Journal of European Public Policy*, *13*(7), 959–974.

Bormann, M., Tranow, U., Vowe, G., & Ziegele, M. (2021). Incivility as a Violation of Communication Norms—A Typology Based on Normative Expectations toward Political Communication. *Communication Theory: CT: A Journal of the International Communication Association*, *32*(3), 332–362.

Boutyline, A., & Willer, R. (2017). The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, *38*(3), 551–569.

Brady, W. J., Crockett, M. J., & Van Bavel, J. J. (2020). The MAD Model of Moral Contagion: The Role of Motivation, Attention, and Design in the Spread of Moralized Content Online. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, *15*(4), 978–1010.

Brady, W. J., Wills, J. A., Burkart, D., Jost, J. T., & Van Bavel, J. J. (2019). An ideological asymmetry in the diffusion of moralized content on social media among political leaders. *Journal of Experimental Psychology. General*, *148*(10), 1802–1813.

Carlson, T. N. (2019). Through the grapevine: Informational consequences of interpersonal political communication. *The American Political Science Review*, *113*(2), 325-339.

Chang, H. C. H., Chen, E., Zhang, M., & Muric, G. (2021). Social bots and social media manipulation in 2020: the year in review. In *Handbook of Computational Social Science*, Volume 1 (pp. 304-323). Routledge.

Chen, C. X., Pennycook, G., & Rand, D. G. (2021). *What Makes News Sharable on Social Media?*

Chen, E., Chang, H., Rao, A., & Lerman, K. (2021). COVID-19 misinformation and the 2020 US presidential election. The Harvard Kennedy School Misinformation Review 1, no. 7.

Conover, M., Ratkiewicz, J., Francisco, M., Goncalves, B., Menczer, F., & Flammini, A. (2011). Political Polarization on Twitter. *Proceedings of the International AAAI Conference on Web and Social Media*, *5*(1), 89–96.

Ditto, P. H., Liu, B. S., Clark, C. J., Wojcik, S. P., Chen, E. E., Grady, R. H., Celniker, J. B., & Zinger, J. F. (2019). At Least Bias Is Bipartisan: A Meta-Analytic Comparison of Partisan Bias in Liberals and Conservatives. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, *14*(2), 273–291.

Druckman, J. N., Gubitz, S. R., & Lloyd, A. M. (2019). How incivility on partisan media (de) polarizes the electorate. *The Journal of Politics*, *81*(1), 291-295.

Druckman, J. N., & Jacobs, L. R. (2015). *Who Governs?* University of Chicago Press.

Druckman, J. N., Klar, S., Krupnikov, Y., Levendusky, M., & Ryan, J. B. (2022). (Mis)estimating affective polarization. *The Journal of Politics*, 84(2), 1106-1117.

Druckman, J. N., Levendusky, M. S., & McLain, A. (2018). No need to watch: How the effects of partisan media can spread via interpersonal discussions. *American Journal of Political Science*, *62*(1), 99–112.

Enders, A., Farhart, C., Miller, J., Uscinski, J., Saunders, K., & Drochon, H. (2022). Are Republicans and Conservatives More Likely to Believe Conspiracy Theories? *Political Behavior*, 1–24.

Ferrara, E., Chang, H., Chen, E., Muric, G., & Patel, J. (2020). Characterizing social media manipulation in the 2020 US presidential election. *First Monday, 25(11)*.

Frimer, J. A., Aujla, H., Feinberg, M., Skitka, L. J., Aquino, K., Eichstaedt, J. C., & Willer, R. (2022). Incivility Is Rising Among American Politicians on Twitter. *Social Psychological and Personality Science*, 19485506221083811.

Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. presidential election. *Science*, *363*(6425), 374–378.

Grossmann, M., & Hopkins, D. A. (2016). *Asymmetric politics: Ideological Republicans and group interest Democrats*. Oxford University Press.

Guay, B., & Johnston, C. D. (2022). Ideological asymmetries and the determinants of politically motivated reasoning. *American Journal of Political Science*, *66*(2), 285–301.

Guay, B., Berinsky, A. J., Pennycook, G., & Rand, D. (2022). *Examining Partisan Asymmetries in Fake News Sharing and the Efficacy of Accuracy Prompt Interventions*.

Guess, A. M. (2021). Experiments using social media data. *Advances in Experimental Political Science*.

Guess, A. M., & Munger, K. (2022). Digital literacy and online political behavior. *Political Science Research and Methods*, 1-19.

Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, *5*(1), eaau4586.

Hemphill, L., & Schöpke-Gonzalez, A. M. (2020). Two Computational Models for Analyzing Political Attention in Social Media. *Proceedings of the International AAAI Conference on Web and Social Media*, *14*, 260–271.

Huszár, F., Ktena, S. I., O'Brien, C., Belli, L., Schlaikjer, A., & Hardt, M. (2022). Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences of the United States of America*, *119*(1).

Jost, J. T. (2021). *Left and Right: The Psychological Significance of a Political Distinction*. Oxford University Press.

Jost, J. T. (2017). Ideological asymmetries and the essence of political psychology. *Political Psychology*, *38*(2), 167–208.

Katz, E., Lazarsfeld, P. F., & Columbia University. Bureau of Applied Social Research. (1955). *Personal Influence; the Part Played by People in the Flow of Mass Communications, by Elihu Katz and Paul F. Lazarsfeld. With a Foreword by Elmo Roper*.

Minozzi, W., Song, H., & Lazer, D. M. J. (2020). The incidental pundit: Who talks politics with whom, and why? *American Journal of Political Science*, *64*(1), 135-151

Mutz, D. C. (2006). *Hearing the Other Side: Deliberative Versus Participatory Democracy*. Cambridge University Press.

Mutz, D. C. (2015). *In-Your-Face Politics*. Princeton University Press.

Neblo, M. A., Esterling, K. M., & Lazer, D. M. J. (2018). *Politics with the People: Building a Directly Representative Democracy*. Cambridge University Press.

Nguyen, D. Q., Vu, T., & Nguyen, A. T. (2020). BERTweet: A pre-trained language model for English Tweets. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 9-14. 2020.

Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., & Petersen, M. B. (2021). Partisan Polarization Is the Primary Psychological Motivation behind Political Fake News Sharing on Twitter. *The American Political Science Review*, *115*(3), 999–1015.

Pennycook, G., & Rand, D. G. (2021). The Psychology of Fake News. *Trends in Cognitive Sciences*, *25*(5), 388–402.

Petrocik, J. R. (1996). Issue Ownership in Presidential Elections, with a 1980 Case Study. *American Journal of Political Science*, *40*(3), 825–850.

Rathje, S., Van Bavel, J. J., & van der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(26).

Riker, W. H. (1993). *Agenda Formation*. University of Michigan Press.

Rogers, N., & Jost, J. T. (2022). Liberals as Cultural Omnivores. *Journal of the Association for Consumer Research*, *7*(3), 255–265.

Schraufnagel, S., Casas, N., Bacharz, T., Holm, N., & Miller, C. (2021). Legislative conflict: are ideologues more uncivil?. In Political Incivility in the Parliamentary, Electoral and Media Arena (pp. 69-86). Routledge.

Sharma, S., & Verma, H. V. (2018). Social Media Marketing: Evolution and Change. In G. Heggde & G. Shainesh (Eds.), *Social Media Marketing: Emerging Concepts and Applications* (pp. 19–36). Springer Singapore.

Simon, A. F. (2002). *The Winning Message: Candidate Behavior, Campaign Discourse, and Democracy*. Cambridge University Press.

Theocharis, Y., Barberá, P., Fazekas, Z., & Popa, S. A. (2020). The Dynamics of Political Incivility on Twitter. *SAGE Open*, *10*(2), 2158244020919447.

Weeks, B. E., Ardèvol-Abreu, A., & de Zúñiga, H. G. (2017). Online influence? Social media use, opinion leadership, and political persuasion. *International Journal of High Risk Behaviors & Addiction*, *29*(2), 214-239.

Table S1: Accuracy and F1 values per topic from deep learning model (using the BERTweet architecture). Each policy topic is labeled and derived from the comparative agenda project.

Label	Policy Topic	Accuracy
0	Not Policy Related	0.666
1	Macroeconomics	0.845
2	Civil Rights	0.883
3	Health	0.948
4	Agriculture	0.893
5	Labor	0.762
6	Education	0.929
7	Environment	0.782
8	Energy	0.834
9	Immigration	0.960
10	Transportation	0.885
12	Law and Crime	0.919
13	Social Welfare	0.885
14	Housing	0.925
15	Domestic Commerce	0.821
16	Defense	0.847
17	Technology	0.808
18	Foreign Trade	0.708
19	International Affairs	0.815
20	Government Operations	0.859
21	Public Land	0.679
Average Accuracy		**0.864**
F1 Score		**0.864**

Table S2: BERTweet classifier accuracy for policy versus not policy.

Label	**Accuracy**
Policy	0.926174
Not Policy	0.824503
F1 Score	0.88518

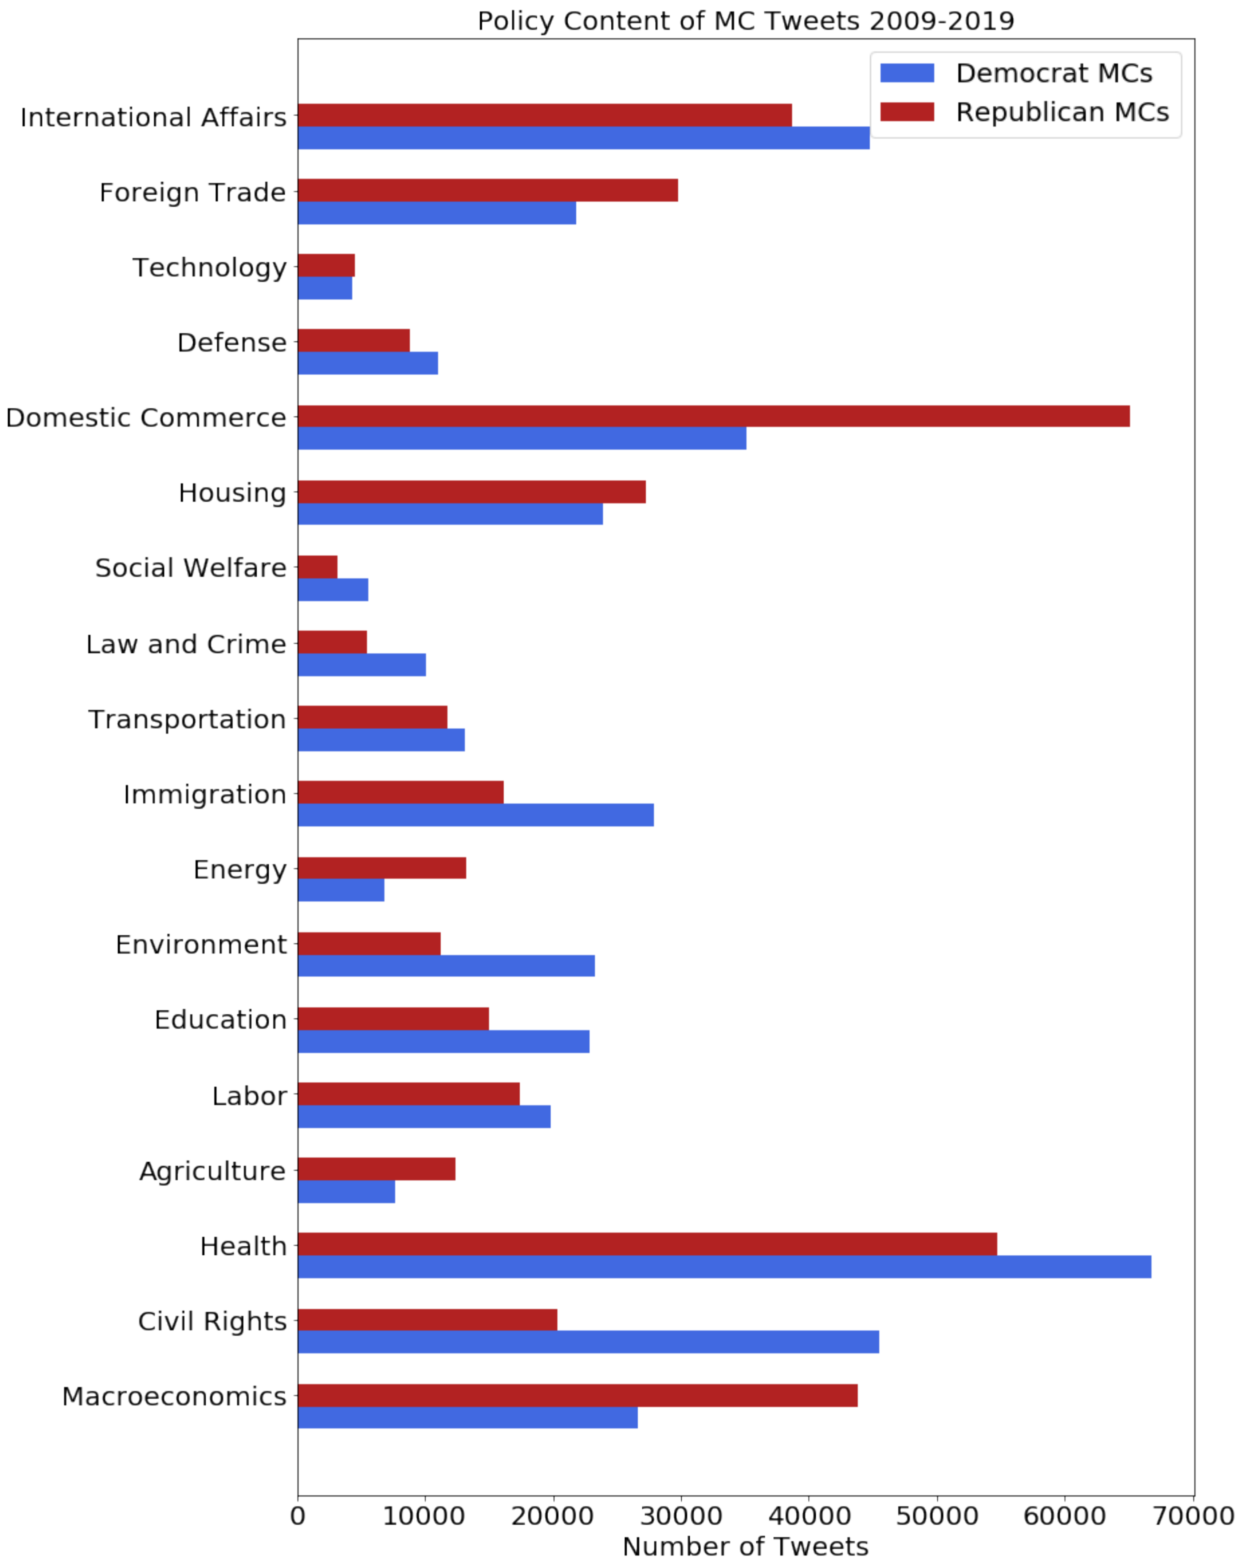

Fig S1: Raw values for the supply in the axes of Figure 2 and Figure 3. Blue denotes number of tweets from Democratic MCs; red denotes number of tweets from Republican MCs.

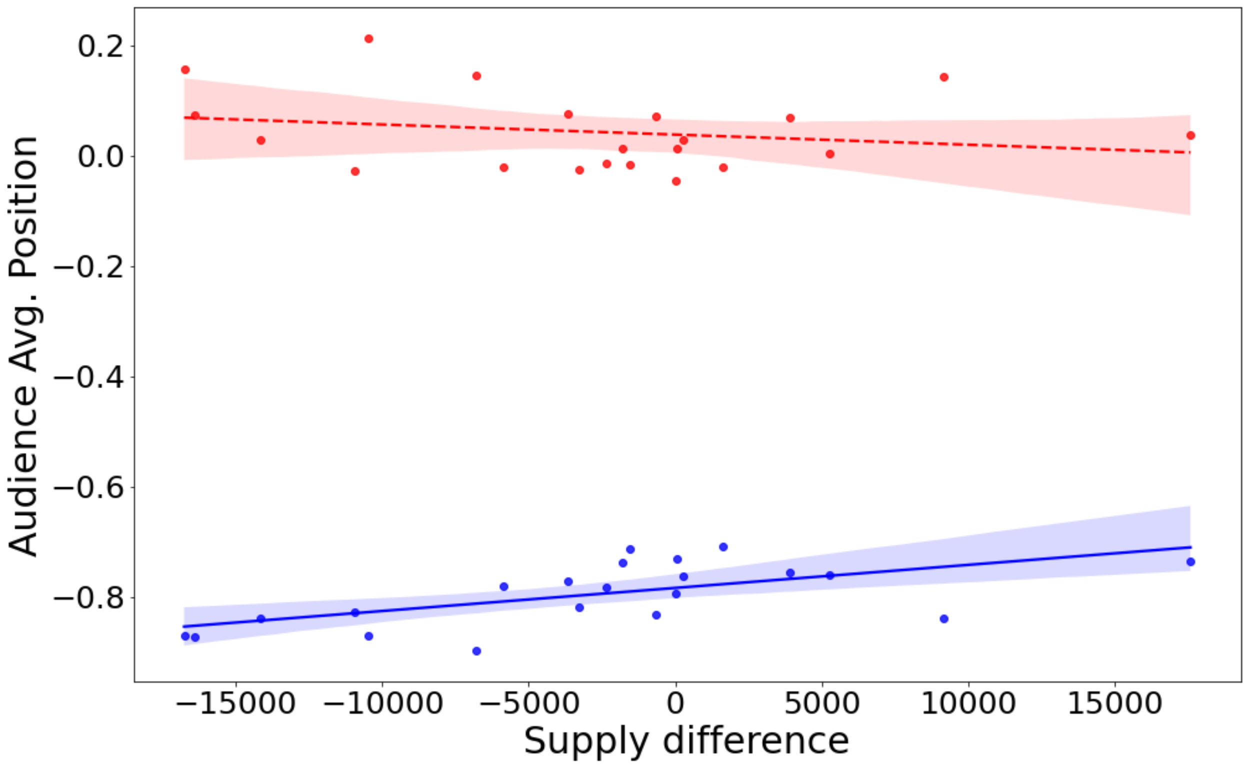

Fig S2: Average political position of audience, by topic and source partisanship (Party of the MC). Validates "starting points" for diversity values in Figure 2.

Table S3: 20 Comparative Agenda Project's Policy topics and sensitivities to toxicity. Beta denotes the slope of toxicity versus retweets (virality)

Names	**Dem. Beta**	**Rep. Beta**	**Supply Diff.**
Defense	1.947447	1.236853	17,578
Macroeconomics	3.134636	0.026845	9,148
International Affairs	0.5516	0.301626	5,246
Energy	2.485579	0.919327	3,931
Agriculture	2.325239	0.644451	1,647
Domestic Commerce	2.665562	0.184382	272
Foreign Trade	3.063597	0.137203	73
Public Lands	2.78943	0.4341	32

Not Policy	1.675464	0.511957	-668
Transportation	2.398488	0.952559	-1,552
Housing	3.103826	-0.391858	-1,796
Technology	3.068483	0.48698	-2,347
Social Welfare	2.276659	0.551307	-3,261
Labor	2.257501	0.172607	-3,648
Education	2.729321	0.821523	-5,869
Government Operations	2.86418	-0.401003	-6,775
Immigration	1.584619	0.347557	-10,484
Environment	2.378217	0.230079	-10,916
Health	1.652244	0.068237	-14,168
Law and Crime	1.559014	-0.316044	-16,394
Civil Rights	1.247989	0.009474	-16,749

For Democratic MCs, the respective percentages are 37% and 28% whereas for Republicans, they are 37% and 18%. In short, Democrats tweet much more about policy and they do so with much more toxicity than Republicans. The overall toxicity is 37% for Democratic MCs and 29% for Republican MCs. This suggests that the supply of content differs between partisan MCs.

Table S4: Policy versus non-toxic / toxic content, across Democrat and Republican MCs. Democrats engage policy with much more toxicity than Republicans (1.8 versus 1.3 non-toxic / toxic ratio). Furthermore, Democrats are in general more toxic (37% versus 29% of content).

(a) Democratic MCs	Not Policy	Policy	**N.Pol + Pol.**
Non-toxic	0.263	0.365	0.628
Toxic	0.093	0.279	0.372
Non-toxic / Toxic Ratio	2.815	1.310	

(b) Republican MCs	Not Policy	Policy	**N.Pol + Pol.**
Non-toxic	0.343	0.37	0.714

Toxic	0.102	0.184	0.286
Non-toxic / Toxic Ratio	3.35	2.01	