Evidence of Demographic rather than Ideological Segregation in News Discussion on Reddit

Corrado Monti
corrado.monti@centai.eu
CENTAI
Turin, Italy

Jacopo D'Ignazi
jacopo.dignazi@isi.it
ISI Foundation
Turin, Italy

Michele Starnini
michele.starnini@centai.eu
CENTAI
Turin, Italy

Gianmarco De Francisci Morales
gdfm@acm.org
CENTAI
Turin, Italy

ABSTRACT

We evaluate homophily and heterophily among ideological and demographic groups in a typical opinion formation context: online discussions of current news. We analyze user interactions across five years in the r/news community on Reddit, one of the most visited websites in the United States. Then, we estimate demographic and ideological attributes of these users. Thanks to a comparison with a carefully-crafted network null model, we establish which pairs of attributes foster interactions and which ones inhibit them.

Individuals prefer to engage with the opposite ideological side, which contradicts the echo chamber narrative. Instead, demographic groups are homophilic, as individuals tend to interact within their own group—even in an online setting where such attributes are not directly observable. In particular, we observe age and income segregation consistently across years: users tend to avoid interactions when belonging to different groups. These results persist after controlling for the degree of interest by each demographic group in different news topics. Our findings align with the theory that affective polarization—the difficulty in socializing across political boundaries—is more connected with an increasingly divided society, rather than ideological echo chambers on social media. We publicly release our anonymized data set and all the code to reproduce our results.[1]

CCS CONCEPTS

• **Information systems** → **World Wide Web**; • **Applied computing** → **Sociology**; • **Human-centered computing** → *Social networking sites*.

KEYWORDS

Homophily, Socio-demographic factors, Polarization, News, Reddit

---

[1]https://github.com/corradomonti/demographic-homophily

1 INTRODUCTION

Socializing across ideological boundaries has been reported to be increasingly difficult [35]. This phenomenon, known as *affective polarization*, has been widely observed in the United States. For instance, the fraction of Republicans and Democrats that would be unhappy if their child married a supporter of the opposing party has increased from 5% in 1960, to more than 30% in 2010 [36]. This kind of polarization contributes to creating a harsh political climate, where citizens believe that other people's opinions are irrational and not based on facts [38]. It has also been associated with the spread of harmful conspiracy theories, which can lead to dehumanization [32], culminating in threats such as the emergence of far-right domestic terrorism [45, 50].

The causes of affective polarization are widely debated, and some studies have pointed at social media as a decisive factor [29]. In this popular view [57], social media—through recommendation algorithms, algorithmic filter bubbles, and self-sorting—would reinforce ideological separation, thus making it harder for an individual to engage with information from the opposite side. While several studies have observed ideological homophily and echo chambers in some social media [15, 16, 18, 27], in others interactions between opposing sides are the norm and not the exception [15, 20, 21, 69]. Some scholars have argued that online media has actually increased chances of engaging with the opposite ideological side [24], and Internet usage explains only a very small share of polarization [11].

An alternative explanation for the increase in affective polarization lies in the role of demographic factors [12, 46, 61]. Underlying material rifts and inequalities in society make gender, race, age, and wealth boundaries increasingly divisive, and also increasingly correlated to party affiliation and ideological standpoints. In their seminal paper, Lazarsfeld et al. [41] distinguished between status homophily, when similarity in social status fosters connections, and value homophily, when similarity in beliefs and ideas tends to do so. In this view, the existing status homophily within these groups would therefore be one of the main drivers behind value homophily and thus contribute to affective polarization [48].

To assess these competing hypotheses, this study focuses on the interactions between different demographic groups in one of the main loci where opinions are formed and challenged: the discussion of news [40]. In particular, the discussion of news on social media is of twofold interest. First, it allows studying how specific demographic groups are connected in a context where they cannot directly observe each other's demographic conditions: any measured effect is entirely due to differences in worldviews, interests, and thoughts expressed through writing. Second, it allows us to measure and compare, at the same time and in the same context, the likelihood of interaction between opposing ideological sides. These reasons lead us to investigate Reddit, one of the most visited websites in the U.S., a public forum widely used to read and discuss news. We focus on the r/news subreddit, the main Reddit community specifically dedicated to discussion of current news, with a focus on U.S. internal affairs.

Therefore, our main research question is *whether social interactions in online news discussions on Reddit tend to be segregated by demographic boundaries, or by ideological echo chambers.*

To this aim, we first distinguish groups of users more likely to be part of a certain demographic group thanks to the method proposed by Waller and Anderson [70]. Then, we measure how much the amount of interactions observed between demographic groups differs from a carefully-crafted null model, where such attributes would not affect interactions. We replicate this analysis separately for five years, from 2016 to 2020.

In all these experiments, we find similar effects. First, that ideological echo chambers are absent: left-wing and right-wing users interact *more* than expected by a null model of random interactions, while within-group interactions are less likely to happen. This result suggests that ideological echo chambers do not necessarily appear in a typical context of opinion formation online. Second, interactions in news discussions on Reddit are unlikely to cross demographic boundaries, e.g., affluent users are more likely to interact among themselves than with less-affluent users. In other words, we observe *segregation* along demographic boundaries, in the sense of "restriction of contacts between various groups" [25], in line with previous literature on social networks [34]. Finally, we demonstrate that demographic homophily is not simply driven by common interests, such as discussions on the same topics. While indeed different demographic groups have different interests in the various news topics, their segregation goes beyond this simple explanation. Most importantly, this segregation happens without users being able to directly observe the demographic attributes, and without the mediating effects of, e.g., geographic segregation.

Our results, therefore, support the idea that not only echo chambers might not be an intrinsic characteristic of social media [15], but that the separation of demographic groups is observable during opinion formation online: fundamental societal divisions might be reflected also in social media.

2 BACKGROUND AND RELATED WORK

The polarization of opinions, with respect to politics and controversial topics in general [60], has been reported to be on the rise in recent years [39]. Online social media, where billions of people express their opinions daily, offer a scalable and inexpensive way to quantify this phenomenon. The literature on this topic is too vast to review here, so we briefly touch on a few landmark studies. In a pioneer study, Conover et al. [16] showed that on Twitter, a network of retweets about the 2010 US midterm elections exhibits a highly segregated partisan structure, with extremely limited connectivity between left- and right-leaning users. Since then, their result has been replicated several times, on different topics, and in various contexts [17, 19, 26, 28, 55]. Opinion polarization has also been detected on Facebook and YouTube, which may foster selective exposure of users toward content they like, thus limiting the exposure to diverse content [3, 4, 7, 32, 62, 66]. Echo chambers, situations where users have their beliefs reinforced due to repeated interactions with like-minded peers, have been proposed as one of the determinants for polarization on social media, and as a threat to democracy [68].

The investigation around the presence of echo chambers on social media has recently spurred a rich literature and a vivid debate among researchers. On Twitter, the content consumption and production of users reveal a similar political leaning, so that users are generally exposed to political opinions that agree with their own [27]. On the same platform, the presence of echo chambers in the debate about the impeachment of former Brazilian President Dilma Rousseff has been shown to alter the information diffusion between the two sides, supporters and opponents of the impeachment [18]. On Facebook, Quattrociocchi et al. [62] showed that users form polarized groups that share two distinct narratives, related to conspiracy theories and science, respectively. On the same platform, the content consumption about vaccines is skewed by echo chambers, while opinion polarization increased over the years [66]. Finally, Cinelli et al. [15] compared the presence of echo chambers around several controversial topics across four social media platforms (Twitter, Facebook, Reddit, and Gab), highlighting that Facebook shows higher segregation of news consumption than Reddit. Interestingly, several works challenge the harmfulness of the phenomenon [67], or the very existence of echo chambers in different contexts [21, 23, 30]. For instance, on Reddit, Trump and Clinton supporters show a preference for cross-cutting political interactions between the two communities rather than within-group interactions, thus contradicting the echo chamber narrative [20]. The current work finds a similar result in a different context (news discussion), and at the same time finds evidence for demographic segregation on social media.

Since online discussion can be a driving factor of opinion formation [40], the search for the causes of polarization has ultimately led to questions being asked on how people interact on the base of their socio-demographic extraction, culture, and ideology. Demographic segregation in physical spaces has been the subject of study for several decades [65], and has been shown to have increased in modern times [44], mostly due to racial sorting. The idea that this sorting is the main cause for increased political polarization has been advanced [8, 43], although it has also been criticized [1]. In this context, affective polarization has been introduced as a measure of the attitude of individuals with opposite ideologies toward each other [36]. Iyengar and Westwood [37] have shown how partisan-wise polarization also leads to increasing animosity (and even loathing) across party lines. Indeed, the social identities of people have grown increasingly aligned with their

partisan identity [35, 46]. Nonetheless, recent research has found how affective polarization can be influenced by factors such as ethnicity [12] and age [61], thus suggesting that socio-demographic and cultural differences might be an important driving force in social interactions. This idea is corroborated by observations of how social interactions are strongly homophilic with respect to socio-demographic extraction [41, 48]. All these findings suggest that more attention is needed to the role of socio-demographic and cultural features, as pivotal factors in influencing individual social choices, and consequently opinion formation and polarization.

To address this gap, we employ graph models that allow studying the interplay between node features and link formation. Our choice of model is driven by a few constraints. For instance, the well-known stochastic block model [56], in its typical form, only allows a node to belong to exactly one class, while we are interested in modeling the presence of more features on nodes at the same time. Also, we need a model able to capture feature-feature relationships that are not limited to homophily. Thus, we opt for the feature-feature model initially proposed by Miller et al. [51] as a way to learn features from data, and then further studied as a generative network model [9], and as a tool to infer feature-feature relations [10, 54]. In the present work, we use this network model, embedded inside a logistic regression, to evaluate the statistical significance of the inferred correlations. Then, we extend it to control for varying degrees of interest in news topics by different demographic groups.

3 DATA

This section describes the data sources we use as input for our analyses. First, Reddit and the r/news interaction graphs we built. Then, the socio-demographic features extracted from Reddit participation data. Finally, the news topics, that we employ as confounders.

3.1 Reddit interaction graph

Reddit is a social content aggregation website that has consistently ranked among the top ten most visited websites in the United States over the past years.[2] It is organized in topical communities, called *subreddits*, centered around a variety of topics. In particular, r/news, the focus of the current work, is dedicated to discussing "news articles about current events in the United States and the rest of the world".[3] Users can pseudonymously post submissions in these subreddits, and comment on other submissions and comments, thus creating a tree structure for the overall discussion.

Differently from other social media such as Facebook or Twitter, Reddit's homepage is organized around subreddits rather than user-to-user relationships. As such, the subreddits chosen by a user represent the main source of the information they consume on the website. We collect public data from Reddit by using Pushift [6].

Since Reddit data is massive—around 10^7 comments each year—we reduce our datasets to what we consider the most relevant interactions. First, we choose users active on the subreddit r/news from 2016 to 2020, included. Then, we select users by applying the following criteria: a minimum activity on r/news, a minimum activity on the whole Reddit, and that they are not likely to be a bot. In particular, a selected user must have written (*i*) at least 25

Table 1: Number of nodes and edges for the considered interaction graph of r/news in each year.

Year	2016	2017	2018	2019	2020
N. nodes	27 976	34 060	31 997	21 225	29 045
N. edges	1 166 076	1 390 243	1 221 779	793 569	1 067 614

between submissions and comments on r/news in a given year,[4] and (*ii*) submissions to at least 5 different subreddits in a given year. These empirically-determined cutoffs serve a twofold purpose: to select the most active users to avoid sparsity issues and to reduce computational costs. Moreover, as we show next, participation in other subreddits is necessary for estimating socio-demographic features. To exclude news bots, we filter out users who either (*i*) made submissions on r/news but no comments, (*ii*) belong to a list of known Reddit bots (e.g., automoderator), (*iii*) have posted in more than 50 subreddits each month, or (*iv*) have the string 'bot' in their username.

Starting from the selected set of users V, we create an interaction graph G by considering the users as nodes and defining an arc $u \rightarrow v$ if a user u replies to a message posted by user v. Table 1 reports the number of selected users and social interactions.

3.2 Socio-demographic features

To extract the socio-demographic features of the users, we use existing subreddit features obtained by embedding social interactions on Reddit. Waller and Anderson [70] assign a score for different demographic axes (such as age, gender, and affluence) and social axes (political leaning) to each subreddit. These scores are based on a set of well-thought pairs of seeds for each axis (e.g., r/teenagers and r/redditforgrownups for age), which are then used to construct the scores via embeddings, and finally validated with external sources of information. We project these scores onto the users, and define the score F_u of user u with respect to feature F as the weighted average of all their submissions in each subreddit,

$$F_u = \frac{\sum_s N_{u,s} F_s}{\sum_s N_{u,s}}, \tag{1}$$

where s is a subreddit, F_s is the feature score of subreddit s, and $N_{u,s}$ is the number of submissions by user u in subreddit s. For instance, if a subreddit is leaning towards 'male' in the gender axis, (i.e., it is more likely to be frequented by males), this scoring method will assign a larger probability of being male to a user who writes often on that subreddit.

Equation (1) allows us to characterize each user by a set of continuous scores that represent their position within each feature axis. To use these features in a logistic regression model, we quantile-normalize these values among the selected users, and—to streamline our analysis—we binarize them, labeling only the ones with the highest and lowest quartile for each axis. For instance, we label a user as male-leaning on the gender axis if their score falls in the lowest quartile of the gender axis (similarly, as female-leaning if in the highest quartile). We stress that a score indicating a female-leaning user does not imply that such a user is necessarily female, as Reddit is known to be participated by more men than women.
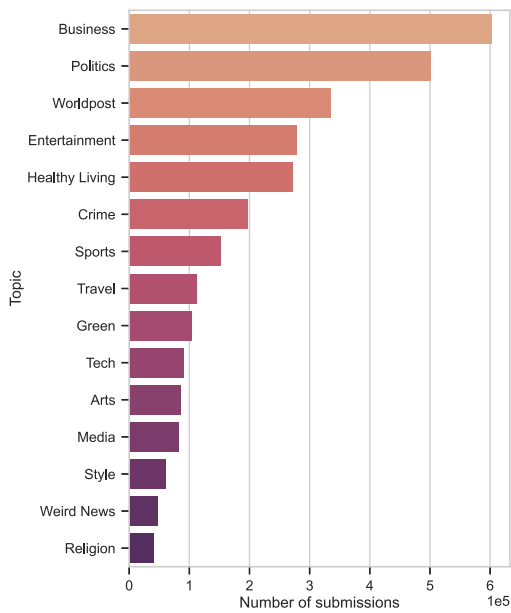
Figure 1: Distribution of topics across the selected submissions of `r/news`.

Rather, it indicates a user that frequents subreddits is more likely to be participated by women. Conversely, a higher score on the age axis simply indicates that the user is more likely to be older relative to the Reddit user base, which is quite young [5]. Similarly, being in the lowest or highest quartile on the partisan axis does not imply ideological extremeness, but rather, statistical confidence about their partisan affiliation given their Reddit activity. Users with very sparse or contradictory activity will instead be placed around the median: as such, we will consider them as a baseline.

3.3 Topic classifier

To check whether the effects of socio-demographic attributes on interactions are simply a byproduct of a varying degree of interest in news topics, we need to classify posts in `r/news` according to their news topic. To do so, first, we train a state-of-the-art supervised topic detection classifier on news article headlines. Then, we apply this classifier to the news appearing in `r/news`, so to obtain a mapping of posts into topics.

To train the supervised classifier we replicate results from the best-performing news topic detection model from Kaggle,[5] which we describe below. As a training set, we use a publicly available data set gathered by Misra [52]. This data set collects news headlines from the Huffington Post, each associated to the topic assigned by the editorial staff according to their topic categorization. Since the posts in `r/news` are links to newspaper articles, this data source is homogeneous with the one we collect. The data set contains 210 294 news headlines published between 2012 and 2022 and 40 topics. Then, we use GloVe embeddings [59] with a dimension of 100 to encode the words identified after tokenization. Finally, we employ a long short-term memory (LSTM) attention neural network to perform classification by using these embeddings. In 10-fold cross-validation on the news data set, this classifier obtains an average

[5]https://www.kaggle.com/code/hengzheng/news-category-classifier-val-acc-0-65/

Table 2: Sample of news headlines found in our Reddit data set for each considered topic.

Topic	News headline
Business	Goldman's Investment Bank to Increase its Hiring of Black Employees
Politics	Trump attacks Schiff: 'He lies awake at night shifting and turning'
Worldpost	From China to Italy, Coronavirus leaves a trail of devastation
Entertainment	Oscar 2020 Winners Full List
Healthy Living	Nursing Homes, Racked by the Virus, Face a New Crisis: Isolation
Crime	Video Shows Missouri Cop Run Over Fleeing Suspect With Car, Then Tackle Him As He Screams
Sports	Moloney Vs Baez - Boxing, June 25, 2020, On ESPN
Travel	Double rainbows form in skies over Delhi, surrounding areas
Green	YouTube Has Been 'Actively Promoting' Videos Spreading Climate Denialism, According to New Report
Tech	Parallels and Google team up to bring Windows apps to Chrome OS
Arts	In Montana, the Art of Crafting Fly-Fishing Rods
Media	Jim Lehrer, Longtime PBS NewsHour Anchor, Dies at 85
Style	New year gifts
Weird News	Cops: McDonald's mask scofflaw breaks window, steals panties
Religion	Join Wajahat Ali to discuss Ramadan during a pandemic

accuracy of 65%, which is considerably high considering that this is a 40-class classification problem.

After this training phase, we apply the resulting classifier to the 3 304 044 news headlines contained in our `r/news` data set and obtain a mapping of posts into topics. Since each social interaction takes place under a post, each link in our interaction network is associated to a single topic.

Figure 1 shows the 15 most popular topics in our data set. These 15 topics account for 90% of our data set; therefore, we exclude the remaining 27 topics from our analysis (the largest excluded topic accounts for 1.1% of the posts). The most popular topic is *Business* (18.2% of posts), followed by *Politics* (15.2%) and *Worldpost* (10.2%), a topic category related to news events taking place outside the United States. The topic *Entertainment* (8.4%) pertains mostly to news about movies and music. *Healthy Living* (8.2%) collects news regarding healthy lifestyles and habits. Other popular topics—beside self-explanatory ones such as *Crime*, *Tech*, and *Religion*—include *Green*, related to environment and climate change; *Media*, about TV and movies; *Style*, about fashion and beauty care; and *Weird News*, collecting strange-but-true headlines, from UFOs to pets. To further exemplify the topics, we report some headlines from our Reddit dataset in Table 2.

4 MODEL

This section introduces the two logistic regression models we develop to estimate the relationships among the socio-demographic groups in Reddit `r/news`. Each one is built upon a network null model where the dependent variable indicates whether a pair of nodes have interacted on the subreddit. The first one, dubbed socio-demographic features model (SD), uses only independent variables derived from the socio-demographic features. The second model, the socio-demographic model with topics (SD+T), investigates whether such relationships hold after controlling for different interest by each demographic group in each news topic.

4.1 Socio-demographic features model

The goal of the SD model is to estimate which pairs of socio-demographic features foster or inhibit interactions among Reddit

users. To do so, we rely on the following logistic regression approach, inspired by the feature-feature graph model [20, 51, 54]. Consider a set of features F and a directed graph $G = (V, E)$. Each node $v \in V$ is characterized by a binary vector of features $x_v \in \{0, 1\}^{|F|}$. Then, for a node pair $(u, v) \in V \times V$, let the dependent variable y be equal to 1 if the two nodes interact (i.e., $(u, v) \in E$) and 0 otherwise. Our model estimates the relationship of the dependent variable y with the vectors of independent variables (features) of the two nodes, x_u and x_v. More specifically, we consider all the possible pairs of features $F \times F$: the hypothesis is that the likelihood of observing an interaction from u to v can be affected by node u having feature h and node v having feature k, for any feature pair $(h, k) \in F \times F$. This model can be expressed by a feature-feature matrix $W \in \mathbb{R}^{|F| \times |F|}$, where each element $W_{h,k}$ indicates the log odds ratio for nodes with feature h to interact with nodes with feature k. A positive weight indicates that nodes with features h are more likely to interact with nodes with features k with respect to the probability of random interactions given by the null model (see below). Conversely, a negative element $W_{h,k}$ indicates that this specific pair of features inhibits interactions.

To convert this framework into a logistic regression model, we apply an outer-product kernel [54] to x_u, x_v. In other words, the independent variables of the logistic regression are all the elements of $x_u \otimes x_v$; that is, each pair of nodes u, v is associated to $|F| \cdot |F|$ variables such that the element h, k is 1 if node u has feature h and node v has feature k. Therefore, the model can be expressed as

$$\text{logit}(y_{u,v}) = \beta_0 + x_u^\intercal W x_v = \beta_0 + \sum_{h,k \in F} x_{u,h} W_{h,k} x_{v,k}. \quad (2)$$

where β_0 indicates the intercept, which adjusts the model according to the average probability of an interaction between two nodes.

In our context, the directed graph $G = (V, E)$ is the graph of interactions in Reddit r/news in a given year, as defined in Section 3.1. The features used in this model are the socio-demographic variables defined in Section 3.2. The set F is therefore composed by features: Young, Old, Male, Female, Poor, Rich, Left-leaning, and Right-leaning. We assign such features to nodes as explained in Section 3.2, so that $x_{u,k}$ is 1 iff node u is in the quartile of considered users more likely to be in socio-demographic group k. This way, the logistic regression model fits a given sequence of pairs $(y_{u,v}, x_u \otimes x_v)$, and obtains an estimate of the feature-feature matrix W. Such a matrix indicates which socio-demographic groups are more likely to interact.

The construction of the sequence of node pairs is crucial, as it represents the null model. We would like this null model to be the random graph obtained by reshuffling the links of the original network, while preserving the in- and out-degree of each node—that is, a configuration model [53]: this way, the activity of each user is properly taken into account. Therefore, we define a balanced sequence of pairs, i.e., we sample a number of negative examples equal to the positive ones. We do so by, first, including all positive examples (i.e., the links E); then, as negative examples, we choose node u with probability proportional to their activity (i.e., the number of posted comments), and node v with probability proportional to their attractiveness (the number of received comments) [20]. If the obtained pair (u, v) is a link, we discard it. This way, the null model expresses the probability of considering a pair of nodes as

the product of two independent probabilities—the probability that a node initiates an interaction, and the probability that a node receives it. The estimated matrix W will therefore represent how the socio-demographic features make the observed links deviate from this null model. Since the estimation of matrix W is obtained from a logistic regression model, it is also accompanied by a statistical test able to tell if each element $W_{h,k}$ is significantly different from zero or if it might be spurious.

4.2 Socio-demographic model with topics

One might wonder whether the status homophily we see is just the result of shared interests. That is, similar people are interested in and comment on similar topics, and therefore are more likely to interact simply because they are drawn towards the same posts. In other words, we wish to assess whether the topic of the post (news) acts as a mediator [58], or conversely, whether there are any residual effects of the socio-demographic features once we control for the topic, which may act as a confounder [58].

To this aim, we augment the SD model to account for the topic of the post by introducing a feature-topic matrix $Q \in \mathbb{R}^{|F| \times |T|}$, where T is the set of possible topics defined in Section 3.3. Each entry $Q_{h,t}$ of the matrix represents the log odds ratio that a user with feature h is interested in topic t (and therefore comments on it). To formalize the model we make use of a binary indicator (one-hot) vector $\tau^p \in \{0, 1\}^{|T|}$ which for each post p (on which the interaction between u and v happens) has value 1 only for the topic of the post, and is 0 otherwise.

We integrate this feature-topic component Q into the model expressed by Equation (2) by assuming that, beside the feature-feature effects expressed by W, the probability of observing a link $u \to v$ on topic t is also affected by the interest of user u in topic t, and the interest of user v in the same topic. Thus, we can formally express this model is as follows:

$$\text{logit}(y_{u,v}) = \beta_0 + x_u^\intercal W x_v + x_u^\intercal Q \tau^p + x_v^\intercal Q \tau^p =$$
$$\beta_0 + \sum_{h,k \in F} x_{u,h} W_{h,k} x_{v,k} + \sum_{h \in F} x_{u,h} Q_{h,t} + \sum_{k \in F} x_{v,k} Q_{k,t},$$
$$(3)$$

where t represents the index of the topic of post p in T.

The matrix Q is the same for both the initiator and the receiver of the interaction, i.e., it is shared between x_u and x_v. In fact, it is also possible to write the topic terms of Equation (3) as $(x_u + x_v)^\intercal Q \tau^p$, where $(x_u + x_v)_k$ indicates therefore the number of participants in the interaction with feature k. As such, Q represents a non-directional effect that captures the interest of socio-demographic groups in each news topic. In addition, the increase in complexity of the model is limited: we need to estimate only an additional $|F| \cdot |T|$ parameters. If the status homophily we measure is due simply to similar interests, we would expect the parameters Q to capture most of the effect, and therefore make W not significant.

5 RESULTS

We now use the data we gathered in combination with these models to answer our research question. We articulate our main research goal into the following research questions:
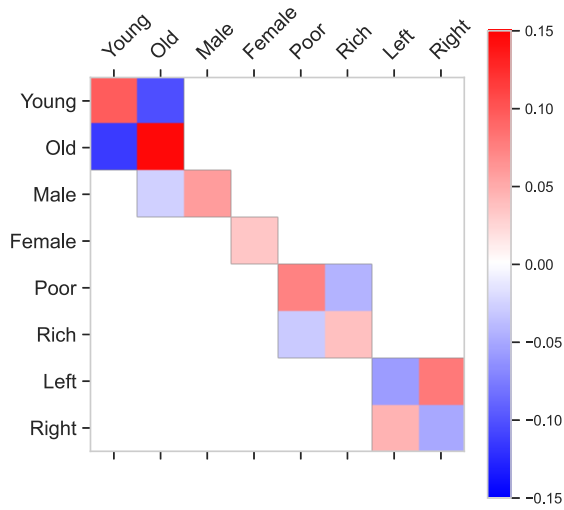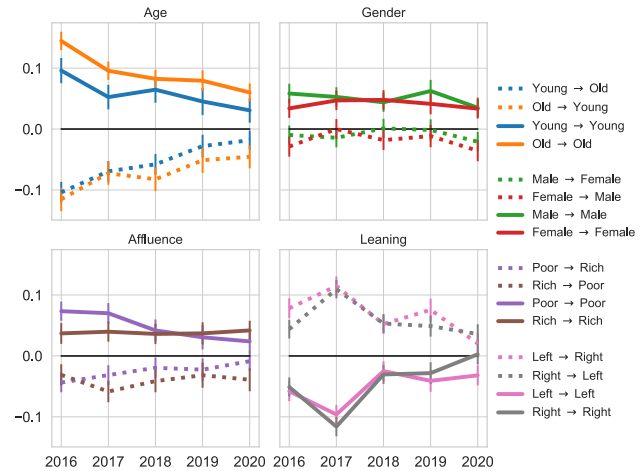
Figure 2: Logistic regression coefficients for interactions between all pairs of demographic features obtained for the `r/news` subreddit in 2016 by the socio-demographic features model. We show only coefficients that are significant in at least 4 years out of 5.

Figure 3: Logistic regression coefficients for interactions between socio-demographic features obtained (independently) for each year in the `r/news` subreddit in the SD model.

RQ1 Are there ideological echo chambers in `r/news`? I.e., are interactions among users on the same political side more likely?

RQ2 Is there evidence of demographic segregation? I.e., are interactions among users within the same demographic group more likely?

RQ3 Do our findings on RQ1 and RQ2 hold after controlling for varying degrees of interests in news topics by different demographic groups?

RQ4 Which news topic is each demographic group more interested in?

First, we use the SD model to assess whether interactions are more homophilic or heterophilic with respect to ideological (RQ1) and demographic attributes (RQ2). Then, we use the SD+T model to investigate whether such results hold after controlling for the varying degree of interest in news topics (RQ3). Finally, we analyze the results of the latter model to study which news topics are more attractive for the different socio-demographic groups (RQ4).

5.1 SD model

We apply the SD model to the interaction graph, using the 8 socio-demographic features we defined, separately for each considered year (2016 to 2020). We obtain an estimate of the feature-feature matrix W for each year. As a first representation of our results, Figure 2 shows the matrix of one year (2016), but shows only the coefficients that are statistically significant in at least four years out of five. More specifically, in this context, we define a coefficient $W_{h,k}$ as significant if the null hypothesis of $W_{h,k} = 0$ is rejected (p-value < 0.05), and the coefficient $W_{h,k}$ has the same sign across all the considered years. The base case for the model is a zero-valued feature vector, i.e., a user who is not in the top or bottom quartile on any socio-demographic axis.

Observing Figure 2, the most striking characteristic is the diagonal—i.e., each socio-demographic group tends to interact significantly more, or less, with members of the same group. In other words, each pair of features display a consistent effect in terms of heterophily and homophily. Remarkably, this effect is heterophily for the political leaning and homophily for the demographic features. Left-leaning and right-leaning users interact significantly more across the boundary, and significantly less within their respective group, w.r.t. the null model. Interactions with the opposite ideological camp are the norm and not the exception in `r/news`: our answer to RQ1 is therefore negative. Conversely, regarding age, gender, and affluence, users tend to interact significantly more *within* their demographic group. For age and affluence, in particular, we observe a consistent tendency of users to segregate: belonging to different groups inhibits interactions between users, thus leading to a positive answer to RQ2.

Since we observe that most of the significant coefficients belong to interactions between pairs of demographic features, we further investigate them in Figure 3. In particular, we show the block-diagonal of the weight matrix for each socio-demographic axis, one per subplot. On the x-axis we report the year, on the y-axis the value of the W coefficient (error bars represent the 95% CI), solid lines represent within-class interactions, and dotted lines represent cross-class ones. Again, the most important global pattern that emerges from the figure is that, for demographic features, within-class interactions are more frequent than expected, while the opposite is true for political leaning. Overall, interactions on `r/news` are characterized by *status homophily* and *value heterophily*.

The pattern is consistent and statistically significant over the years, with the exception of the gender axis, where the within-class interactions are not far from the null model. We also see a general trend towards convergence as we get closer to 2020 (especially pronounced in the age and leaning axes). This effect might be due to the highly polarizing 2016 US presidential election, which saw an important participation on Reddit, while in 2020 Trump supporters
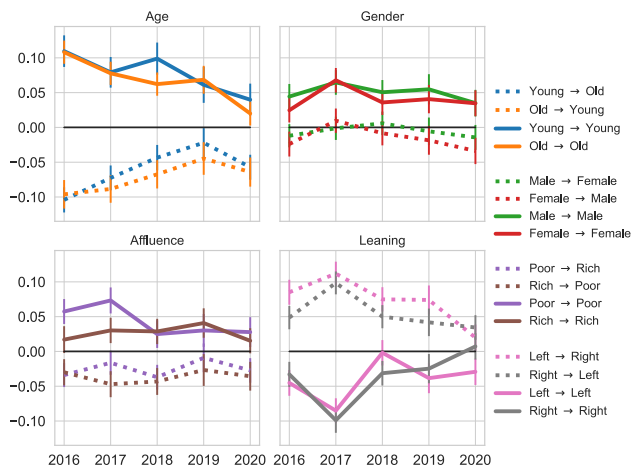
Figure 4: Logistic regression coefficients for interactions between demographic features obtained (independently) for each year in the `r/news` subreddit in the SD+T model.

were banned from Reddit, and resorted to alternative platforms such as Gab and Parler [2, 71].

5.2 SD+T model

We now move to the socio-demographic model with topics, which controls for the topic of the discussion where the interaction takes place. Figure 4 shows the weights of the matrix W for the logistic regression model of socio-demographic model with topics (again, the block-diagonal for each axis). The main result is that the patterns do not change much from what we observed in Figure 3, which leads to a positive answer to RQ3. Indeed, we observe again status homophily and value heterophily, as within-class interactions along the age, gender, and affluence axes are more likely than expected, while cross-class ones are more likely for political leaning. The same trends are present also in this case: convergence towards zero as we approach 2020, and non-significance of the cross-class gender terms. Overall, controlling for the topic does not alter the results of the model, although it considerably improves its fit. This fact indicates that the topic acts as a confounder and not as a mediator: while the homophily and heterophily effects are not caused solely by the topic of the post, interactions are nonetheless partially driven by the different interests of demographic groups in different topics.

Therefore, we look at the coefficients of the topic-class interactions encoded in matrix Q of socio-demographic model with topics in order to answer RQ4. Figure 5 shows the weights of the matrix for one year (2016, the other years are qualitatively similar). Each row represents a topic, each column an axis, and the value for each coefficient is full if it is statistically significant at least 4 out of 5 years, otherwise, it is hollow. The error bars represent the 95% CI of the estimates. A few interesting observations can be made.

On the age axis, younger people are more interested in arts, technology, and entertainment, while older users are more interested in business and politics. Older users are also more interested in green/environment topics, which might seem counter-intuitive, but can be understood by looking at the demographic of Reddit's user base, where older users are in their 30's or 40's, and younger

ones are in their teens. On the gender axis, typically male users show more interest in technology and business, but also lifestyle. More female users are interested in religion, weird news, healthy living, and lifestyle. On the affluence axis, low-income users are interested in religion and world news, while in general, they seem less interested in environment, travel, healthy living, business, and sports. Conversely, richer users are interested in technology, travel, business, and to a lesser degree in healthy living and entertainment.

Finally, on the political leaning axis, we see a strong effect for right-wing users who are clearly less interested in arts and lifestyle news, and also show disinterest in several other topics including technology, environment, healthy living, entertainment, and business. The only sizeable positive coefficients are for world news and crime. Left-leaning users, instead, show an interest in arts, media, and entertainment.

6 DISCUSSION

We have studied social interactions among Reddit users in an opinion formation context: the discussion of news on the subreddit `r/news`. By looking at the process through a socio-demographic lens, we have uncovered informative and consistent patterns. Specifically, the interaction network displays clear status homophily, whereby users similar in age, gender, and affluence, are more likely to interact. In addition, we find value heterophily, as users with opposing political leaning are also more likely to interact. Importantly, these patterns are robust to controlling for the interest in the topic of the news: the preference of a socio-demographic group for a news topic is not sufficient to explain the bias we observe. Even inside the same news topic, users' interactions are biased toward demographic homophily and ideological heterophily.

Overall, these results paint a picture where ideological echo chambers are not evident in news discussions, while we find evidence of demographic segregation. Strikingly, this demographic segregation happens even though the demographic traits at play are not immediately available to the discussants. Indeed, as the discussions happens online and in writing, the age, gender, and affluence of the user is a piece of latent information. Nevertheless, these traits affect the people with whom one chooses to interact.

The most likely explanation for this result supports the theory that one's worldview is affected by the demographic group they belong to. These worldviews, in turn, are evident in the opinions expressed through writing on online social media, which in turn drives the homophilic interaction process. This result runs against the common narrative of the Web and the Internet as a 'global village' [47]. Note that Harteveld [33] showed that higher levels of demographic segregation are associated with higher levels of affective polarization. As such, our results support the hypothesis that the increase in affective polarization observed over the last decades might be more associated with an increasingly divided society than with online echo chambers.

Our findings also confirm that there is no ideological selective exposure in news discussions on Reddit. While it is known that more niche subreddits might form echo chambers, with users moving within ideologically uniform communities [63], our results strengthen previous analyses about the lack of ideological echo chambers in generalist Reddit communities [20]. Conversely, echo
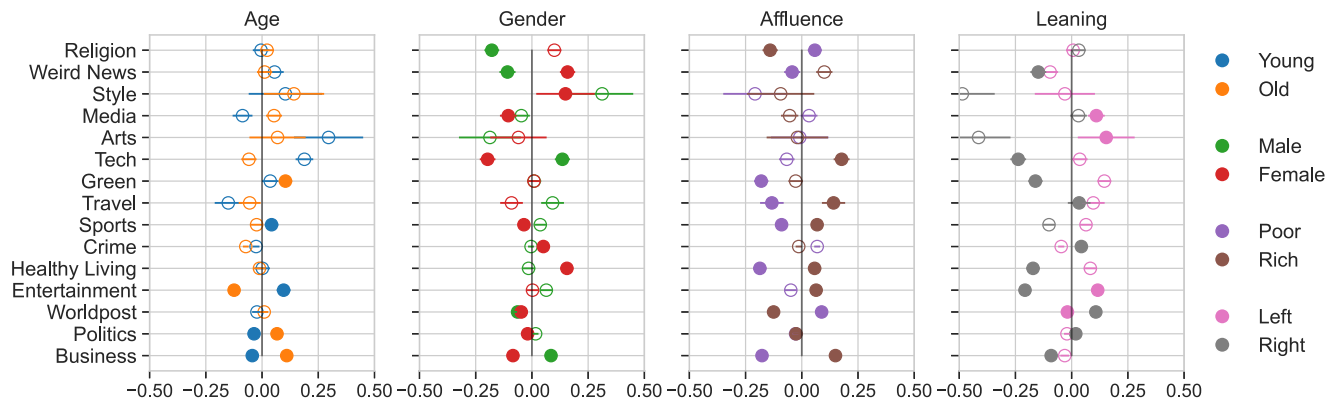
Figure 5: Logistic regression coefficients (and 95% CI) for interactions between demographic features and topics obtained for the `r/news` subreddit in 2016 by the SD+T model. Filled disks indicate associations significant in at least 4 years out of 5.

chambers have been reportedly observed on social media such as Twitter and Facebook [3, 4, 16, 27]. Cinelli et al. [15] have advanced the hypothesis that this difference between Reddit and other social media derives from their different structure: interest-based forum versus interpersonal social networks. Future work should investigate if, on these social networks, demographic segregation also contributes to affective polarization.

Our analysis also revealed some interesting patterns about the preferences of different socio-demographic groups toward certain news topics. It is already known that newspapers and media producers target news content to their audience, given their advertising-based financial model [22]. Some of these preferences have also been reported in the literature. For instance, Meeter et al. [49] find that gender influences news retention about business news, both age and educational attainment affect political news, and educational attainment also had an effect on science news. It is reasonable to expect that also news consumption would be affected, and indeed, the business topic shows a marked distinction on the gender axis (with males more interested in the topic), and age has an effect on the interest in political news (older users are more interested).

In a similar spirit, Lee and Chyi [42] investigate the relationship between demographic variables and the noteworthiness of news. They find that local news is perceived as more noteworthy by women, older, and less affluent people. In contrast, political and international news is favored by older and more educated people. Business and sports are quite different, with the former perceived as more noteworthy by older, more educated, and more affluent people, while the latter is more noteworthy for men and white people. Finally, entertainment is one category where younger adults have more interest than older people. On all the demographic variables that we also measure, our results are in line with what was found by Lee and Chyi [42], who used a completely different methodology. It is well known that different news sources have different attractiveness for different socio-demographic groups [64]. For instance, male and wealthy audiences prefer business news, such as the ones published on the Economist and the Wall St. Journal [14]. We see this effect in our results.

Conversely, entertainment news is often associated with a young and female audience [31]. We see more mixed evidence for this pattern, with, for instance, entertainment news being more engaging for younger users, but without any effect for gender, and sports news with a similar effect for age but being more engaging for a male audience. Healthy living and weird news, arguably two of the main soft news categories, show the opposite pattern, with a more female audience but no difference in age.

These effects suggest that, through demographic segmentation of the audience, online news might be contributing to widening the rift between demographic groups. For instance, demographic segmentation has been observed even in political advertising on social media, where different demographic groups witness different messages [13]. However, this work shows that the separation between demographic groups goes beyond their interests in different news topics. Social media users are self-sorting [33] by demographic groups, whereby interactions across groups are less likely. The widening rift between demographic groups, whose worldviews and material conditions are growing apart, is a phenomenon that should be taken more into consideration, even in social media analysis.

This finding suggests a few research directions to overcome the limitations of the present work. On a methodological level, this study is limited to a few demographic traits, estimated via a black-box approach. Firstly, it would be important to develop a white-box, probabilistic assessment of demographic groups online: having a Bayesian estimate would help to quantify the statistical significance of segregation patterns. Secondly, the inclusion of more demographic attributes beyond gender, age, and affluence, could highlight other important phenomena in online social interactions. Furthermore, we restrict our analysis to `r/news`: other venues for online discussions might present different characteristics. Beyond improving the methodology and broadening the scope, in future works one should investigate the causal pathways that connect demographic segregation and affective polarization. In this work, we have quantitatively compared two competing hypotheses to explain affective polarization and found that demographic segregation is more prominent than ideological echo chambers on Reddit news discussions. More analysis of the causal connections between these important aspects of our society is urgently needed.

REFERENCES

[1] Samuel J Abrams and Morris P Fiorina. 2012. "The big sort" that wasn't: A skeptical reexamination. *PS: Political Science & Politics* 45, 2 (2012), 203–210. (Cited on 2)

[2] Max Aliapoulios, Emmi Bevensee, Jeremy Blackburn, Barry Bradlyn, Emiliano De Cristofaro, Gianluca Stringhini, and Savvas Zannettou. 2021. An early look at the parler online social network. *arXiv preprint arXiv:2101.03820* (2021). (Cited on 7)

[3] Jisun An, Daniele Quercia, and Jon Crowcroft. 2014. Partisan Sharing: Facebook Evidence and Societal Consequences. In *COSN'14: ACM Conference on Online Social Networks*. 13–24. (Cited on 2, 8)

[4] Eytan Bakshy, Solomon Messing, and Lada A Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132. (Cited on 2, 8)

[5] Michael Barthel, Galen Stocking, Jess Holcomb, and Amy Mitchell. 2016. Reddit news users more likely to be male, young and digital in their news preferences. *Pew Research Center* (2016). (Cited on 4)

[6] Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. The pushshift reddit dataset. *arXiv preprint arXiv:2001.08435* (2020). (Cited on 3)

[7] A. Bessi, F. Zollo, M. Del Vicario, M. Puliga, A. Scala, G. Caldarelli, B. Uzzi, and W. Quattrociocchi. 2016. Users polarization on Facebook and Youtube. *PLOS ONE* 11, 8 (2016), e0159641. (Cited on 2)

[8] Bill Bishop. 2009. *The big sort: Why the clustering of like-minded America is tearing us apart.* Houghton Mifflin Harcourt. (Cited on 2)

[9] Paolo Boldi, Irene Crimaldi, and Corrado Monti. 2016. A network model characterized by a latent attribute structure with competition. *Information Sciences* 354 (2016), 236–256. (Cited on 3)

[10] Paolo Boldi and Corrado Monti. 2016. Llamafur: learning latent category matrix to find unexpected relations in wikipedia. In *Proceedings of the 8th acm conference on web science*. 218–222. (Cited on 3)

[11] Levi Boxell, Matthew Gentzkow, and Jesse M Shapiro. 2017. Greater Internet use is not associated with faster growth in political polarization among US demographic groups. *Proceedings of the National Academy of Sciences* 114, 40 (2017), 10612–10617. (Cited on 1)

[12] Max Bradley and Simon Chauchard. 2022. The Ethnic Origins of Affective Polarization: Statistical Evidence from Cross-National Data. *Frontiers in Political Science* (2022), 64. (Cited on 1, 3)

[13] Arthur Capozzi, Gianmarco De Francisci Morales, Yelena Mejova, Corrado Monti, André Panisson, and Daniela Paolotti. 2021. Clandestino or Rifugiato? Anti-immigration Facebook Ad Targeting in Italy. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15. (Cited on 8)

[14] Pew Research Center. [n. d.]. Demographics and Political Views of News Audiences — pewresearch.org. https://www.pewresearch.org/politics/2012/09/27/section-4-demographics-and-political-views-of-news-audiences/. [Accessed 12-Feb-2023]. (Cited on 8)

[15] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The Echo Chamber Effect on Social Media. *PNAS: Proceedings of the National Academy of Sciences* 118, 9 (2021), e2023301118. https://doi.org/10.1073/pnas.2023301118 (Cited on 1, 2, 8)

[16] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on twitter. In *Proceedings of the international aaai conference on web and social media*, Vol. 5. 89–96. (Cited on 1, 2, 8)

[17] Alessandro Cossard, Gianmarco De Francisci Morales, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2020. Falling into the Echo Chamber: The Italian Vaccination Debate on Twitter. In *ICWSM '20: Fourteenth International AAAI Conference on Web and Social Media*. 130–140. (Cited on 2)

[18] Wesley Cota, Silvio C. Ferreira, Romualdo Pastor-Satorras, and Michele Starnini. 2019. Quantifying echo chamber effects in information spreading over political communication networks. *EPJ Data Science* 8, 1 (2019), 35. https://doi.org/10.1140/epjds/s13688-019-0213-9 (Cited on 1, 2)

[19] Kareem Darwish. 2019. Quantifying polarization on twitter: the kavanaugh nomination. In *International conference on social informatics*. Springer, 188–201. (Cited on 2)

[20] Gianmarco De Francisci Morales, Corrado Monti, and Michele Starnini. 2021. No Echo in the Chambers of Political Interactions on Reddit. *Scientific Reports* 11 (2021), 2818. https://doi.org/10.1038/s41598-021-81531-x (Cited on 1, 2, 5, 7)

[21] Elizabeth Dubois and Grant Blank. 2018. The echo chamber is overstated: the moderating effect of political interest and diverse media. *Information, communication & society* 21, 5 (2018), 729–745. (Cited on 1, 2)

[22] Erick Elejalde, Leo Ferres, and Rossano Schifanella. 2019. Understanding news outlets' audience-targeting patterns. *EPJ Data Science* 8, 1 (2019), 1–20. (Cited on 8)

[23] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly* 80, S1 (2016), 298–320. Publisher: Oxford University Press tex.date-added: 2019-11-13 17:02:52 +0000

[24] Richard Fletcher and Rasmus Kleis Nielsen. 2018. Are people incidentally exposed to news on social media? A comparative analysis. *New media & society* 20, 7 (2018), 2450–2468. (Cited on 1)

[25] Linton C Freeman. 1978. Segregation in social networks. *Sociological Methods & Research* 6, 4 (1978), 411–429. (Cited on 2)

[26] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2016. Quantifying Controversy in Social Media. In *WSDM '16: 9th ACM International Conference on Web Search and Data Mining*. 33–42. (Cited on 2)

[27] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. In *Proceedings of the 2018 World Wide Web Conference*. 913–922. (Cited on 1, 2, 8)

[28] Venkata Rama Kiran Garimella and Ingmar Weber. 2017. A long-term analysis of polarization on Twitter. In *Eleventh international AAAI conference on web and social media*. (Cited on 2)

[29] R Kelly Garrett. 2009. Echo chambers online?: Politically motivated selective exposure among Internet news users. *Journal of Computer-Mediated Communication* 14, 2 (2009), 265–285. (Cited on 1)

[30] Andrew Guess, Brendan Nyhan, Benjamin Lyons, and Jason Reifler. 2018. Avoiding the echo chamber about echo chambers. *Knight Foundation* (2018). (Cited on 2)

[31] James T Hamilton. 2011. All the news that's fit to sell. In *All the News That's Fit to Sell*. Princeton University Press. (Cited on 8)

[32] Tal Orian Harel, Jessica Katz Jameson, and Ifat Maoz. 2020. The normalization of hatred: Identity, affective polarization, and dehumanization on Facebook in the context of intractable political conflict. *Social Media+ Society* 6, 2 (2020), 2056305120913983. (Cited on 1, 2)

[33] Eelco Harteveld. 2021. Ticking all the boxes? A comparative study of social sorting and affective polarization. *Electoral Studies* 72 (2021), 102337. (Cited on 7, 8)

[34] Bas Hofstra, Rense Corten, Frank Van Tubergen, and Nicole B Ellison. 2017. Sources of segregation in social networks: A novel approach using Facebook. *American Sociological Review* 82, 3 (2017), 625–656. (Cited on 2)

[35] Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. 2019. The origins and consequences of affective polarization in the United States. *Annual Review of Political Science* 22, 1 (2019), 129–146. (Cited on 1, 3)

[36] Shanto Iyengar, Gaurav Sood, and Yphtach Lelkes. 2012. Affect, not ideology: a social identity perspective on polarization. *Public opinion quarterly* 76, 3 (2012), 405–431. (Cited on 1, 2)

[37] Shanto Iyengar and Sean J Westwood. 2015. Fear and loathing across party lines: New evidence on group polarization. *American journal of political science* 59, 3 (2015), 690–707. (Cited on 2)

[38] Chris Jackson. 2018. Fake News, Filter Bubbles, and Post-Truth Are Other People's Problems. *Global Market Research and Public Opinion Specialist (IPSOS), Washington DC* (2018). (Cited on 1)

[39] Gary C Jacobson. 2016. Polarization, gridlock, and presidential campaign politics in 2016. *The ANNALS of the American Academy of Political and Social Science* 667, 1 (2016), 226–246. (Cited on 2)

[40] Joohan Kim, Robert O Wyatt, and Elihu Katz. 1999. News, talk, opinion, participation: The part played by conversation in deliberative democracy. *Political communication* 16, 4 (1999), 361–385. (Cited on 2)

[41] Paul F Lazarsfeld, Robert K Merton, et al. 1954. Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society* 18, 1 (1954), 18–66. (Cited on 1, 3)

[42] Angela M Lee and Hsiang Iris Chyi. 2014. When Newsworthy is Not Noteworthy: Examining the value of news from the audience's perspective. *Journalism studies* 15, 6 (2014), 807–820. (Cited on 8)

[43] Matthew Levendusky. 2009. *The partisan sort: How liberals became Democrats and conservatives became Republicans.* University of Chicago Press. (Cited on 2)

[44] Trevon D Logan and John M Parman. 2017. The national rise in residential segregation. *The Journal of Economic History* 77, 1 (2017), 127–170. (Cited on 2)

[45] Wesley Lowery, Kimberly Kindy, and Andrew Ba Tran. 2018. In the United States, right-wing violence is on the rise. *Washington Post* 25 (2018). https://www.washingtonpost.com/national/in-the-united-states-right-wing-violence-is-on-the-rise/2018/11/25/61f7f24a-deb4-11e8-85df-7a6b4d25cfbb_story.html (Cited on 1)

[46] Lilliana Mason and Julie Wronski. 2018. One tribe to bind them all: How our social group attachments strengthen partisanship. *Political Psychology* 39 (2018), 257–277. (Cited on 1, 3)

[47] Marshall McLuhan. 1994. *Understanding media: The extensions of man.* MIT press. (Cited on 7)

[48] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* (2001), 415–444. (Cited on 1, 3)

[49] Martijn Meeter, Dirkjan JC Ochtman, Steve MJ Janssen, and Jaap MJ Murre. 2010. Of sports and politics: Predicting category-specific retention of news events from demographic variables. *European Journal of Cognitive Psychology* 22, 1 (2010), 117–129. (Cited on 8)

[50] George Michael. 2017. The rise of the alt-right and the politics of polarization in America. *Skeptic (Altadena, CA)* 22, 2 (2017), 9–18. (Cited on 1)

[51] Kurt Miller, Michael Jordan, and Thomas Griffiths. 2009. Nonparametric latent feature models for link prediction. *Advances in neural information processing systems* 22 (2009). (Cited on 3, 5)

[52] Rishabh Misra. 2022. News category dataset. *arXiv preprint arXiv:2209.11429* (2022). (Cited on 4)

[53] Michael Molloy and Bruce Reed. 1995. A critical point for random graphs with a given degree sequence. *Random structures & algorithms* 6, 2-3 (1995), 161–180. (Cited on 5)

[54] Corrado Monti and Paolo Boldi. 2017. Estimating latent feature-feature interactions in large feature-rich graphs. *Internet Mathematics* 1, 1 (2017), 2640. (Cited on 3, 5)

[55] Alfredo Jose Morales, Javier Borondo, Juan Carlos Losada, and Rosa M Benito. 2015. Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 25, 3 (2015), 033114. (Cited on 2)

[56] Krzysztof Nowicki and Tom A B Snijders. 2001. Estimation and prediction for stochastic blockstructures. *Journal of the American statistical association* 96, 455 (2001), 1077–1087. (Cited on 3)

[57] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you.* penguin UK. (Cited on 1)

[58] Judea Pearl. 2009. *Causality.* Cambridge University Press. (Cited on 5)

[59] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP).* 1532–1543. (Cited on 4)

[60] Pew Research Center. 2017. *The partisan divide on political values grows even wider.* Technical Report. https://www.pewresearch.org/politics/2017/10/05/the-partisan-divide-on-political-values-grows-even-wider (Cited on 2)

[61] Joseph Phillips. 2022. Affective Polarization: Over Time, Through the Generations, and During the Lifespan. *Political Behavior* (2022), 1–26. (Cited on 1, 3)

[62] Walter Quattrociocchi, Antonio Scala, and Cass R Sunstein. 2016. Echo chambers on Facebook. *Available at SSRN 2795110* (2016). (Cited on 2)

[63] Cesare Rollo, Gianmarco De Francisci Morales, Corrado Monti, and André Panisson. 2022. Communities, Gateways, and Bridges: Measuring Attention Flow in the Reddit Political Sphere. In *Social Informatics: 13th International Conference, SocInfo 2022, Glasgow, UK, October 19–21, 2022, Proceedings.* Springer, 3–19. (Cited on 7)

[64] Stacy Rosenberg. Pew Research Center. Who is Listening, Watching, Reading - and Why — pewresearch.org. https://www.pewresearch.org/politics/2010/09/12/section-4-who-is-listening-watching-reading-and-why. [Accessed 12-Feb-2023]. (Cited on 8)

[65] Thomas C Schelling. 1969. Models of segregation. *The American economic review* 59, 2 (1969), 488–493. (Cited on 2)

[66] Ana Lucía Schmidt, Fabiana Zollo, Antonio Scala, Cornelia Betsch, and Walter Quattrociocchi. 2018. Polarization of the vaccination debate on Facebook. *Vaccine* 36, 25 (2018), 3606–3612. (Cited on 2)

[67] Cass R Sunstein. 2005. *Why societies need dissent.* Vol. 9. Harvard University Press. (Cited on 2)

[68] Cass R Sunstein. 2009. *Republic.com 2.0.* (Cited on 2)

[69] Kathie Treen, Hywel Williams, Saffron O'Neill, and Travis G Coan. 2022. Discussion of Climate Change on Reddit: Polarized Discourse or Deliberative Debate? *Environmental Communication* (2022), 1–19. (Cited on 1)

[70] Isaac Waller and Ashton Anderson. 2021. Quantifying social organization and political polarization in online platforms. *Nature* 600, 7888 (2021), 264–268. (Cited on 2, 3)

[71] Yuchen Zhou, Mark Dredze, David A Broniatowski, and William D Adler. 2018. Gab: The alt-right social media platform. *John Hopkins University* (2018). (Cited on 7)