

Most users do not follow political elites on Twitter; Those who do, show overwhelming preferences for ideological congruity.*

Magdalena Wojcieszak ^{†‡} Andreu Casas [§] Xudong Yu ^{*} Jonathan Nagler [¶]
Joshua A. Tucker [¶]

May 2022

Few users follow politicians, pundits, and news media on Twitter. Those who do, strongly prefer ideological congruity.

Abstract

We offer comprehensive evidence of preferences for ideological congruity when people engage with politicians, pundits, and news organizations on social media. Using four years of data (2016-2019) from a random sample of 1.5 million Twitter users, we examine three behaviors studied separately to date: (a) following of in-group vs. out-group elites, (b) sharing in-group vs. out-group information (retweeting), and (c) commenting on the shared information (quote tweeting). We find the majority of users (60%) do not follow any political elites. Those who do, follow in-group elite accounts at much higher rates than out-group accounts (90% vs. 10%), share information from in-group elites 13 times more frequently than from out-group elites, and often add negative comments to the shared out-group information. Conservatives are twice as likely as liberals to share in-group vs. out-group content. These patterns are robust, emerge across issues and political elites, and regardless of users' ideological extremity.

Introduction

Social media platforms are the primary source of political information for a growing number of citizens (1, 2). This profoundly shifts the ways users encounter information and places individuals in unique media environments characterized by information flows curated by the users themselves and filtered through their social contacts. Although these changes facilitate exposure to different perspectives

and connections with diverse people, they may also lead to the emergence of insular online communities where users follow members of their political group and share information consistent with their views (3, 4). Such insular communication is feared to fuel extremity, exacerbate inter-party hostility, and ultimately thwart consensual governance (5, 6, 7, 4, 8).

Given their democratic consequences, burgeoning research aims to describe such political congruity—which we refer to concisely as *political biases* in this manuscript—in users' behaviors on social media (9). Existing evidence regarding the prevalence of these political biases, however, is inconclusive. Some studies show that bloggers primarily connect to sources from their ideologically in-group (10, 11) and that social media users exchange political information with co-partisans (12), disproportionately follow politicians (13) and other users (14) from within their ideological group, and tend to share messages from their own party rather than those from other parties (15). Yet, others suggest that opposing partisans largely follow the same political accounts (16) and discuss *non-political* topics across party lines (12). The extent of ideological asymmetries in these behaviors is also unclear: sometimes liberals are found to engage in more like-minded following than conservatives (16, 17), other times they are shown to share more information across party lines (12, 18), and yet other studies show that these asymmetries depend on the behaviors examined (14). This mixed evidence may be due to distinct samples and different behaviors studied in past work, such as who users follow (16) versus whether they retweet information about a few selected policies (12, 18). Lacking is a comprehensive approach that integrates these distinct behaviors in one analysis on a large random sample.

Here, we advance past work on political biases in users' online behaviors in three key ways. First, we fo-

*M.W. corresponding author: mwojcieszakucdavis.edu.

[†]University of California-Davis

[‡]University of Amsterdam

[§]Vrije Universiteit Amsterdam

[¶]New York University, Center for Social Media and Politics

cus on users’ engagement with information produced by political elites — arguably the most influential and politically active users — testing if these engagements are in fact motivated by political bias. Unlike ordinary users, politicians, pundits, and news media contribute the overwhelming majority of political content and dominate online discussions (19). On social media platforms, politicians engage with constituents, “broadcast” information about their activities, and shape the political agenda (19). In turn, journalists and news media increase the reach of their stories and play a central role in the content that gets shared on social media platforms (20). The online activities of journalists, pundits, and news outlets are especially important as many citizens receive news through social network sites, not from producers directly (21).

Elite communication is also central to attitudes and behaviors of the electorate and may exacerbate unprecedented partisan conflicts in America (5). Elite cues can distort citizens’ policy preferences (22, 23, 24) and — by making inter-party divisions clearer — polarize their attitudes (22, 25). Elite communication makes people’s partisan identities more salient and casts politics as us-versus-them conflict, intensifying out-group hostility (26). Therefore, accounting for whether and how users engage with politicians, pundits, and news organizations on social media has clear societal implications.

Second, we integrate three distinct behaviors studied separately in past work and — and in doing so — offer a more complete portrayal of political biases in users’ behavior on social media. We examine (a) the following of in-group versus out-group political elites, (b) the sharing of their messages (i.e., retweeting), and (c) adding comments to the shared messages of in- versus out-group elites (i.e., quote tweets). These three behaviors reflect distinct affordances of social media platforms and have different implication for users themselves and online discourse at large. Following — although important — does not guarantee exposure to and interactions with the elite accounts one follows. It is also rather passive and “private” in nature (apart from mere exposure, which we cannot examine). In turn, sharing and commenting represent more active engagements with elite messages, are more public in nature, and — as such — have a greater impact on the online public sphere. Diffusion through sharing increases information reach and shows how messages spread on social media (16). Also, sharing often indicates trust in the message and its source, agreement with the message (27), and users’ motivation to strategically construct and present their political identity to their online networks (21).

Sharing or retweeting information, moreover, affords users an under-analyzed ability of adding comments to the shared messages. Theoretical frameworks of social identity establish that once individuals identify with a group, they aim to maintain and enhance their identity by positively distinguishing themselves and their in-group from the out-group (28, 29, 30). In the context of partisan communication on social media, this need for distinctiveness — or a unique positive social identity — can be met by sharing in-group content with positive comments. However, another way to enhance in-group status and achieve the need for distinctiveness is to attack or derogate the out-group (31). This phenomenon is visible in partisan media, which do cover the out-party, but often in derogatory ways, referring to it as Nazis or communists (32, 33, 34).

This phenomenon, although consequential, has been largely overlooked in the online environment. In the context studied, users may well retweet messages from across the political aisle, but do so to criticize the message or its source. Assessing the *sentiment* of the comments added to the shared elite information can show if the assumed “endorsement by sharing” is undercut by mocking or criticism. If it is (e.g., conservatives retweeting Biden only to mock him), encountering information from out-group elites could actually have the effect of reinforcing political biases online.

Our third core contribution lies in putting in perspective extant concerns about political biases in online communication by examining a large random sample of Twitter users that contains both politically engaged as well as politically disengaged users. Most past work *a priori* explores online behaviors among politically engaged citizens, such as those who use political hashtags (10, 11, 18, 35, 36). Those politically active users are far from representative of social media users at large, however (37). Those citizens are more strongly partisan (38), and — as such — driven by confirmation bias or the need to “be right” and protect their viewpoints and political identity (39). Accordingly, studies on politically engaged users typically find political biases in their online behaviors, namely following and engaging with one’s partisan in-group. The impression that political biases on social media are prevalent is further reinforced by the fact that the highly partisan users discuss salient policies online substantially more than other groups of users (19). Furthermore, those who are ideologically and affectively polarized are the ones who also amplify highly polarizing and sometimes misinformative content online (40, 41). As a result, partisan and politically active users increase the visibility of political

and polarizing information among their less partisan and engaged friends and followers and send the signal to recommendation algorithms to further amplify this content on social media platforms (42, 43). In short, extant focus on this small group portrays an important yet incomplete picture of allegedly prevalent political biases in online behaviors.

In contrast, a solid majority of social media users are unlikely to engage in such biased, partisan, and polarizing online behaviors. Growing evidence shows that news and politics constitute a small fraction of people’s information diet. This is the case on social media: news makes up only 4% of News Feed on Facebook (44), public affairs more broadly comprise 1.8% of the average News Feed of college students (45), and only about 1 in 300 outbound clicks from Facebook correspond to substantive news (46). This is also the case online more broadly: only between 2% (47) and 7%–9% (48) of all URLs visited by large samples of Americans are news domains and - across mobile and desktop - news comprises only 4.2% of total online consumption (49).

As such, the aforementioned biases in online communication are likely to follow the power law distribution, in that the majority of users are likely to be politically disengaged and not follow politicians, pundits, and news media organizations, and the small group that does engage with political elites is likely to be vocal, visible, and politically biased. Such power law distribution is increasingly detected with regard to (problematic) political behaviors online. On social media, a small share of highly active users produces the vast majority of content (e.g., the most active 25% of US Twitter users produces 97% of all tweets (37)), a small fraction of people share fake news online (e.g., 1% of Twitter users accounted for 80% of exposures to fake news sources and 0.1% accounted for 80% of all fake news sources shared (50)), and small groups of extreme partisans generate a majority of views to and engagements with partisan media on platforms (16). Because we start with an incredibly large random sample, we can offer generalizable evidence on these power laws in engagement with politicians, pundits, and news media among diverse, politically inclined and not, ordinary Twitter users.

To examine whether users engage with in- versus out-group political elites in ways that reinforce political biases, we rely on four years of data (2016-2019) from a random sample of about 1.5 million Twitter users. We study whether they follow over 2,500 American political elite accounts and also examine instances in which they shared or quoted tweets from these accounts, which includes about 20 million

retweets in total. After using a validated method for estimating the ideology of regular and elite accounts (12) and a Convolutional Neural Network for classifying the sentiment of the quotes (51), we use these data to address five progressively specific questions. (1) What is the proportion of users who follow political elites? Among those who do: (2) What proportion follow in-group vs. out-group elites; (3) What is the proportion of in-group vs. out-group elite information shared by users?; (4) What is the sentiment - positive, neutral, or negative - of the commentary added to the tweets they share from in-group vs. out-group elites?; and (5) Are there ideological asymmetries in these online behaviors?

Our analyses yield three clear pictures. The first is of a political vacuum on Twitter. The majority of American Twitter users follows no politicians, pundits, or news media (59.6% of our sample) and only 23% follow more than 3 political elites. Second, when focusing on this small group, the second picture that emerges is one of pronounced political biases in users’ behaviors *vis a vis* political and media elites. In-group elites are followed at much higher rates than out-group accounts (around 90% vs. 10%) and tweets from in-group elites are shared overwhelmingly more frequently than out-group tweets (at about a 13:1 ratio). Indeed, the sharing of out-group elite is extremely limited, accounting for 7% of the retweets in our sample. Moreover, the sentiment of the comments on *out-group* tweets works to introduce further bias to this cross-party sharing: quote tweets (or the added commentary) not only constitute a much larger percentage of out-group than in-group shares (about 1/3 *v.* less than 1/10) but also are negative two thirds of the time when accompanying out-group tweets. Across 20 million retweets of elite content, only 5% were of out-group elites *without* any negative commentary. Third, we find important ideological asymmetries: conservative users are roughly twice as likely as liberals to share in-group *vs.* out-group content, as well as to add negative commentary to out-group shares. These patterns hold when accounting for the proportion of in-group vs. out-group elites followed (i.e., users share in-group elites *not* only because they follow them more), across elite actors (but especially for politicians) and numerous issues, ranging from the economy to civil rights.

Results

The Materials and Methods section provides additional details on the data and methodology, and the Supplementary Information (SI) offer detailed de-

scription of all materials and methods used within this study as well as additional robustness checks, extended discussion of the machine learning classifiers as well as alternative classifications. For four years (2016-2019), we tracked the messaging activity of a random sample of 1,437,774 Twitter users as well as an extensive set of 2,624 political elite accounts. In this project, we focus on national political elites due to their visibility and national-level influence on public opinion and the political process. This focus is also in line with recent findings, which suggest that American political behavior is increasingly more nationalized, with voters being more engaged with and knowledgeable about national than local level issues and politics (52).

Our basic descriptives offer an important corrective to extant concerns about political biases on social media. We find that only 40.4% of the users (580,921 out of 1,437,774) follow 1 of these elites or more. The remaining 59.6% follow no political actors whatsoever, even though the list includes key politicians (e.g. Donald Trump, Joe Biden), prominent pundits (e.g. Rachel Maddow, Sean Hannity), and the most popular media outlets (e.g. MSNBC, FoxNews) (see SI.10 for the full distribution of how many elite accounts users follow). Only 23% of users follow 3 or more elite accounts. Overall, the majority of American Twitter users are not sufficiently interested in politics to follow even a single political or media elite from our list (RQ1). To put this finding in perspective, we compare the following of political elites to that of celebrities related to music (e.g. Katy Perry), sports (e.g. LeBron James), TV-film (Oprah Winfrey, Kim Kardashian, Tom Hanks), literature (Paulo Coelho), and fashion (Kendall Jenner), among other popular celebrities. We rely on a publicly available list of top 1000 Celebrity Twitter accounts (<https://gist.github.com/mbejda/9c3353780270e7298763>.) We find that 70.7% of the users in our full sample follow at least one of the celebrities in the list, and over 53.2% follow at least 3 of them, a stark difference compared to how many follow political elites. On average the full set of users we study follow about 10.7 celebrities, but only 3.35 of the journalists, 1.52 of the politicians, and 1.13 of the media accounts on our list. Further details on these differential patterns are shown in SI.12. To examine whether the politically engaged users follow in- *versus* out-group elite accounts (RQ2), we first use the Bayesian Spatial Following model estimated by Barbera (12) to assign ideological scores, on the same continuous scale, to political elites and ordinary Twitter users. We obtained an ideology score for 180,203 users. Although this group of politically

engaged users is small relative to our total random sample (13% of 1,437,774), it accounts for 86% of all the shares of elite accounts sent by all ordinary users in our sample during the four years analyzed, and is also politically consequential, as we note in the discussion.

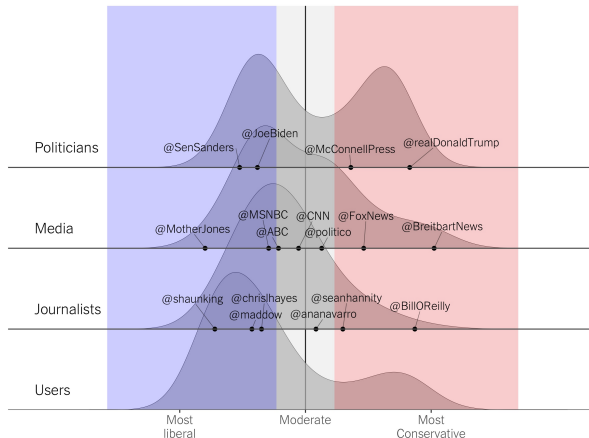


Figure 1: **Ideology estimates.** Distribution of the estimated ideology of the actors and ordinary users under study.

We then split the continuous ideology scale into a liberal, moderate, and conservative space, and classify the accounts accordingly. Users with a score lower than 0 were classified as liberal, between 0 and 1.2 as moderates, and higher than 1.2 as conservative. Figure 1 shows the distribution of the ideology and the face validity of our measure. We focus on liberal and conservative elites and ordinary users (i.e., the blue and red areas in Figure 1), excluding moderates, for whom a clear political in-group and out-group cannot be determined (keeping 1,721 elite accounts and 151,063 ordinary users). Importantly, SI.13 shows that our findings hold when including and classifying moderates into liberals and conservatives (using the vertical line in Figure 1 as cut-off point). In total, we examine 407 politicians (193 liberal and 214 conservative; 82 moderate accounts were excluded), 1,234 pundits (969 liberal and 265 conservative, 782 moderate accounts were excluded), and 78 media organizations (51 liberal and 27 conservative, 39 moderate accounts were excluded). We also examine 115,589 liberal users and 35,474 conservative users (excluding 29,140 moderate ones from the 180,203 for which we obtained ideology scores).

Pronounced differences depicted in the left panel of Figure 2.A reveal clear biases among politically active users: In-group politicians, pundits, and media are followed at much higher rates than out-group political

elites (around 90% vs. 10%). We do not observe any major difference between conservative (88.5% in-group vs 11.5% out-group elites followed) and liberal (89.5% in-group vs. 10.9% out-group) users.

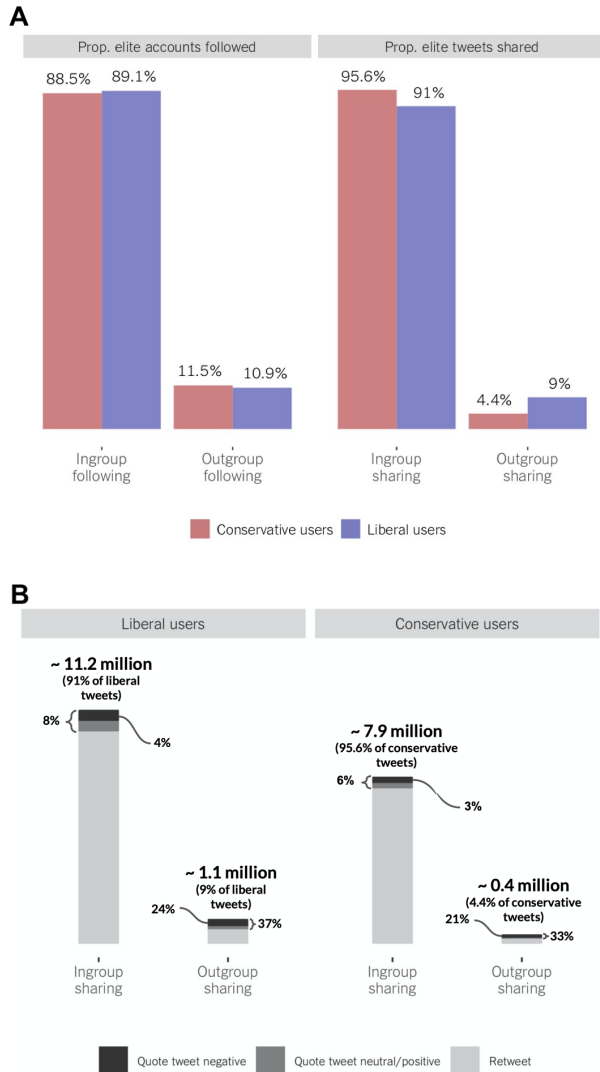


Figure 2: Amount of in-group (v. out-group) following and sharing. (A) Proportion (of all elite accounts users follow) that are in- *versus* out-group elite accounts; and Proportion (of all elite tweets users shared) that are from in- *versus* out-group elites (retweets and quote tweets pooled together). (B) How often ordinary users share (retweet *v.* quote tweet) messages from in-group *v.* out-group elites (providing further details about the right-panel in 2.A.)

Because – as aforementioned – following is rather “private,” we take the next step, asking whether users

actively share content from in-group vs. out-group elites (RQ3). We collect all tweets quoting (i.e., retweets with a comment) or retweeting (i.e., tweets shared without any commentary) messages sent by elite accounts, 20,731,455 messages shares total. We find strong evidence that it does. Of all elite tweets users share (with or without commentary), about 93% are from elites consistent with the users’ ideology. In general, for every out-group tweet a user shares, the user shares around 13 tweets from the in-group elite. Comparing Figure 2.A to 2.B suggests that this bias is greater for sharing than following: conditional on the number of in/out-group elite accounts followed, users have yet higher propensity to share in-group messages. The pattern is consistent across elite actors, although users are more likely to share tweets from in-group vs. out-group media and pundits than politicians (see [SI.5](#)). These biases are *not* a function of a few extreme users sharing information from few extreme elite sources: the levels of in-group sharing are also very high (although not as high) among users with low ideological extremity (see [SI.7](#)).

We find clear ideological asymmetries. Although in-group elite content represents similar proportion of all elite sharing by conservative (95.6%) and liberal (91%) users, the ratio of in- to out-group shares is dramatically different: 10:1 for liberals (11.2 mln in-group/1.1 mln out-group shares) and 20:1 for conservatives (7.9 mln in-group/0.4 mln out-group). As Figure 2.A shows, this cannot be explained by conservative users following fewer out-group elites.

The extent and nature of political biases further depends on the commentary added to the shared elite messages. Do users add negative comments when retweeting out-group messages (RQ4)? In order to establish the sentiment of the comments, we manually annotated a random sample of quotes for whether they were negative, neutral or positive toward the original tweet (see Materials and Methods for the details on the annotation and [SI.8](#) for examples of labeled quote tweets). We used those annotated data train a Convolutional Neural Net (CNN) to predict the tone of the remaining quotes (see [SI.8](#)). We emphasize that the findings are robust to using sentiment predictions generated by a Support Vector Machine and an Ensemble of several ngram-based models (see [SI.3](#)).

Two patterns emerge. First, users add comments at a higher rate when sharing out-group tweets (37% for liberals; 33% for conservatives) than when sharing in-group messages (8% and 6% respectively), suggesting that many users retweet out-group elites to express their stance, rather than endorse uncritically

(see Figure 2.B). Second, the machine learning sentiment predictions find that relative to all messages shared, users are more likely to add a negative comment to an out-group tweet (24% of all out-group shares for liberals; 21% for conservatives) than to an in-group tweet (4% of all in-group shares for liberals; 3% for conservatives). This translates into users adding negative commentary to an out-group elite tweet 6 times (liberals) and 7 times (conservatives) more often than to an in-group tweet. In short, on the rare occasions users share tweets from across the aisle, they do so to promote the in-group perspective: 63% of all quote tweets from out-group elites are shared with negative annotation.

We now offer more stringent evidence on these biases in engagement and on ideological differences by accounting for potential confounders (RQ5). First, we estimate a logistic regression model predicting the likelihood that users share in- vs. out-group elite tweets as a function of the ideology of the user (liberal vs. conservative), the type of the elite (politicians, pundits, news media), and the ideological extremity of the elite actor (i.e., folded continuous ideology score). To minimize the threat that our findings are driven by some users following many and others following few or no out-group elites, we control for the number of out-group accounts followed by each user (see [SI.1](#) for point estimates and 95% CIs). As Figure 3 shows, compared to conservative users, liberals are 39% *less* likely to share messages from in-group elites, even after controlling for these factors. We also find lower in-group sharing rates for politicians than pundits, indicating that tweets by out-group politicians are more likely to be shared (likely to be criticized, as we find above). Also, the more extreme the political elite actor, the more likely their tweets are shared by in-group users, suggesting the appeal of polarizing content and elites.

Second, to better understand the conditions in which users negatively comment on out-group messages and to explore partisan differences, we estimate a set of ordinal logistic regression models predicting the sentiment of the comments on the shared tweets as a function of whether the retweet was from an in- vs. out-group elite and all the covariates from the model above. Figure 4 reports the results of six models (see [SI.2](#) for the coefficient tables). Across all messages (top row), a tweet of an out- vs. in-group elite is substantially more likely to be shared with a negative vs. a neutral or positive comment, even after accounting for potential confounders. These patterns emerge among liberals and conservatives alike (2.18 times more likely for liberals, and 2.05 for conservatives). Crucially, elite type and the ideology of or-

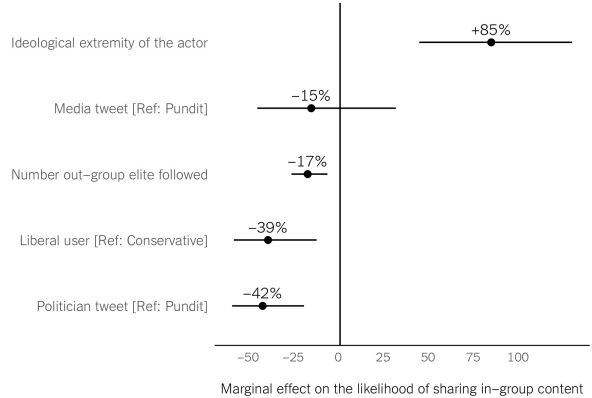


Figure 3: **Logistic regression predicting whether users share in- v. out-group elite tweets.** Marginal effects, on in-group sharing, of the elite account being a politician or media (v. pundit), the ideological extremity of the elite actor, the ideology of the users, and the number of out-group elite accounts a users follows. For continuous variables (ideological extremity and number of out-group elites followed) we report the marginal effect of a one standard deviation change.

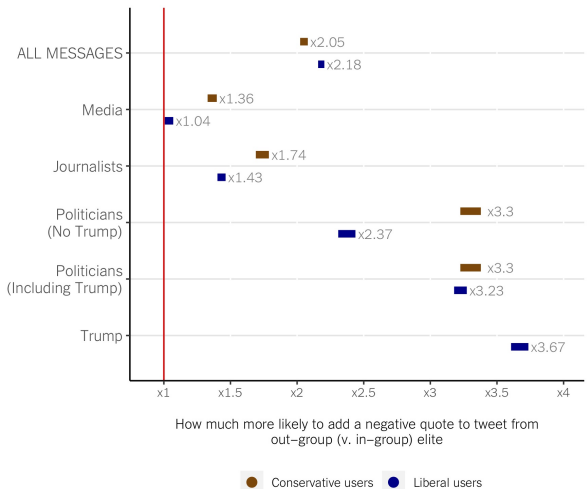


Figure 4: **Six multinomial models predicting the likelihood of attaching negative comments when sharing tweets from out-group (v. in-group) elite accounts.** Marginal effects for the elite tweet being from an out-group (rather than in-group) account.

dinary users matter (rows below). Tweets from out-group politicians are most likely to be shared with a negative commentary, followed by out-group pundits and media. Despite the aggregate similarity, conservatives are more likely to add negative comment to

retweets from out-group politicians (3.3 times more likely vs. 2.37), journalists (1.74 vs. 1.43), and especially out-partisan media (1.36; liberals are only 4% more likely to negatively comment on out- than on in-group media retweets). Quote tweets of Donald Trump are the reason for the aggregate similarity between liberals and conservatives. These tweets represent a large portion of the quote tweets in our data (around 20%) and liberal users overwhelmingly share them with negative comments. These models are robust to sentiment predictions from alternative machine learning models (see [SI.3](#)). Overall, although both groups are similarly biased in their following patterns, conservative users exhibit greater bias by sharing messages from in-group elites and also – except for Trump for liberals – have a higher propensity to negatively comment on the out-group messages they share.

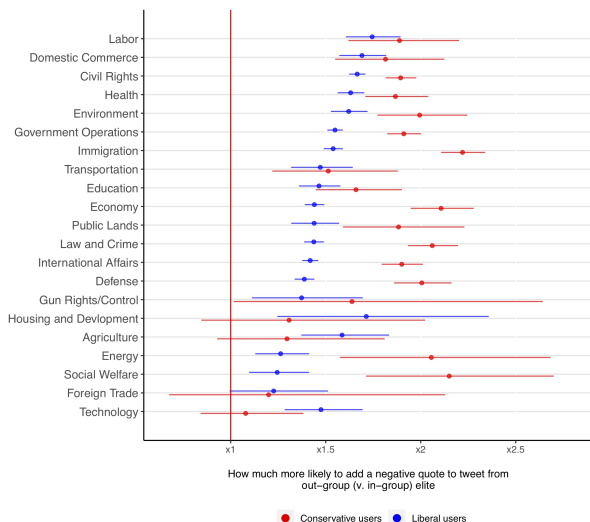


Figure 5: **Twenty-one multinomial models predicting the likelihood of attaching negative comments when sharing tweets from out-group (v. in-group) elite accounts.** For each of the policy areas discussed in the original elite tweets, we estimate the marginal effect of the tweet being from an out-group (rather than in-group) account.

Lastly, to assess whether these patterns depend on specific topics (e.g., when elites discuss hot-button issues vs. complex policies), we trained a Convolutional Neural Net (CNN) to predict the presence of topics from the Comparative Agendas Project in the tweets (see [SI.9](#)) and estimated multinomial models predicting the sentiment of the commentary on quote tweets about each topic, independently of the type of the elite. The political biases detected in our analyses

are reinforced with negative commentary on divisive issues in American politics, such as immigration or civil rights, more than on the technical ones, such as technology or foreign trade. Yet, with a couple isolated exceptions, regardless of the policy discussed, users from both ideologies are always more likely to add a negative comment to out-group rather than in-group messages (see Figure 5).

Discussion

Our findings offer comprehensive evidence on political biases in people’s engagement with politicians, pundits, and news organizations on social media. Given the different affordances of Twitter and the potentially distinct effects that various user behaviors have on the users themselves, their social networks, and the online public discourse, we attend to the following and the sharing of in- versus out-group elites as well as to the addition of positive or negative commentary to the shares. We asked whether a large random sample of users engages with in-group (rather than out-group) politicians, pundits, and news media on Twitter in ways that reinforce and exacerbate the feared insular communication patterns online. We offer three big take-aways.

First, most Twitter users do not follow or engage with any political elites online. This demonstrates a dichotomy between elite use of Twitter - politicians, pundits, and media (and also academics) – and mass use of Twitter. The elite discussion on the platform is important - but it is not necessarily observed directly by the masses. Given that Twitter users *are* more politically engaged than the general population to begin with ([37](#)), this finding of very low political elite following is surprising. In our case, 59.6% of a random sample of users (856,853 out of 1,437,774) were insufficiently politically interested to follow the accounts of the president, key senators, or major news media organizations. This bleak finding adds to some other evidence that many Twitter users do not follow news media ([16](#)) or Members of Congress ([13](#)). It also aligns with the aforementioned work showing low absolute levels of news consumption online ([47](#), [48](#)) and on social media more specifically ([44](#)), which users use primarily for entertainment ([37](#)). In our data, the following of celebrities is greater than that of any politician, pundit, or a news media organization: 70.7% of users follow at least one celebrity (athlete, musician, actor, etc.) compared to the 40.6% that follow at least one politically-relevant elite account (see [SI.12](#)).

Second, those who engage with political elites do

so in an overwhelmingly one-sided way, displaying clear political biases in their behaviors. Users disproportionately follow and disseminate messages by like-minded politicians, pundits, and news media, rarely following and yet more rarely sharing cross-cutting elites. We counter the hope that these biases are confined to a small group of extreme users: our patterns are robust (albeit naturally less pronounced) when examining all the users in our sample, including those ideologically moderate (see [SL.13](#)), and are not driven by a few extreme users (see [SL.7](#)). In addition, users are not only more likely to add a commentary to the out-group content they (rarely) share (i.e., quote tweets) but also add negative commentary to these shares. The negative sentiment of the commentary added to *out-group* retweets works to reinforce the ideological bubble. In sum, across the approximately 20 million shares of elite content we analyzed, only 5% were of out-group elites *without* any negative commentary. Analyzing following, sharing, and commenting, an understudied feature of social media platforms, is one of the ways this project extends past work and shows that Twitter users do exhibit strong and previously unaddressed political biases when engaging with politicians, pundits, and news media on Twitter.

Although some work suggests that people’s information diets partly overlap, online ([48](#)) and on social media ([16](#)), the evidence presented above suggests there is little overlap between conservatives and liberals when interacting with political elites on Twitter. To speak more directly to the past studies, we checked this overlap in the following, sharing, and annotating of elite accounts (see [SL.11](#)). Conservatives and liberals rarely follow elite accounts of the *opposing* ideology even though both groups jointly follow moderate elite accounts to some extent (about 35% of all elite accounts followed by liberal users are moderate accounts, and about 20% for conservatives), denoting some overlap. This overlap is smaller for users with a more extreme ideology, consistent with the idea that it is precisely the strong partisans who are most likely to engage in the most politically biased behaviors on social media. In short, the overlap between between conservatives and liberals is mostly confined to moderate users following moderate politicians, pundits, and news media. Tellingly, these overlaps gradually decrease the more “public” the analyzed behavior becomes. That is, compared to the amount of moderate elite accounts liberal and conservatives follow, retweets of moderate accounts represent a smaller proportion of their elite shares (about 30% for liberals, 6% for conservatives). In addition, although liberals are more likely to add a

positive or neutral comment to tweets from moderate elite accounts, conservatives are more likely to add a negative commentary, emphasizing an even smaller overlap when it comes to commenting on tweets from moderate accounts.

Given these patterns of increasing political biases (and decreasing overlaps) in more “public” behaviors on social media, we speculate that these biases are partly due to *perceived* polarization of one’s peer group. If citizens see others – be it the general public, their social network, or the imagined audiences of their tweets – as more polarized than they actually are ([53](#)), users may never share content from the other side, even if they are sympathetic to that content. In other words, social pressure to conform to (perceived) dominant group opinions may lead citizens to engage in “performative” sharing and commenting, which may explain and further exacerbate the detected biases. Future work should systematically test this idea and attend to disassociations between reading, following, sharing, and commenting, as each behavior entails different costs, sends different signals to one’s network, is subject to different pressures, and – as such – generates distinct political biases.

The third key finding regards ideological asymmetries – a key area for this research ([54](#)). Both conservative and liberal users are much more likely to follow in- vs. out-group elites, and both groups do so at similar rates. Also, although both groups are disproportionately more likely to retweet in-group than out-group elites, conservatives engage in cross-ideological diffusion substantially *less*. Also, apart from tweets from Donald Trump, conservatives tend to annotate out-group tweets with negative commentary more often than liberals do.

These asymmetries, consistent with prior research on political biases in users’ behaviors on social media ([12](#), [55](#), [18](#)), can be due to two interrelated factors. The work on distinct cognitive styles of political ideologues suggests that conservatives manifest cognitive styles such as dogmatism, rigidity, or uncertainty avoidance ([56](#)), which might predispose conservatives to shield away from and be more negative toward cross-cutting views (see ([12](#), [57](#))). In addition, these asymmetries can be due to a broader social network ecosystem, such as the actions of friends and followers in right-leaning groups and also conservative users following and being targeted by more inauthentic accounts ([17](#)). Again, we encourage researchers to systematically attend to these differences and their underlying reasons.

These findings, although important, naturally do not offer a complete picture of political biases in all

various information and communication behaviors on social media. For one, we do not examine interactions among ordinary users. Our theoretical and practical focus was on political elites, who dominate political discussions on platforms, have disproportionate influence on public, media, and policy agenda (19), and can further exacerbate – or mitigate – polarization. In addition, we do not analyze replies to elite tweets, focusing on retweeting and quote tweeting, behaviors that are more visible to one’s own followers than replies. Accounting for whether users replied to elite messages, and – if so – for the tone of the reply, would have offered a more complete portrayal of user engagement with elites on Twitter.

In a related vein, we cannot capture the content merely seen by the users in our sample, instead focusing on the more “active” behaviors of following, sharing (retweeting), and annotating (quote tweeting). Our focus may be underestimating the number of people exposed to elite messages. That is, some disengaged users may be exposed to the studied elites indirectly, through the retweets of their more engaged friends. Inasmuch as conservative/liberal users follow other conservative/liberal users ((58) see also (17)), this indirect exposure would be mostly to in-group not out-group elites, thus introducing additional political biases on social media. This indirect exposure, moreover, could further exacerbate perceived or false polarization (53). The most partisan users are most likely to share political elites *and* the extreme elites are most likely to be shared (59). As a result, apolitical users who do not themselves follow any elites would encounter content that is hyper-partisan. Furthermore, if out-group content is retweeted with an added commentary, those apolitical users would see messages that derogate the other side. This could create the perception that politics is divisive and polarized, further disengaging some citizens from the political process (60, 61). Testing these indirect, inadvertent exposures to elite communications and analyzing their effects on the users is a worthwhile direction for future work.

We also encourage researchers to extend our work to a local level of Twitter discussions. In this project, we offered foundational overview of political biases in following, sharing, and annotating on a large national scale, looking at the most powerful and – therefore potentially most frequently followed – political elites (52). Yet, some recent work suggests that contentious political issues are also discussed at state-level and that patterns of media use, political talk, and policy attitudes differ between localities within states (62, 63). Future work should test whether similar biases in the following, retweeting, and quote

tweeting of in- versus out-group politicians and news media organizations emerge on local levels. We also hope that scholars will extend our approach to international contexts to examine whether the U.S. political ecosystem online is unique or whether Twitter users follow, share, and annotate in similar ways in multi-party and potentially less polarized systems.

Our findings have important implications for research and democracy at large. Despite the hopes that social media would reinvigorate American democracy by lowering information costs and access barriers and directly connecting representatives with their constituencies, most citizens do not engage with politicians, pundits, and news organizations on social media platforms. The unprecedented choice in the online environment reinforces the divide between the politically withdrawn and more politically active citizens. This latter group can now easily selectively engage with like-minded sources and information, disparage the out-group and its messages, and become yet more polarized.

When we witness a growing radicalization of certain groups in American society (and globally), decreasing support for democratic norms, and rising support for political violence (64), concerns about political biases in online behaviors are ever more pressing, no matter how small the groups engaging in those behaviors may be. Because these small groups are disproportionately more vocal, participatory, and used by mainstream news media to represent public opinion (65), they amplify the general public perception of ideological extremity, political biases, and unprecedented polarization. And yet, scholars and public observers need to keep in mind that these political biases are removed from the everyday information and communication ecosystem of most American citizens and that pulling these less engaged and more moderate citizens back into the democratic process may decrease political polarization online and offline.

Materials and Methods

Data Collection. The data were collected by the Center for Social Media Politics at New York University in between 2016 and 2019. We generated a random sample of Twitter users in two ways. First, (a) before Twitter switched to 64-bit IDs in 2016, we automatically generated 32-bit random numbers and checked for whether they were existing Twitter users (first about 100,000 users). After the introduction of the longer Twitter IDs, (b) we increased the size of the sample: we collected tweets mentioning a set of English stopwords (i.e. “the”) for numer-

ous short amounts of time selected at random and then pulled the authors of those tweets and information about how frequently they tweeted, and we subsampled a set of authors with a tweeting distribution similar to that of the users in the list created using the first approach (a). For the second approach (b), we selected times in the day when users from other English-speaking countries were the least likely to tweet so that to assure that the tracked users are in the U.S. We do not impose other additional geographic restrictions, but given the U.S.-centric character of the English-speaking Twitter at the time we collected the data and the fact that we randomly sampled users messaging at times when users from other English-speaking countries are least likely to tweet, we expect most of these users to be located in the U.S. In fact, our findings hold when only looking at a set of 24,328 users we have confidently located in the U.S. (using the method described in SI B in (19)) as shown in SI.10.

In sum, we tracked the following, tweeting, retweeting, and quote tweeting activity of a random sample of 1,437,774 Twitter users for four years (2016-2019) by regularly pulling their timelines using the Twitter REST API. Specifically, we examined whether the users followed an extensive set of 2,624 political elite accounts (489 politicians, 2,016 pundits, and 119 media organizations), if they have retweeted messages from these accounts, and have used quote tweets to add comments to the retweeted messages. This full sample is used in our initial descriptive analyses addressing RQ1.

To address our subsequent RQs that focus on ideological in-groups and out-groups, we apply the ideology classification developed by Barbera and colleagues (12). We note that this classification has been extensively validated using external indicators on the aggregate and individual levels ((12) SI Figure 2, 4, 5). We do not train a new model from scratch, and use the model in (12) to generate ideology score for the users and actors in our sample. Because the list of elite accounts used by (12) to generate ideology estimation for ordinary users was slightly more restrictive than ours, there are some users who do not follow enough elites in the original list to estimate their ideology. Ultimately, as aforementioned in the main text, we were able to classify the ideology of 180,203 users in our sample.

In the analyses examining whether users actively share content from in-group vs. out-group elites (RQ3), we include those users who retweeted or quote tweeted messages sent by the elite accounts on our list. In particular, we use 20,731,455 shares of elite accounts classified as liberal or conservative (moder-

ates excluded), that were shared by 151,063 politically active users classified as liberal or conservative (moderates excluded as well). We included any quote tweet, even a quote tweet of someone a user does not follow her/himself.

Lastly, to address RQ4 regarding the sentiment (positive, neutral, or negative) of the commentary added to the shared elite tweets (i.e., quote tweets), we needed to remove quote tweets too short for sentiment predictions (< 5 words, after pre-processing the text). The final sample for these analyses is 1,469,708 tweets sent by 85,849 users (about 57% of the 151,063 politically active users classified as liberals or conservatives) quoting 1,668 elites: 563,689 (38%) tweets quoting 402 politicians, 391,433 (27%) quoting 78 media organizations, and 514,586 (35%) quoting 1,188 pundits. In general, 668,248 (46%) quoted conservative actors and 801,460 (54%) quoted liberal actors.

Classifiers. Sentiment Classifier. To determine the sentiment of the commentary added to the shared elite messages, whether positive, neutral, or negative, we trained a Convolutional Neural Net (CNN) classifier predicting the sentiment of the quote tweets. First, we randomly sampled 8,351 tweets from our full dataset of quote tweets. Four trained research assistants manually coded them for whether the quote was negative, neutral or positive toward the message and/or the political actor, independently of the tone of the original message (Krippendorff’s $\alpha = .816$). We used those annotated data train the following 5 types of machine learning models predicting whether the commentaries were positive, neutral, or negative (multi-class models): (a) a Decision tree (TREE), (b) a K-neighbors model, (c) a Support vector machine, (d) a majority-based ensemble model that took into account the output of the three previous ones, and (e) a four-layer Convolutional Neural Net (CNN). For training (a), (b), (c), and (d), we transformed all text to lowercase, removed stopwords, and lemmatized the remaining tokens to finally create a TF-IDF matrix that we used as model input. For the CNN model, we transformed all text to lower case and used 300-dimension GloVe embeddings as inputs. We tested the accuracy of each algorithm using 5-fold cross-validation and a 80/20 train-test split. As seen in SI.8, Figure S5, the CNN proved to be the most accurate of the five classifiers (see colored SI.8 for additional details and validations. Given the superior performance, we use the CNN classifier to predict the tone of all quote tweets in our dataset. We also emphasize that the findings are robust to using sentiment predictions generated by a Support Vector Machine and an Ensemble of several ngram-based models (see SI.3).

Topic classifier. In addition, we automatically classified the topic of the content of the original tweets from elite accounts that we study. Given the large number of tweets, manual coding was not practical for the full corpus. To reliably and at-scale predict the topic of the original elite tweets, we trained a Convolutional Neural Net (CNN) predicting whether each tweet discussed one of the 20 topics of the Comparative Agendas Project (CAP) (66). In [SI.9](#) we provide detailed information about the model architecture, how it was trained, as well as its performance.

Acknowledgements

The Center for Social Media and Politics at New York University is supported by funding from the John S. and James L. Knight Foundation, the Charles Koch Foundation, Craig Newmark Philanthropies, the William and Flora Hewlett Foundation, the Siegel Family Endowment, the Bill and Melinda Gates Foundation, and the National Science Foundation. **Author Contributions:** M.W., X.Y. and A.C. designed research; J.T. and J.N. oversaw the data collection, A.C. contributed the classifiers/analytic tools; A.C. and X.Y. analyzed data; and M.W. and A.C. wrote the paper. M.W. and A.C. contributed equally to the paper. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper, and the replication code, are available in a public GitHub repository (https://github.com/CasAndreu/ingroup_filtering) as well as on Harvard Dataverse (<https://doi.org/10.7910/DVN/G5APVR>).

References

- [1] Richard Fletcher and Rasmus Kleis Nielsen. Are people incidentally exposed to news on social media? A comparative analysis. *New media & society*, 20(7):2450–2468, 2018.
- [2] A.W. Geiger. Key findings about the online news landscape in America, September 2019.
- [3] Eli Pariser. *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. Penguin Books, reprint edition, April 2012.
- [4] Cass R. Sunstein. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press, Princeton, N.J. Oxford, updated edition, April 2018.
- [5] Eli J. Finkel, Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Klar, Lilliana Mason, Mary C. McGrath, Brendan Nyhan, David G. Rand, Linda J. Skitka, Joshua A. Tucker, Jay J. Van Bavel, Cynthia S. Wang, and James N. Druckman. Political sectarianism in America. *Science*, 370(6516):533–536, October 2020. Publisher: American Association for the Advancement of Science Section: Policy Forum.
- [6] Marc J. Hetherington and Thomas J. Rudolph. *Why Washington Won't Work: Polarization, Political Trust, and the Governing Crisis*. University of Chicago Press, illustrated edition edition, September 2015.
- [7] Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. The origins and consequences of affective polarization in the united states. *Annual Review of Political Science*, 22:129–146, 2019.
- [8] Nathaniel Persily and Joshua A Tucker, editors. *Social Media and Democracy: The State of the Field, Prospects for Reform*. Cambridge University Press, 2020.
- [9] Pablo Barberá. Social media, echo chambers, and political polarization. In Nathaniel Persily and Joshua A. Tucker, editors, *Social Media and Democracy: The State of the Field, Prospects for Reform*. Cambridge University Press, 2020.
- [10] L. A. Adamic and N. Glance. The political blogosphere and the 2004 US election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery*, pages 36 – 43, New York, 2005. ACM.
- [11] Eric Lawrence, John Sides, and Henry Farrell. Self-segregation or deliberation? blog readership, participation, and polarization in american politics. *Perspectives on Politics*, 8(1):141–157, 2010.
- [12] Pablo Barberá, John T. Jost, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau. Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science*, August 2015. Publisher: SAGE PublicationsSage CA: Los Angeles, CA.
- [13] Gonzalo Rivero. Preaching to the choir: ideology and following behaviour in social media.

- Contemporary Social Science*, 14(1):54–70, Jan 2019.
- [14] Elanor Colleoni, Alessandro Rozza, and Adam Arvidsson. Echo chamber or public sphere? predicting political orientation and measuring political homophily in twitter using big data. *Journal of Communication*, 64(2):317–332, Apr 2014.
 - [15] Antoine Boutet, Hyounghick Kim, and Eiko Yoneki. What’s in your tweets? i know who you supported in the uk 2010 general election. *Proceedings of the International AAAI Conference on Web and Social Media*, 6(11):411–414, 2012.
 - [16] Gregory Eady, Jonathan Nagler, Andy Guess, Jan Zilinsky, and Joshua A. Tucker. How Many People Live in Political Bubbles on Social Media? Evidence From Linked Survey and Twitter Data. *SAGE Open*, 9(1):2158244019832705, January 2019. Publisher: SAGE Publications.
 - [17] Wen Chen, Diogo Pacheco, Kai-Cheng Yang, and Filippo Menczer. Neutral bots probe political bias on social media. *Nature Communications*, 12(1):5580, Dec 2021.
 - [18] Michael D. Conover, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. Partisan asymmetries in online political activity. *EPJ Data Science*, 1(11):1–19, Dec 2012.
 - [19] Pablo Barberá, Andreu Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost, and Joshua A. Tucker. Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data. *American Political Science Review*, 113(4):883–901, November 2019. Publisher: Cambridge University Press.
 - [20] Galen Stocking, Michael Barthel, and Elizabeth Grieco. *Sources Shared on Twitter: A Case Study on Immigration*. Jan 2018.
 - [21] Patrick W Kraft, Yanna Krupnikov, Kerri Milita, John Barry Ryan, and Stuart Soroka. Social media and the changing information environment: Sentiment differences in read versus recirculated news content. *Public Opinion Quarterly*, 84:195–215, Mar 2020.
 - [22] James N. Druckman, Erik Peterson, and Rune Slothuus. How Elite Partisan Polarization Affects Public Opinion Formation. *The American Political Science Review*, 107(1):57–79, 2013. Publisher: American Political Science Association, Cambridge University Press.
 - [23] Kevin J. Mullinix. Partisanship and Preference Formation: Competing Motivations, Elite Polarization, and Issue Importance. *Polit Behav*, 38(2):383–411, June 2016.
 - [24] Stephen P. Nicholson. Polarizing Cues. *American Journal of Political Science*, 56(1), 2012.
 - [25] Matthew Levendusky. *How Partisan Media Polarize America*. University of Chicago Press, Chicago; London, illustrated edition, September 2013.
 - [26] Matthew Levendusky and Neil Malhotra. Does Media Coverage of Partisan Polarization Affect Political Attitudes? *Political Communication*, 33(2):283–301, April 2016. Publisher: Routledge eprint: <https://doi.org/10.1080/10584609.2015.1038455>.
 - [27] Panagiotis Metaxas, Eni Mustafaraj, Kily Wong, Laura Zeng, Megan O’Keefe, and Samantha Finn. What do retweets indicate? results from user survey and meta-review of research. *Proceedings of the International AAAI Conference on Web and Social Media*, 9(11):658–661, 2015.
 - [28] Leonie Huddy. Contrasting theoretical approaches to intergroup relations. *Political Psychology*, 25(6):947–967, 2004.
 - [29] Henri Tajfel. Social Psychology of Intergroup Relations. *Annual Review of Psychology*, 33(1):1–39, 1982.
 - [30] Henri Tajfel and John Turner. An Intergrative Theory of Intergroup Conflict. In Mary Jo Hatch and Majken Schultz, editors, *Organization Identity: A Reader*, pages 56 – 66. Oxford University Press, 1979.
 - [31] Marilynn B. Brewer. Ingroup Identification and Intergroup Conflict. In Richard D. Ashmore, Lee Jussim, and David Wilder, editors, *Social Identity, Intergroup Conflict, and Conflict Reduction*, pages 17 – 41. Oxford University Press, 2001.
 - [32] Matthew A. Baum and Tim Groeling. Shot by the Messenger: Partisan Cues and Public Opinion Regarding National Security and War. *Polit Behav*, 31(2):157–186, June 2009.
 - [33] Riccardo Puglisi and James M. Snyder. Newspaper Coverage of Political Scandals. *The Journal*

- of *Politics*, 73(3):931–950, July 2011. Publisher: The University of Chicago Press.
- [34] Jeffrey M. Berry and Sarah Sobieraj. *The Outrage Industry: Political Opinion Media and the New Incivility*. Oxford University Press, December 2013.
 - [35] Albert Feller, Matthias Kuhnert, Timm Sprenger, and Isabell Welp. Divided they tweet: The network structure of political microbloggers and discussion topics. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(11):474–477, 2011.
 - [36] Elanor Colleoni, Alessandro Rozza, and Adam Arvidsson. Echo chamber or public sphere? predicting political orientation and measuring political homophily in twitter using big data: Political homophily on twitter. *Journal of Communication*, 64(2):317–332, Apr 2014.
 - [37] Colleen McClain, Regina Widjaya, Gonzalo Rivero, and Aaron Smith. The behaviors and attitudes of u.s. adults on twitter, Nov 2021.
 - [38] Lilliana Mason. *Uncivil Agreement: How Politics Became Our Identity*. University of Chicago Press, Apr 2018. Google-Books-ID: R29RDwAAQBAJ.
 - [39] Charles S. Taber and Milton Lodge. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769, 2006.
 - [40] Mathias Osmundsen, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann, and Michael Bang Petersen. Partisan polarization is the primary psychological motivation behind political fake news sharing on twitter. *American Political Science Review*, 115(3):999–1015, Aug 2021.
 - [41] Steve Rathje, Jay J. Van Bavel, and Sander van der Linden. Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences*, 118(26):1–9, Jun 2021.
 - [42] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239):1130–1132, June 2015. Publisher: American Association for the Advancement of Science Section: Report.
 - [43] Nathan Bartley, Andres Abeliuk, Emilio Ferrara, and Kristina Lerman. Auditing algorithmic bias on twitter. In *13th ACM Web Science Conference 2021*, page 65–73. ACM, Jun 2021.
 - [44] Mark Zuckerberg. Facebook, January 2018. <https://www.facebook.com/zuck/posts/10104445245963251>.
 - [45] Chris Wells and Kjerstin Thorson. Combining big data and survey techniques to model effects of political content flows in facebook. *Social Science Computer Review*, 35(1):33–52, Feb 2017.
 - [46] Seth Flaxman, Sharad Goel, and Justin M. Rao. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1):298–320, Jan 2016.
 - [47] Magdalena Wojcieszak, Sjifra de Leeuw, Ericka Menchen-Trevino, Seungsu Lee, Ke M. Huang-Isherwood, and Brian Weeks. No polarization from partisan news: Over-time evidence from trace data. *The International Journal of Press/Politics*, page 19401612211047194, Nov 2021.
 - [48] Andrew Guess. (Almost) Everything in Moderation: New Evidence on Americans’ Online Media Diets. *American Journal of Political Science*, 65(4):1007–1022, 2021.
 - [49] Jennifer Allen, Baird Howland, Markus Mobius, David Rothschild, and Duncan J. Watts. Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14):eaay3539.
 - [50] Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. Fake news on twitter during the 2016 U.S. presidential election. *Science*, 363(6425):374–378, Jan 2019.
 - [51] Yoon Kim. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar, October 2014. Association for Computational Linguistics.
 - [52] Daniel J. Hopkins. *The Increasingly United States: How and Why American Political Behavior Nationalized*. University of Chicago Press, May 2018. Google-Books-ID: ZYIZDwAAQBAJ.
 - [53] Philip M. Fernbach and Leaf Van Boven. False polarization: Cognitive mechanisms and potential solutions. *Current Opinion in Psychology*, 43:1–6, Feb 2022.

- [54] Joshua A. Tucker, Andrew Guess, Pablo Barbera, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature. SSRN Scholarly Paper ID 3144139, Social Science Research Network, Rochester, NY, March 2018.
- [55] William J. Brady, Julian A. Wills, Dominic Burkart, John T. Jost, and Jay J. Van Bavel. An ideological asymmetry in the diffusion of moralized content on social media among political leaders. *Journal of Experimental Psychology: General*, 148(10):1802–1813, 2019.
- [56] John T. Jost, Jack Glaser, Arie W. Kruglanski, and Frank J. Sulloway. Political conservatism as motivated social cognition. *Psychological Bulletin*, 129(3):339–375, 2003.
- [57] Christopher A. Bail, Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221, 2018.
- [58] Mohsen Mosleh, Cameron Martel, Dean Eckles, and David G. Rand. Shared partisanship dramatically increases social tie formation in a twitter field experiment. *Proceedings of the National Academy of Sciences*, 118(7):e2022761118, Feb 2021.
- [59] Michael W. Wagner and Mike Gruszczynski. Who gets covered? ideological extremity and news coverage of members of the u.s. congress, 1993 to 2013. *Journalism Mass Communication Quarterly*, 95(3):670–690, Sep 2018.
- [60] Samara Klar, Yanna Krupnikov, and John Barry Ryan. Affective polarization or partisan disdain?: Untangling a dislike for the opposing party from a dislike of partisanship. *Public Opinion Quarterly*, 82(2):379–390, Jun 2018.
- [61] Samara Klar and Yanna Krupnikov. *Independent Politics*. Cambridge University Press, Jan 2016. Google-Books-ID: KNA9CwAAQBAJ.
- [62] Chris Wells, Lewis A. Friedland, Ceri Hughes, Dhavan V. Shah, Jiyoun Suk, and Michael W. Wagner. News media use, talk networks, and anti-elitism across geographic location: Evidence from wisconsin. *The International Journal of Press/Politics*, 26(2):438–463, Apr 2021.
- [63] Lewis A. Friedland, Dhavan V. Shah, Michael W. Wagner, Katherine J. Cramer, Chris Wells, and Jon Pevehouse. Battleground: Asymmetric communication ecologies and the erosion of civil society in wisconsin. *Elements in Politics and Communication*, Feb 2022.
- [64] Nathan P. Kalmoe and Lilliana Mason. Lethal mass partisanship:. page 41.
- [65] Shannon C. McGregor and Logan Molyneux. Twitter’s influence on news judgment: An experiment among journalists. *Journalism*, 21(5):597–613, May 2020. Publisher: SAGE Publications.
- [66] Frank R. Baumgartner and Bryan D. Jones. *Agendas and instability in American politics*. University of Chicago Press, 2010.
- [67] Zhanna Terechshenko, Fridolin Linder, Vishakh Padmakumar, Fengyuan Liu, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau. A comparison of methods in political science text classification: Transfer learning language models for politics. *Presented at the XXXVII PolMeth Annual Meeting*, 2020.
- [68] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, 2013.
- [69] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [70] Annelise Russell. Us senators on twitter: Asymmetric party rhetoric in 140 characters. *American Politics Research*, 46(4):695–723, 2018.
- [71] Jan Zilinsky, Cristian Vaccari, Jonathan Nagler, and Joshua A. Tucker. Don’t republicans tweet too? using twitter to assess the consequences of political endorsements by celebrities. *Perspectives on Politics*, 18(1):144–160, 2020.