Troll and Divide: The Language of Online Polarization

Almog Simchon^{1,2*} William J. Brady³ Jay J. Van Bavel^{4,5*}

¹Department of Psychology, Ben-Gurion University of the Negev, Beer Sheva, Israel

²School of Psychological Science, University of Bristol, Bristol, UK

³Department of Psychology, Yale University, New Haven, CT, USA

⁴Department of Psychology, New York University, New York, NY, USA

⁵Center for Neural Science, New York University, New York, NY, USA

**Corresponding Authors:*

Almog Simchon, Ben-Gurion University of the Negev, Department of Psychology, POB 653, Beer Sheva, 8410501, Israel, almogsi@post.bgu.ac.il

Jay J. Van Bavel, New York University, Department of Psychology, New York University, New York, NY 10003, USA, jay.vanbavel@nyu.edu

Author contribution:

A.S. and W.J.B. conceived and designed the experiments; A.S. and W.J.B. performed the experiments; A.S. and W.J.B. analyzed the data; A.S., W.J.B., and J.V.B. contributed materials/analysis tools; A.S., W.J.B., and J.V.B. wrote the paper.

Acknowledgments: The authors would like to thank Dr. Michael Gilead for resource assistance, and to the Civiqs company and Dr. Kate Starbird for data sharing. We would like to acknowledge members of the NYU Social Identity and Morality Lab for comments on a previous version of this manuscript. This work was partially supported by research grants from the John Templeton Foundation to J.V.B.

Abstract

The affective animosity between the political left and right has grown steadily in many countries over the past few years, posing a threat to democratic practices and public health. There is a rising concern over the role that ‘bad actors’ or trolls may play in the polarization of online networks. In this research, we examined the processes by which trolls may sow intergroup conflict through polarized rhetoric. We developed a dictionary to assess online polarization by measuring language associated with communications that display partisan bias in their diffusion. We validated the polarized language dictionary in four different contexts and across multiple time periods. The polarization dictionary made out-of-set predictions, generalized to both new political contexts (#BlackLivesMatter) and a different social media platform (Reddit), and predicted partisan differences in public opinion polls about COVID-19. Then we analyzed tweets from a known Russian troll source ($N = 383,510$) and found that their use of polarized language has increased over time. We also compared troll tweets from three countries ($N = 798,33$) and found that they all utilize more polarized language than regular Americans ($N = 1,507,300$) and trolls have increased their use of polarized rhetoric over time. We also find that polarized language is associated with greater engagement, but this association only holds for politically engaged users (both trolls and regular users). This research clarifies how trolls leverage polarized language and provides an open-source, simple tool for exploration of polarized communications on social media.

Keywords: polarization, trolls, social media

Troll and Divide: The Language of Online Polarization

A growing body of research suggests that the American public has become more polarized over the past few decades (Klein, 2020; Mason, 2018). These attitudes are mirrored in rising partisan antipathy; dislike toward members of the opposing ideology--a phenomenon known as 'affective polarization' (Boxell et al., 2020; Finkel et al., 2020; Iyengar et al., 2019). The consequences of polarization include growing political radicalism (Warner, 2010), increased ingroup bias (Amira et al., 2019), and even different behavioral reactions to deadly pandemics (Gollwitzer et al., 2020). The alarming consequences of polarization are by no means limited to America: India, Poland, Columbia, Bangladesh, Israel, Indonesia, Britain, and Brazil, are just some of the many countries facing growing levels of political polarization (Carothers & O'Donohue, 2019), and some research attributes this intergroup conflict to the rise of social media (Allcott et al., 2020; Levy, 2021; but see Boxell et al., 2017; Van Bavel, Rathje, et al., 2021). In the current paper, we examine the language of online polarization employed by regular citizens and internet trolls.

On social media, polarization is often defined as emerging clusters of like-minded individuals who engage in confirmation bias and curate narratives congruent with their beliefs (Del Vicario et al., 2016; Sikder et al., 2020). The formation of like-minded social networks is particularly salient in social media platforms that deploy a news-feed algorithm (e.g., Facebook), or a computational formula that favors some pieces of content over others (Cinelli et al., 2021), creating filtered feeds of personalized content (Pariser, 2011). The combination of ideological or partisan groups joining like-minded networks and receiving algorithmically determined political content may be amplifying polarization (Allcott et al., 2020; Barberá et al., 2015; Brady et al., 2017; Yardi & Boyd, 2010). This trend has raised concerns that people may lose a shared sense of reality.

Although recent evidence suggests that general information consumption on social media might not be an echo-chamber for many users (Guess, 2020; Mukerjee et al., 2020), there is nevertheless substantial evidence supporting the argument that segregated online communities emerge around politically contentious topics (Baumann et al., 2020; Brady et al., 2017; Evans & Fu, 2018; Jasny et al., 2018; Starnini et al., 2016; Sunstein, 2018; Yardi & Boyd, 2010). Moreover, exposure to out-group partisans may even increase polarization (Bail et al., 2018). One damaging effect of ideology-based homophily is enabling and fostering the spread of

misinformation (Brady et al., 2020; Del Vicario et al., 2016). Falsehoods appear to spread farther, faster, deeper, and more broadly than the truth on social media, especially for political news (Vosoughi et al., 2018). As billions of people have opened social media accounts and use these platforms to get their news, it has also exposed them to a hotbed of conspiracy theories, misinformation, and disinformation (Lazer et al., 2018; Van Bavel, Harris, et al., 2021). The rise of misinformation has fueled an international health crisis during the COVID-19 pandemic, leading the World Health Organization to declare this an “infodemic” of misinformation.

There has also been growing concern over the role bad actors may play in online polarization and the spread of misinformation (e.g., anti-quarantine messages during COVID-19; Benson, 2020). For the past several years, cyberspace has been affected by organized groups of social media users, commonly referred to as ‘trolls’, who intentionally pollute online discourse. Since 2018, Twitter has been releasing the Twitter Transparency Report, archives of tweets authored by state-affiliated information operations¹. The most famous of these operations is the Internet Research Agency (IRA), also known as a Russian ‘Troll Farm’, The IRA has engaged in online political tactics to sew intergroup conflict and influence U.S. citizens during the 2016 presidential election (Badawy et al., 2018) and British citizens prior to the Brexit vote (Llewellyn et al., 2018). Similarly, other state-affiliated influence operations have been found in numerous countries, including Iran, Bangladesh, Venezuela, China, Saudi Arabia, Ecuador, the United Arab Emirates, Spain, and Egypt¹. In the current paper, we developed and validated a polarization dictionary and examined whether the rhetoric used by these troll operations was highly polarized.

Some evidence suggests that trolls tend to take on far-right topics and stances, spreading hate speech and islamophobia (Pintak et al., 2019). However, it would be inaccurate to say that trolls are only far-right leaning, and spreading conservative ideology may not even be their ultimate goal. Instead, their main goal appears to be creating polarization and fostering social conflict within democracies. For instance, during #BlackLivesMatter discourse on Twitter Russian trolls were heavily engaged in spreading messages from the two ends of the debate; both anti-BLM and pro-BLM (Arif et al., 2018). The same pattern was observed during online Anti-

¹ <https://transparency.twitter.com/en/information-operations.html>

Vaccine debates: trolls were found to echo both positions (pro and against vaccines; Broniatowski et al., 2018). Taken together, these data suggest that online trolls are attempting to polarize social media users during political discourse.

Overview

The current research had two goals: (i) to create a dictionary of polarized language (i.e., linguistic expressions that are associated with political polarization) and (ii) to examine how this language has been used by trolls around the world. We began by building a simple tool to measure polarized language. Previous work studied polarization through network analysis or by exploring topics known to be polarized (Demszky et al., 2019). These methodologies have several advantages (Garimella et al., 2018; Guerra et al., 2013) but can be computationally expensive, create a barrier for adoption for behavioral scientists who lack the required technical expertise, and are most likely context-dependent which can undercut replicability (Van Bavel et al., 2016). Here, we sought to validate a dictionary of polarized language that would be applicable across numerous contexts. In what follows we describe how the dictionary was constructed, its validation using different topics and time periods, and how it tracks dynamic changes in partisan opinions during a time of national polarization (the COVID-19 pandemic).

Next, we examined the online rhetoric of trolls and regular citizens using the polarization dictionary. We conducted a high-powered study using nearly 2,300,000 tweets from trolls in multiple countries and compared results to a random sample of American Twitter users. To help determine if trolls were using polarized rhetoric more than the average American (Broniatowski et al., 2018; Cosentino, 2020), we examined the levels of polarized language in their tweets when compared to a control group, and explored how levels of polarized language changed over time within each group. These studies suggest that polarized rhetoric was weaponized by online trolls during political discourse.

Method

Data collection

We used the SCI lab twitter database at Ben-Gurion University (Simchon et al., 2020). Tweets were collected from all 50 states in the United States and the District of Columbia. We extracted tweets between November 2017 and December 2019. Trolls' data was taken from the

Twitter Transparency Report (Twitter, 2018, 2019a, 2019b). Additional data collection was done using Twitter API 2.0 and the ‘academictwitterR’ R package (Barrie & Ho, 2021).

All research was conducted in accordance with the Departmental IRB committee at Ben-Gurion University and was ruled “exempt”.

Preprocessing

Our sample size consisted of 2,306,233 original tweets in the English language (retweets were filtered out): 383,510 by Russian trolls, 329,453 by Iranian trolls, 85,970 by Venezuelan trolls, and 1,507,300 by American Controls (random sample from our Twitter database with no specific text search). Following the exclusion of retweets, English tweets constituted 34% of the Russian trolls dataset, 15% of the Iranian trolls dataset, and 1.25% of the Venezuelan trolls dataset.

For our content-matched analysis, we extracted the 20 most-frequent hashtags that appeared on politically engaged Russian trolls tweets (#MAGA, #tcot, #BlackLivesMatter, #PJNET, #news, #top, #mar, #topl, #Trump, #2A, #IslamKills, #WakeUpAmerica, #FAKENEWS!, #GOPDebate, #NowPlaying, #TCOT, #ccot, #amb, #sports, #TrumpTrain) and searched for tweets posted in the USA with the same hashtags. After the exclusion of retweets, politically-engaged Russian trolls sample size was 55,726, and so was their politically-matched American controls (55,726).

We could not use our sample of American Controls for Study 4 as it lacked engagement metrics. Therefore, we collected a new control sample, matched in time and without a specific text search (1,144,767).

All tweets had links, tags, and emoticons removed prior to any linguistic analysis. Text mining was done using the ‘quanteda’ package (Benoit et al., 2018) using R (Versions 3.6.3 and 4.0.3).

Study 1: Development and Validation of a Polarization Dictionary

To develop a polarization dictionary, we synthesized data-driven methods and domain expertise. Specifically, we (i) explored the language associated with polarization in a data-driven fashion; (ii) manually pruned the dictionary; (iii) expanded the dictionary by using GloVe word-embeddings (Pennington et al., 2014) and (iv) employed manual trimming. The dictionary contained 205 words (e.g., *corruption*, *kill*, *lie*, *terrorists*, *political*, *stupid*; see online materials

for the full list) and its full development and psychometric properties are reported in the Supplementary Information. All the materials are publicly available on OSF <https://osf.io/bm8uy>.

Dictionary Validation

We first validated our dictionary on a subset of the original database used in its construction (Brady et al., 2017). The database included tweets about contentious political topics that showed a range of ingroup bias in their spread through social networks (i.e., they either were shared with only the political ingroup or spread to one or more outgroup members). We built the dictionary on a randomly selected 80% of the original dataset (N training set = 19,841) and tested it on the remaining 20% (N test set = 5,008). This out-of-sample testing was conducted to ensure the predictive performance of the model and to avoid overfitting. Data preprocessing included removing all duplicates from the data and automatically deleting links and emojis from the text. A polarization score was calculated based on the count of dictionary words in the text, normalized by the tweet length. The means reported below represent the average percentage of the text that was found in the dictionary (for a similar approach see LIWC; (Pennebaker et al., 2015). Our analysis found that the dictionary successfully discriminated between polarized and non-polarized tweets from the test set ($M_{\text{polarized}} = 6.70$, $SD_{\text{polarized}} = 9.08$, $N = 3696$; $M_{\text{non-polarized}} = 4.39$, $SD_{\text{non-polarized}} = 6.60$, $N = 1312$), $t(3156) = 9.79$, $p < .001$, *Cohen's d* = 0.27. In other words, our dictionary was able to determine which corpus was more likely to include polarized communications compared to another corpus.

To evaluate generalizability, we validated the polarization dictionary with a different political topic (i.e., different from the original research). We examined the effectiveness of the polarization dictionary in the context of the online #BlackLivesMatter (#BLM) discourse between December 2015 and October 2016, which focused on issues of racial justice in the USA. Prior work had studied the flow of information in #BLM tweets by using a machine learning clustering technique to identify distinct Twitter communities and quantifying the spatial retweet flow within and between clusters (Arif et al., 2018). The original dataset included 58,698 tweets², and we were able to retrieve 24,747 tweets out of the original sample from Twitter's API. Like in the prior validation, messages were categorized with regard to the spread of information; whether the tweets showed ingroup bias (retweeted within one political cluster), or

² <https://github.com/leo-gs/ira-reproducibility>

not (retweeted by a user from the other cluster, as classified by the authors). We applied our dictionary to the posts we were able to retrieve, and again we observed that ingroup bias messages contained more polarized language than messages that diffused between clusters ($M_{\text{ingroup bias}} = 5.54$, $SD_{\text{ingroup bias}} = 5.68$, $N = 24,077$; $M_{\text{diffused}} = 4.83$, $SD_{\text{diffused}} = 5.09$, $N = 670$), $t(716.06) = 3.58$, $p < .001$, *Cohen's d* = 0.13. This helped establish the generalizability of our dictionary to a novel political topic.

Beyond testing out-of-sample generalizability, we also tested cross-platform generalizability. We tested the polarization dictionary on the platform *Reddit* using a wider range of political topics. Reddit is an online social media platform that consists of many discussion forums, or communities, called *subreddits*, including several communities devoted to politics (Soliman et al., 2019). We extracted up to 1,000 messages from 36 political communities with established ideologies (18 from each political side). As a control group, we sampled up to 1,000 messages from 18 other communities, randomly sampled from a list of popular subreddits³. We collected 53,859 posts between June 2015 and December 2018 from the Pushshift Reddit API (Baumgartner et al., 2020). Following data cleaning, our sample size consisted of 49,230 original posts. We applied the polarization dictionary on the Reddit sample and conducted a one-way between-group ANOVA. A planned comparison between the political groups revealed a significant difference between the control and the other political communities ($M_{\text{left}} = 2.38$, $SD_{\text{left}} = 4.61$, $N = 17,005$; $M_{\text{right}} = 2.57$, $SD_{\text{right}} = 5.34$, $N = 15,859$; $M_{\text{control}} = 0.97$, $SD_{\text{control}} = 3.44$, $N = 16,366$), $t(49,227) = 34.81$, $p < .001$, *Cohen's d* = 0.31. More information is reported in the Supplementary Information. In other words, the rhetoric in political Reddit groups was more polarized than apolitical Reddit groups.

As a more stringent sensitivity test, we replaced the randomly sampled control group with a “neutral” reference of contentious topics. We extracted messages from the popular subreddit *NeutralPolitics* (www.reddit.com/r/NeutralPolitics), a reddit community devoted to factual and respectful political discourse. This sample consisted of 9,984 posts between April 2016 and December 2018 (9,772 after data cleaning). A planned comparison between the political groups revealed a significant difference in polarized rhetoric between *NeutralPolitics* and the other political communities ($M_{\text{left}} = 2.38$, $SD_{\text{left}} = 4.61$, $N = 17,005$; $M_{\text{right}} = 2.57$, $SD_{\text{right}} = 5.34$, $N = 15,859$; $M_{\text{neutral}} = 2.24$, $SD_{\text{neutral}} = 4.49$, $N = 9,772$), $t(42,633) = 4.12$, $p < .001$, *Cohen's d* = 0.04.

³ <https://github.com/saiarcot895/reddit-visualizations>

See Supplementary Information for more details. This suggests that polarized rhetoric was reduced among the reddit community focused on respectful political discourse (although we note that the effect size here is very small).

To determine if our dictionary would track dynamic changes in polarized public opinions over time, we compared polarized language with polls about U.S. citizens' concern about the COVID-19 pandemic. The data were collected from a representative panel by Civiqs⁴, an online polling and analytics company. Recent polls have revealed clear partisan differences between Democrats and Republicans in reported concerns about the COVID-19 pandemic--such that Democrats are consistently more concerned about the pandemic than Republicans (Van Bavel, 2020). We tested whether the language in tweets about coronavirus was associated with the partisan discrepancy in public opinion about COVID-19. We calculated a “partisan difference score” from February 25th until April 14th, 2020 by subtracting the daily Republican net concern from the daily Democratic net concern, as reported by Civiqs (the specific question was ‘how concerned are you about a coronavirus outbreak in your local area?’). The poll was based on responses from 22,256 respondents and included measures to avoid demographic and ideological biases.

To compare Twitter language to partisans' concern, we collected 553,876 Twitter messages from the United States within these dates that used the terms “covid” or “coronavirus”. We then applied the polarization dictionary to the tweets and aggregated by date. We found that polarized language on social media, measured by the mean % of words from our dictionary contained in the tweets, was positively associated with partisan differences in concern about the COVID-19 pandemic over time, $r(48) = .45, p = .001$, see Figure 1. A *post-hoc* analysis revealed that the correlation between poll responses and twitter language was strongest when Twitter language was lagged by eight days (i.e., $\text{poll}_{t0}, \text{twitter}_{t8}$) $r(40) = .67, p < .001$ (for a full lag table of 16 days, see supplementary Table S1). In other words, polarized rhetoric about COVID-19 mirrored polarization in public opinion over the early phase of the pandemic. This also suggests that the polarization dictionary may be useful in detecting future patterns of public opinion.

⁴ https://civiqs.com/results/coronavirus_concern

Figure 1. Dynamic polarization changes in polls of COVID-19 concern and polarized language on Twitter. The solid line represents partisan differences in COVID-19 concern ($N = 22,256$), and the dashed line represents the degree of polarized discourse on Twitter ($N = 553,876$, dashed line). Values on the X-axis represent the time, and values on the Y-axis represent standardized scores of the variables. The functions have gone through a locally estimated scatterplot smoothing (span = 0.33, degree = 1). Shaded areas around the regression line denote 95% CI.

Taken together, these four sets of analyses (cross-validation, out of set validation, cross-platform validation and predictive validation) provide converging validity for the dictionary, showcasing its ability to capture political polarization in language across four different contexts. For a summary of all validation steps, see Table 1.

Table 1. Summary of validation steps. Effect sizes correspond to Cohen's d' or Pearson's r . All tests are significant at $p < .001$.

Validation Type	N	Effect size
Cross Validation	5008	$d = 0.27$
Out of Set (BLM)	24,747	$d = 0.13$
Cross Platform (Reddit)	49,230	$d = 0.31$
Predictive Validation (COVID)	553,876	$r = 0.45$

Study 2

Study 2a: Polarization in Russian Trolls

Russian trolls, or anonymous social media accounts that are affiliated with the Russian government, were active around highly contentious political topics around the world, including in the United States and Britain (Badawy et al., 2018; Llewellyn et al., 2018). With the release of the Twitter Transparency Report, a sample of the Russian and other countries' operations were officially disclosed and used to study the role of trolls in amplifying political polarization (Arif et al., 2018; Broniatowski et al., 2018; Walter et al., 2020). Therefore, we hypothesized that state-affiliated trolls would use more polarized language on social media compared to ordinary Twitter users. We also examined how polarized language may have changed over time. For instance, if trolls' levels of polarized language are increasing over time, it would imply that trolls are spending increased energy toward tactics that sow discontent and aim to influence polarized discourse. On the other hand, levels of polarized language might be increasing among American Twitter users as well, similar to trends of affective polarization (Iyengar et al., 2019).

Results

We compared Twitter messages posted by trolls to an American control sample (collected from across the United States through the Twitter API). We only used original tweets that were posted in the English language and were most likely aimed for an international/American audience. The comparison was matched for the same time range (November 23, 2016 - May 30, 2018). We applied the polarization dictionary, which was generated from and validated on different datasets (see Study 1) to extract polarization scores. First, we found that Russian trolls ($M = 2.37$, $SD = 5.14$, $N = 61,413$) used significantly more polarized language than tweets sent by the control sample ($M = 1.47$, $SD = 5.35$, $N = 516,525$), $t(78,081) = 40.96$, $p < .001$, *Cohen's d* = 0.17. These results suggest that trolls are leveraging polarized language to push conflict among U.S. citizens in the context of political discourse. For the top 25 most used words adjusting for their frequency (tf-idf), see Figure S1 in the Supplementary Information.

However, not all trolls are equal. Research suggests that Russian trolls could be classified into five distinct types: Right, Left, News, Hashtag Gamers and Fearmongers (Linville & Warren, 2020). It could be argued that a cleaner analysis would only constitute Left and Right trolls, and should be contrasted with a politically engaged American sample. Therefore, we used the

Russian Troll classification⁵ (Linville & Warren, 2020), and matched an American sample for their content (via hashtag use, see Method section), posting time (January 2015-May 2018) and quantity. Again, we find that politically-oriented Russian trolls use significantly more polarized language than their politically matched American sample (Russian trolls: $M = 5.16$, $SD = 8.00$, $N = 55,726$; American controls: $M = 2.91$, $SD = 6.84$, $N = 55,726$), $t(108,836) = 50.61$, $p < .001$, *Cohen's d* = 0.30 (for a robustness check, see Supplementary Materials).

To determine if polarized language is increasing over time, we sampled 1,507,300 tweets that were posted between November 2016 and December 2019 in the United States. These tweets were pulled randomly from BGU's SCI lab twitter database (sampling approach described in the Method section), with no specific text search. We applied the polarization dictionary to the text and aggregated by months. We conducted a weighted linear regression with monthly observations as the weighting factor. We found that Russian trolls used far more polarized language as time progressed ($b = 0.03$), $R^2 = .46$, $F(1, 69) = 58.85$, $p < .001$. Moreover, this was a strikingly large effect size. We did not find the same pattern among American control users, ($b = -0.001$) $R^2 = .05$, $F(1, 35) = 1.90$, $p = .178$ (see Figure 2). This suggests that trolls are increasing the use of polarized language much faster than ordinary users, independent groups correlation comparison $z = 5.06$, 95% CI [.55, 1.21], $p < .001$.

⁵ Shared in partnership with FiveThirtyEight on <https://github.com/fivethirtyeight/russian-troll-tweets>

Figure 2. Scatter plot of the average polarized score by Twitter sample. We examined monthly polarized language in American controls ($N = 1,507,300$; blue), and trolls from Russia ($N = 383,510$; red), Venezuela ($N = 85,970$; yellow), and Iran ($N = 329,453$; green). Values on the Y-axis represent the average percent of polarized language in the month. The size of the dots corresponds to the monthly sample size. Shaded areas around the regression line denote 95% CI. Note that the Y-axis is fixed to 0-5, data points exceeding this limit are not shown in the figure; the regression lines take these observations into account. Results indicate that trolls from Russia and Venezuela have been increasing their use of polarized rhetoric, but Americans have not.

This finding suggests Russian trolls have increased their use of polarized rhetoric, but the average U.S. Twitter does not show evidence of mirroring the type of language used by the trolls. This could be because trolls are only reaching and influencing the most politically active Twitter users, or that the average user expresses polarized attitudes in different ways. However, we note that the time frame for trolls and controls is not identical. As such, any differences in these trends should be treated as tentative. That said, in a *post-hoc* analysis conducted on the same time frame (November 23, 2016 - May 30, 2018), Again, Russian trolls used far more

polarized language as time progressed ($b = 0.03$), $R^2 = .51$, $F(1, 17) = 17.48$, $p < .001$ while American control users did not, ($b = 0.006$), $R^2 = .16$, $F(1, 17) = 3.14$, $p = .094$ (however note the small sample sizes in this analysis).

Study 2b: Polarization in Venezuelan and Iranian Trolls

We next sought to see if this pattern of polarized language generalized to other political contexts and countries. Given Russia's effort at online political warfare (Jensen et al., 2019), we also tested whether polarization attempts extended to other political actors. Russia, Iran, and Venezuela all hold anti-American views and share warm relationships with each other (Hakimzadeh, 2009; Katz, 2006; Moore, 2014). Therefore, these countries may have incentives to meddle with American politics. We analyzed trolls from these nations to see if they were using similar polarized rhetoric to sow conflict with Americans.

Results

We compared Twitter messages posted by Venezuelan and Iranian trolls (identified by Twitter¹) to a neutral American control sample. Again, we only used original tweets that were posted in the English language which were most likely aimed for an international/American audience. The paired comparisons were again matched for the same time range. In both countries we examined, the tweets sent by trolls used significantly more polarized language than tweets sent by American control samples ($ps < .001$), see Table 2. For the top 25 most used words adjusting for their frequency (tf-idf), see Figure S1 in the Supplementary Information.

Table 2. Means, SDs, sample sizes, and time range for each troll group comparison with American controls. The table consists of t statistics, degrees of freedom, and *Cohen's d*'. All the t-tests are significant at $p < .001$.

	Trolls		American Controls		Date Range	t	df	<i>Cohen's d</i>
	<i>Mean (SD)</i>	N	<i>Mean (SD)</i>	N				
Iran	2.15	220,628	1.46	929,908	2016-11-23 -	61.32	366,496	0.13
	(4.66)		(5.26)		2018-11-28			
Venezuela	1.80	30,987	1.45	953,197	2016-11-23 -	12.87	33,588	0.07
	(4.67)		(5.25)		2018-12-07			

Following the same analysis as in Study 2a, we conducted a weighted linear regression with monthly observations as the weighting factor, and found a diverging pattern between

populations of trolls: Whereas trolls based in Venezuela used more polarized language as time progressed ($b = 0.02$), $R^2 = .19$, $F(1, 91) = 20.91$, $p < .001$, Iranian trolls used less polarized language ($b = -0.03$), $R^2 = .33$, $F(1, 85) = 41.06$, $p < .001$, see Figure 2. Therefore, any trends in polarized language might be specific to the foreign nation involved.

Study 3: Exploratory Topics of Polarization

In Study 2 we showed that foreign agents from various countries strategically used polarized language in social media communications, and in a majority of cases we see an increase over time in these attempts. One remaining question is what specific forms of polarized rhetoric are leveraged by different groups. Investigating the themes associated with polarized language could shed light on the social-psychological processes capitalized by trolls, and generate better intuition on their strategies. Therefore, we conducted an exploratory analysis in which we decomposed the dictionary into different factors to determine whether these factors could contribute to our understanding of the trends in online polarized rhetoric.

As in Study 1, we used GloVe word embeddings as a high-dimensional representation of the words in the polarized dictionary (Pennington et al., 2014) and conducted hierarchical clustering analysis. We observed that in the highest level of division, the algorithm clustered the words in a manner that is somewhat consistent with the theoretical separation between *issue* and *affective* polarization (See Figure 3. and Figure S2.) See Supplementary Information for more information.

We applied the two subsets of the polarization dictionary on the social media messages posted by trolls and a random sample of American users. As in Study 2, we compared polarization levels between the groups (paired comparisons matched for the same time range). In all countries we examined, the tweets sent by trolls used significantly more polarized language than tweets sent by American control samples ($ps < .005$), both on affective and issue polarization, see Table 3 and Figure S3. Temporal analyses are reported in the Supplementary Information.

Table 3. Means, SDs, sample sizes, and time range for each troll group comparison with American controls by Issue and Affective polarization components. The table consists of t statistics, degrees of freedom, and *Cohen's d*. All the t -tests are significant at $p < .005$.

Trolls		American Control		Date Range	t	df	$Cohen's d$	
$Mean$ (SD)	N	$Mean$ (SD)	N					
Issue Polarization								
Russia	0.51 (2.35)	61,413	0.10 (1.00)	516,525	2016-11-23 - 2018-05-30	42.68	64,104	0.34
Iran	0.33 (1.70)	220,628	0.11 (1.08)	929,908	2016-11-23 - 2018-11-28	58.55	264,182	0.18
Venezuela	0.39 (1.99)	30,987	0.11 (1.07)	953,197	2016-11-23 - 2018-12-07	24.59	31,575	0.25
Affective Polarization								
Russia	1.86 (4.62)	61,413	1.37 (5.24)	516,525	2016-11-23 - 2018-05-30	24.53	81,500	0.09
Iran	1.82 (4.29)	220,628	1.34 (5.14)	929,908	2016-11-23 - 2018-11-28	44.69	386,294	0.09
Venezuela	1.41 (4.19)	30,987	1.34 (5.14)	953,197	2016-11-23 - 2018-12-07	2.84	34,087	0.01

In the current exploratory study, we showed that the polarization dictionary is composed of two subcomponents that map onto theoretical elements of polarization (Issue and Affective). In addition, we showed that all troll groups use more polarized language than a random sample of American social media users and that this holds for both affective and issue polarization (although effect sizes of issue polarization are substantially larger).

Study 4: Polarized Language and Engagement

Studies 2 and 3 demonstrated a link between polarized language use and troll accounts on social media. Indeed, our findings are consistent with the idea that trolls sow discord among Americans by using polarized language in conversations with others. Yet our results are agnostic to whether polarized language creates divisions vs. merely reflects an existing polarized state. To address this ambiguity, in the current study we investigate the extent to which polarized language is associated with increased engagement. If polarized language used by trolls draws more engagement from ordinary social media users, it would demonstrate that when trolls seed polarized language online, users become active agents in spreading the polarized messaging among groups.

Engagement of polarized language is an important metric because engagement with political content online is generally associated with high levels of ingroup bias; that is, it is far more likely to be shared within the political ingroup than in the outgroup (Barberá et al., 2015; Brady et al., 2017; Cinelli et al., 2021). For example, one of the key predictors of engagement on Twitter and Facebook in the political context is outgroup animosity (Rathje et al., 2021). Whether intentional or not, if polarized language used by Trolls is associated with increased engagement, it would suggest that Trolls' language use has potential to exacerbate division among political users (even if users were already divided).

Results

To test whether polarized language is associated with selective engagement in political discourse, we utilized the Russian Trolls Classification data (Linville & Warren, 2020) from Study 2. We only included observations for which we had an engagement metric (i.e., retweet count, $N = 118,215$). For ease of interpretation, we counted the number of words in the dictionary and used it to predict the retweet count in a negative binomial generalized linear model. We find that for every polarized word in a tweet, retweets increase by 72%, $IRR = 1.72$, $CI [1.64, 1.79]$, $p < .0001$. Next, we added an interaction by troll category. We conducted an Analysis of Deviance and found both main effects and interaction term significant (Polarized Language: $\chi^2(1) = 5263$, $p < .0001$; Troll Category: $\chi^2(7) = 373,563$, $p < .0001$; Polarized Language * Troll Category: $\chi^2(7) = 2223$, $p < .0001$). In a planned comparison we find that this effect stems directly from political trolls, Political vs. non Political trolls ratio = $\exp(1.35 \cdot 10^{15})$, 95% CI [$\exp(5.03 \cdot 10^{14})$, $\exp(2.2 \cdot 10^{15})$], $p < .0001$, see Figure 4.

Figure 4. Polarized language predicts retweets in political Russian trolls. The graph depicts the number of retweets predicted for a given tweet as a function of polarized language present in the tweet and type of troll. Bands reflect 95% CIs. For varying Y-axes, see Figure S5.

We take these results as evidence that polarized language is indeed polarizing, however, there is no reason to assume this effect applies strictly to trolls. We conducted the same analysis on samples of politically engaged controls (Study 2, $N = 55,726$), and a new sample of American controls for which we obtained engagement metrics ($N = 1,144,767$). Again, we find that in the politically engaged controls there is a positive association between polarized language and retweets, such that for every polarized word in a tweet, retweets increase by 39%, $IRR = 1.39$, 95% CI [1.35, 1.48], $p < .0001$. However, in a random sample of Americans we do not find a significant association $IRR = 1.19$, 95% CI [0.80, 1.77], $p = .390$.

We should note that these analyses are usually done with the number of followers as a covariate, yet retrospective information was only available for the trolls' dataset. For transparency, we show here the analysis controlling for the covariate. After adding followership in the trolls analysis we find the same pattern of results, however the effect size diminishes: $IRR = 1.61$, 95% CI [1.57, 1.67], $p < .0001$; planned contrasts: Political vs. non Political trolls ratio = $\exp(7.6 \cdot 10^9)$, CI [$\exp(3.6 \cdot 10^9)$, $\exp(1.21 \cdot 10^{10})$], $p < .0001$.

Overall, these results indicate that polarized language is associated with greater traction on social media, but only in political contexts. Since the probability of a political message to be retweeted within the political ingroup is far greater than the outgroup, we take this as evidence

that polarized language is not only a marker for a static polarized state, but contributes to the polarization process.

Discussion

We developed and validated a dictionary of polarized language used on social media. We validated this dictionary using three strategies and showed it consistently detected polarized discourse on Twitter and Reddit on multiple topics and corresponded well to the dynamics of partisan differences in attitudes towards the COVID-19 pandemic. We found that state-affiliated trolls from Russia and other countries use more polarized language than a random sample of American users and that while the language of Russian and Venezuelan trolls have used more polarized rhetoric with time, levels of polarized language in American controls did not increase. We found that our data-driven dictionary taps into distinct theoretical elements of polarization, and that trolls from all tested countries use more polarized rhetoric in both issue and affective factors (broadly denied). Lastly, we showed that polarized language is associated with more traction on social media, but only in political contexts; this finding suggests that polarized language advances polarization and not merely reflects it.

These results expand on prior work documenting trolls' attempts to pollute the online environment with polarized content and sow discord among Americans (Golovchenko et al., 2020). We provide novel evidence that this mission spans several countries that hold anti-American views. Prior research has revealed that when exploring the clusters of polarized topics, trolls are often found in the centroids of these clusters, driving the partisan discourse on both ends (Arif et al., 2018; Broniatowski et al., 2018; Walter et al., 2020). Our research extends these findings; we found that trolls share controversial content and engage in highly polarized issues, but that they also use higher levels of polarized language as a tool in their discourse. In addition, we found that polarized language is associated with greater engagement, however, this association only holds for politically engaged users – both trolls and controls. This is consistent with a view that trolls' use of polarized language is intended and weaponized in order to sow polarization, however our methods are not sufficient to draw such causality.

Questions remain as to the extent of influence of trolls' social media presence on real people. However, it is important to note that even a small number of agents with aggressive attitudes can have a substantial influence on the majority view, a process called “information gerrymandering” (Stewart et al., 2019). Exposure to polarizing attitudes even produced by a

small number of agents can have a devastating effect on political compromise in a social network; such findings suggest that trolls have the ability to influence many of the users on social networks. Furthermore, recent evidence suggests that troll's messages propagate to mainstream media and are represented as 'the voice of the people' (Lukito et al., 2019). This way, trolls win twice: once when they share the polarized content, and then again when it is being echoed on other media platforms, creating a polarizing loop.

However, some are skeptical of the change trolls may impose on people's attitudes. A recent paper followed over 1,200 American Twitter users for the course of one month in late 2017. The authors found that only a small fraction of users interacted with Russian trolls, and they did not observe any change in partisan attitude during that time among these users (Bail et al., 2020). In a study that explored the domestic effect of Russian trolls (i.e., messages that were targeted inwards to Russian users), it was found that trolls were trying to promote a pro-government agenda and dissolve government criticism (Sobolev, 2018); nevertheless, trolls were only successful at the latter, suggesting their influence is restricted in scope. While our results cannot speak to causal factors, we do find that while levels of polarized language were rising in Russian trolls, this was not the case among American users. Future research is required to understand the precise impact trolls have in reference to specific political events.

Given the evidence on the growing polarization and partisan antipathy in the American public (Iyengar et al., 2019), we also explored whether polarized discourse on social media would increase with time among a sample of American users. We did not find evidence to support this hypothesis; levels of polarization did not increase across time, suggesting that polarized discourse among average American users did not grow between November 2016 and December 2019. These results are consistent with other findings that do not find evidence for increased polarization during this brief time frame (Westwood et al., 2019). This could suggest that polarized discourse has not changed, that it has reached a plateau, or that American users' way of expressing polarized language has changed slightly over time. Discerning between these possibilities is an important endeavor for future research.

This paper also introduced the polarization dictionary and showcases its validation and application in studying political polarization. The dictionary is easy to use and can be utilized externally with LIWC (Pennebaker et al., 2015), or with the example code provided in the Supplementary Information for *R*. Having a quantifiable measure of polarized language in social

media messages is a quick way to estimate polarization levels that aligns with other current practices wherein researchers relied on computationally extensive network analyses, or narrowed down to a specific partisan topic to carry out their studies.

The current study has several limitations. The polarization dictionary has been built on data collected in 2015 and on three polarized topics. Therefore, it is subjected to bias about topics that were timely in 2015 and is potentially restricted in its scope. We attempt to get around this limitation by expanding the lexicon using word-embeddings and testing its validation over multiple time periods. Nonetheless, language is highly dynamic on social media and our dictionary should always be validated when applied to a new context. Given its data-driven development, it also includes some terms that may not seem strictly polarized (e.g., *people*). Therefore, if being used by other researchers, we recommend using it comparatively by having a baseline corpus and measuring amounts of polarized language between groups to get a relative estimate.

One potential issue is with the authenticity of early social media accounts identified as trolls. Some countries use hacked, purchased, or stolen accounts. Early data, therefore, may not have originated with the nation in question. While this was probably not the case with the Russian trolls dataset, it could be the case with some Venezuelan or Iranian content, and may have biased our polarization over time analyses. That said, we employed a weighted regressions analysis that takes into account the relatively sparse nature of early messages (and therefore downweights their importance). These analyses complement the Russian sample and provide a wider, descriptive view of how different troll populations use polarized language.

In addition, this work has focused primarily on quasi-experimental manipulations or correlational methodology. Future work should examine if there are causal factors that increase or decrease polarization. For instance, given the potential influence that the design of social media can have on moralized language (Brady et al., 2020), it is possible that specific design feature changes could impact polarization language. For instance, down-weighting polarized language on social media news feeds might influence attitudes such as partisan antipathy.

Conclusion

Taken together, this research offers a tool to detect and understand the use of polarized rhetoric on social media. In times when it seems like we have reached toxic levels of polarization

in America it is increasingly important to continually develop tools to study and combat the potentially polarizing influence of foreign agents in American politics.

References

- Allcott, H., Braghieri, L., Eichmeyer, S., & Gentzkow, M. (2020). The Welfare Effects of Social Media. *The American Economic Review*, 110(3), 629–676.
- Amira, K., Wright, J. C., & Goya-Tocchetto, D. (2019). In-Group Love Versus Out-Group Hate: Which Is More Important to Partisans and When? *Political Behavior*, 1–22.
- Arif, A., Stewart, L. G., & Starbird, K. (2018). Acting the part: Examining information operations within# BlackLivesMatter discourse. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), 20.
- Badawy, A., Ferrara, E., & Lerman, K. (2018). Analyzing the Digital Traces of Political Manipulation: The 2016 Russian Interference Twitter Campaign. *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 258–265.
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences of the United States of America*, 115(37), 9216–9221.
- Bail, C. A., Guay, B., Maloney, E., Combs, A., Hillygus, D. S., Merhout, F., Freelon, D., & Volfovsky, A. (2020). Assessing the Russian Internet Research Agency’s impact on the political attitudes and behaviors of American Twitter users in late 2017. *Proceedings of the National Academy of Sciences of the United States of America*, 117(1), 243–250.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science*, 26(10), 1531–1542.

- Barrie, C., & Ho, J. (2021). academictwitteR: an R package to access the Twitter Academic Research Product Track v2 API endpoint. *Journal of Open Source Software*, 6(62), 3272.
- Baumann, F., Lorenz-Spreen, P., Sokolov, I. M., & Starnini, M. (2020). Modeling Echo Chambers and Polarization Dynamics in Social Networks. *Physical Review Letters*, 124(4), 048301.
- Baumgartner, J., Zannettou, S., Keegan, B., Squire, M., & Blackburn, J. (2020). The Pushshift Reddit Dataset. *Proceedings of the International AAAI Conference on Web and Social Media*, 14, 830–839.
- Benoit, K., Watanabe, K., Wang, H., Nulty, P., Obeng, A., Müller, S., & Matsuo, A. (2018). quanteda: An R package for the quantitative analysis of textual data. *Journal of Open Source Software*, 3(30), 774.
- Benson, T. (2020, April 24). Trolls and bots are flooding social media with disinformation encouraging states to end quarantine. *Business Insider*.
<https://www.businessinsider.com/trolls-bots-flooding-social-media-with-anti-quarantine-disinformation-2020-4>
- Boxell, L., Gentzkow, M., & Shapiro, J. M. (2017). Greater Internet use is not associated with faster growth in political polarization among US demographic groups. *Proceedings of the National Academy of Sciences of the United States of America*, 114(40), 10612–10617.
- Boxell, L., Gentzkow, M., & Shapiro, J. M. (2020). *Cross-Country Trends in Affective Polarization* (No. 26669). National Bureau of Economic Research.
<https://doi.org/10.3386/w26669>
- Brady, W. J., Crockett, M. J., & Van Bavel, J. J. (2020). The MAD Model of Moral Contagion: The Role of Motivation, Attention, and Design in the Spread of Moralized Content Online.

- Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 15(4), 978–1010.
- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(28), 7313–7318.
- Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., Quinn, S. C., & Dredze, M. (2018). Weaponized Health Communication: Twitter Bots and Russian Trolls Amplify the Vaccine Debate. *American Journal of Public Health*, 108(10), 1378–1384.
- Carothers, T., & O'Donohue, A. (2019). *Democracies divided: The global challenge of political polarization*. Brookings Institution Press.
- Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences of the United States of America*, 118(9). <https://doi.org/10.1073/pnas.2023301118>
- Cosentino, G. (2020). Polarize and Conquer: Russian Influence Operations in the United States. In G. Cosentino (Ed.), *Social Media and the Post-Truth World Order: The Global Dynamics of Disinformation* (pp. 33–57). Springer International Publishing.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences of the United States of America*, 113(3), 554–559.
- Demszky, D., Garg, N., Voigt, R., Zou, J., Shapiro, J., Gentzkow, M., & Jurafsky, D. (2019). Analyzing Polarization in Social Media: Method and Application to Tweets on 21 Mass Shootings. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*

(*Long and Short Papers*), 2970–3005.

Evans, T., & Fu, F. (2018). Opinion formation on dynamic networks: identifying conditions for the emergence of partisan echo chambers. *Royal Society Open Science*, 5(10), 181122.

Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., McGrath, M. C., Nyhan, B., Rand, D. G., Skitka, L. J., Tucker, J. A., Van Bavel, J. J., Wang, C. S., & Druckman, J. N. (2020). Political sectarianism in America. *Science*, 370(6516), 533–536.

Garimella, K., Morales, G. D. F., Gionis, A., & Mathioudakis, M. (2018). Quantifying Controversy on Social Media. *ACM Transactions on Social Computing*, 1(1), 1–27.

Gollwitzer, A., Martel, C., Brady, W. J., Pärnamets, P., Freedman, I. G., Knowles, E. D., & Van Bavel, J. J. (2020). Partisan differences in physical distancing are linked to health outcomes during the COVID-19 pandemic. *Nature Human Behaviour*, 4(11), 1186–1197.

Golovchenko, Y., Buntain, C., Eady, G., Brown, M. A., & Tucker, J. A. (2020). Cross-Platform State Propaganda: Russian Trolls on Twitter and YouTube during the 2016 U.S. Presidential Election. *The International Journal of Press/Politics*, 1940161220912682.

Grimmer, J., King, G., & Superti, C. (2014). You Lie! Patterns of Partisan Taunting in the US Senate. *Society for Political Methodology*.
https://scholar.harvard.edu/files/gking/files/polmethposter_1.pdf

Guerra, P. C., Meira, W., Jr, Cardie, C., & Kleinberg, R. (2013). A measure of polarization on social media networks based on community boundaries. *Seventh International AAAI Conference on Weblogs and Social Media*.
<https://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/viewPaper/6104>

Guess, A. (2020). (Almost) everything in moderation: New evidence on Americans' online media diets. *American Journal of Political Science*.

https://www.dropbox.com/s/3rjsnp8k3im7377/AGuess_OMD_AJPS.pdf?dl=0

Hakimzadeh, K. (2009). Iran & Venezuela: The“ Axis of Annoyance.” *Military Review*, 89(3), 78.

Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The Origins and Consequences of Affective Polarization in the United States. *Annual Review of Political Science*, 22(1), 129–146.

Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, Not Ideology A Social Identity Perspective on Polarization. *Public Opinion Quarterly*, 76(3), 405–431.

Jasny, L., Dewey, A. M., Robertson, A. G., Yagatich, W., Dubin, A. H., Waggle, J. M., & Fisher, D. R. (2018). Shifting echo chambers in US climate policy networks. *PloS One*, 13(9), e0203463.

Jensen, B., Valeriano, B., & Maness, R. (2019). Fancy bears and digital trolls: Cyber strategy with a Russian twist. *Journal of Strategic Studies*, 42(2), 212–234.

Katz, M. N. (2006). The Putin-Chavez Partnership. *Problems of Post-Communism*, 53(4), 3–9.

Klein, E. (2020). *Why we're polarized*.

Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096.

Levy, R. 'ee. (2021). Social Media, News Consumption, and Polarization: Evidence from a Field Experiment. *The American Economic Review*, 111(3), 831–870.

Linville, D. L., & Warren, P. L. (2020). Troll Factories: Manufacturing Specialized Disinformation on Twitter. *Political Communication*, 37(4), 447–467.

- Llewellyn, C., Cram, L., Favero, A., & Hill, R. L. (2018). Russian troll hunting in a brexit Twitter archive. *Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries*, 361–362.
- Lukito, J., Suk, J., Zhang, Y., Doroshenko, L., Kim, S. J., Su, M.-H., Xia, Y., Freelon, D., & Wells, C. (2019). The Wolves in Sheep's Clothing: How Russia's Internet Research Agency Tweets Appeared in U.S. News as Vox Populi. *The International Journal of Press/Politics*, 1940161219895215.
- Mason, L. (2018). *Uncivil agreement: How politics became our identity*. University of Chicago Press.
- Moore, E. D. (2014). *Russia-Iran relations since the end of the Cold War*. Routledge.
- Mukerjee, S., Jaidka, K., & Lelkes, Y. (2020). *The Ideological Landscape of Twitter: Comparing the Production versus Consumption of Information on the Platform*.
<https://doi.org/10.31219/osf.io/w98ms>
- Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). *The Development and Psychometric Properties of LIWC2015*.
https://repositories.lib.utexas.edu/bitstream/handle/2152/31333/LIWC2015_LanguageManual.pdf
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Pintak, L., Albright, J., Bowe, B. J., & Pasha, S. (2019). *#Islamophobia: Stoking Fear and Prejudice in the 2018 Midterms*. Social Science Research Council.

<https://www.ssrc.org/publications/view/islamophobia-stoking-fear-and-prejudice-in-the-2018-midterms/>

- Rathje, S., Van Bavel, J. J., & van der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences of the United States of America*, 118(26). <https://doi.org/10.1073/pnas.2024292118>
- Sikder, O., Smith, R. E., Vivo, P., & Livan, G. (2020). A minimalistic model of bias, polarization and misinformation in social networks. *Scientific Reports*, 10(1), 5493.
- Simchon, A., Guntuku, S. C., Simhon, R., Ungar, L. H., Hassin, R. R., & Gilead, M. (2020). Political depression? A big-data, multimethod investigation of Americans' emotional response to the Trump presidency. *Journal of Experimental Psychology. General*, 149(11), 2154–2168.
- Sobolev, A. (2018). How pro-government “trolls” influence online conversations in Russia. *University of California Los Angeles*. [7]. <http://www.wpsanet.org/papers/docs/2019W-Feb-Anton-Sobolev-Trolls-VA.pdf>
- Soliman, A., Hafer, J., & Lemmerich, F. (2019). A Characterization of Political Communities on Reddit. *Proceedings of the 30th ACM Conference on Hypertext and Social Media*, 259–263.
- Starnini, M., Frasca, M., & Baronchelli, A. (2016). Emergence of metapopulations and echo chambers in mobile agents. *Scientific Reports*, 6, 31834.
- Stewart, A. J., Mosleh, M., Diakonova, M., Arechar, A. A., Rand, D. G., & Plotkin, J. B. (2019). Information gerrymandering and undemocratic decisions. *Nature*, 573(7772), 117–121.
- Sunstein, C. R. (2018). *# Republic: Divided democracy in the age of social media*. Princeton University Press.
- Twitter. (2018). Internet Research Agency (October 2018) [Data set]. In *Twitter Elections*

- Integrity Datasets*. <https://transparency.twitter.com/en/reports/information-operations.html>
- Twitter. (2019a). Iran (January 2019) [Data set]. In *Twitter Elections Integrity Datasets*. <https://transparency.twitter.com/en/reports/information-operations.html>
- Twitter. (2019b). Venezuela (January 2019, set 1) [Data set]. In *Twitter Elections Integrity Datasets*. <https://transparency.twitter.com/en/reports/information-operations.html>
- Van Bavel, J. J. (2020, March 23). In a pandemic, political polarization could kill people. *The Washington Post*. <https://www.washingtonpost.com/outlook/2020/03/23/coronavirus-polarization-political-exaggeration/>
- Van Bavel, J. J., Harris, E. A., Pärnamets, P., Rathje, S., Doell, K. C., & Tucker, J. A. (2021). Political psychology in the digital (mis)information age: A model of news belief and sharing. *Social Issues and Policy Review*, 15(1), 84–113.
- Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *Proceedings of the National Academy of Sciences of the United States of America*, 113(23), 6454–6459.
- Van Bavel, J. J., Rathje, S., Harris, E., Robertson, C., & Sternisko, A. (2021). How social media shapes polarization. *Trends in Cognitive Sciences*, 25(11), 913–916.
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151.
- Walter, D., Ophir, Y., & Jamieson, K. H. (2020). Russian Twitter Accounts and the Partisan Polarization of Vaccine Discourse, 2015-2017. *American Journal of Public Health*, e1–e7.
- Warner, B. R. (2010). Segmenting the Electorate: The Effects of Exposure to Political Extremism Online. *Communication Studies*, 61(4), 430–444.
- Westwood, S. J., Peterson, E., & Lelkes, Y. (2019). Are there Still Limits on Partisan Prejudice?

Public Opinion Quarterly, 83(3), 584–597.

Wilson, A. E., Parker, V., & Feinberg, M. (2020). Polarization in the contemporary political and media landscape. *Current Opinion in Behavioral Sciences*, 34, 223–228.

Yardi, S., & Boyd, D. (2010). Dynamic Debates: An Analysis of Group Polarization Over Time on Twitter. *Bulletin of Science, Technology & Society*, 30(5), 316–327.

Supplementary Information for

Troll and Divide: The Language of Online Polarization

This document includes:

Supplementary Text

Tables. S1, S2

Figures. S1- S5

References

For the dictionary, full code and analysis see OSF repository:

<https://osf.io/bm8uy>

Study 1: Development of a Polarization Dictionary

Dictionary Development

We built on the work by Brady et al. (2017) which studied the diffusion of controversial political content on Twitter during discussions of climate change, same-sex marriage, and gun control. Their dataset included 24,849 tweets with available information on whether the tweets were polarized (retweeted within one political community), or not (retweeted by a user from the opposing ideology). We performed a differential language analysis, a procedure in which two groups are compared in their frequency of word use (Schwartz et al., 2013), on 80% of their data (the rest was kept for validation). We compared the word use of the polarized cluster vs. the non-polarized cluster by calculating a chi-square statistic for every word in the data set; resulting in a shortlist of words that were significantly associated with polarization.

In the second step, we manually pruned the list of words (i.e., the dictionary) by filtering out names of individuals (e.g. *Bernie*) and topical words (e.g., *antarctic*) that would be unlikely to generalize to other contexts outside the original research (Brady et al., 2017). The full dictionary ($N = 256$) was judged independently for pruning by two of the authors (A.S. and W.J.B.) and agreement reached a Cohen's κ of 0.61, $z = 9.93$, $p < .001$, 95% CI [.51, .72]. Remaining disagreements were discussed to reach convergence. Words that were associated with depolarization were removed. The pruned version of the dictionary consisted of 57 words.

Next, we used word embeddings, a vectorized representation of words that encompasses semantic fields, to expand the pruned dictionary. This process helped the dictionary capture a greater linguistic space while staying close to the semantic space implied from the dictionary. The GloVe algorithm (Pennington et al., 2014) utilizes word co-occurrence in large corpora to create embeddings of 200 dimensions. We used a pre-trained GloVe model by Stanford NLP which was built on 2 billion tweets (<https://github.com/stanfordnlp/GloVe>) to extract the five most semantically-related words to each of the “seed” words from the prior step. For example, the word *threat* was expanded by the words *threats*, *attacks*, *terrorism*, *targets*, and *threatening*. The fully expanded dictionary contained 232 words.

In the final step, we trimmed proper names (e.g., *Obama*) and nonsensical additions (e.g., *prettylittleliars*). This time there was perfect agreement between the raters in applying the two rules, which resulted in the removal of 27 words (the final dictionary contained 205 words; the word lists with raters agreement are found on the Online Repository).

Internal Consistency

Conducting psychometric assessments of dictionaries is a well-known issue in text analysis (Pennebaker et al., 2007). Especially in the context of social media and even more so when using Twitter data, it is important to understand what is the unit of analysis in the psychometric evaluation. To conduct an analysis of internal consistency, we grouped together tweets of the same authors. Originally our training set consisted of 19,841 tweets. After grouping tweets together by authors, the training corpus consisted of 7,963 observations. To assess internal consistency in the *binary method* (Pennebaker et al., 2007), we calculated a binary occurrence matrix of the dictionary elements wherein each word in the dictionary is considered an item in the “questionnaire” (i.e., the dictionary), and calculated Cronbach’s *alpha* of 0.75, 95% CI [0.75,0.76].

Dictionary Validation

Reddit Analysis

We extracted reddit comments from 36 politically mapped subreddit (Soliman et al., 2019). The list of subreddits and their political orientation is shown in Table S2.

Since many comments on Reddit do not contain more than a title, we combined the title and the body of the message into a unified text variable. We then removed links and emoticons and filtered out deleted or removed messages. Messages in languages other than English were removed as well. Reddit messages were collected through the Pushshift API and using the *rreddit* R package (Kearney, 2019).

Results. We applied the dictionary on the Reddit sample (political left, political right and control group) and conducted a one-way between-group ANOVA. Results show a significant effect of political group $F(2,49227) = 610.65, p < .001, \eta_p^2 = .024$, which was followed by a planned comparison reported in the main text. The second analysis included a neutral sample (*NeutralPolitics*) instead of control messages collected from a random sample of popular communities. We applied a one-way between-group ANOVA. As before, results show a significant effect of political group $F(2,42633) = 14.51, p < .001, \eta_p^2 < .001$, which was followed by a planned comparison reported in the main text.

Studies 2 and 3: Term frequency-inverse document frequency analysis

To better understand the type of language that drives differences in polarized language between trolls and American controls, we conducted a term frequency-inverse document frequency (tf-idf) analysis; a statistical procedure that marks the word importance in a corpus, based on comparing a word's frequency with its base-rate usage. We selected only the words that appear in the polarization dictionary and ranked them by their tf-idf value. Figure. S1 displays the top 25 polarized words based on their tf-idf value.

Results

We used the Russian Troll classification ([Linvill & Warren, 2020](#)), and matched an American sample for their content (via hashtag use, see Method section), posting time (January 2015 - May 2018) and quantity. A closer look reveals that one of the popular hashtags which was used for content-matching may not be a good sampling decision (#NowPlaying). Therefore, as a robustness check we decided to analyze the data with the exclusion of tweets containing this hashtag. The results continue to show more polarized language in politically-oriented Russian trolls vs. politically-matched Americans $t(103,528) = 39.48$, $p < .001$, *Cohen's d* = 0.24.

Study 4: Exploratory Topics of Polarization

Hierarchical Clustering

To conduct a thematic clustering of the polarization dictionary, we extracted GloVe word embeddings (Pennington et al., 2014). We then lemmatized the words in the dictionary, and for every word that shared a lemma, we took the average embedding of that lemma, resulting in 170 words in total for clustering. Next, we conducted hierarchical clustering analysis and cut the clustered at the highest level of division (2). See dendrogram in Figure S2.

Results

We applied the two subsets of the polarization dictionary on the social media messages posted by trolls and a random sample of American users across time. As in Studies 2 and 3, we calculated monthly polarization scores and conducted a weighted linear regression predicting polarized language as a function of time, dictionary subcomponent and their interaction with monthly observations as the weighting factor. We were interested in whether the slope of the two

dictionary components differ in each group. While there were no significant interactions in the Russian or Venezuelan groups, we found that in American controls, issue polarization had a positive slope, however not significant $b = 0.0004$, $SE = 0.0005$, 95% CI [-0.0005,0.0016], while affective polarization had a significant negative slope $b = -0.001$, $SE = 0.0005$, 95% CI [-0.0024,-0.0003], resulting in a significant slope difference $b = -0.001$, $SE = 0.0072$, $t(70) = -2.55$ $p = .013$. We also found that in Iranian trolls, both affective $b = -0.0207$, $SE = 0.0026$, 95% CI [-0.0259,-0.0155] and issue polarization $b = -0.0059$, $SE = 0.0026$, 95% CI [-0.0111,-0.0007] had significant negative slopes, which differ significantly from each other $b = 0.0148$, $SE = 0.0036$, $t(170) = -4.015$ $p < .001$, see Figure S3.

In the current exploratory study, we showed that the polarization dictionary is composed of different subcomponents that map onto theoretical elements of polarization. In addition, we show that the lack of significant polarization trend in American controls, could be attributed to the different trends in affective and issue polarization. On a closer look, affective polarization showed a significant negative trend, however further inspection revealed the trend is driven by a relatively high value which was given the most weight, namely August 2017. When omitted from the analysis, the negative trend was no longer significant $b = -0.001$, $SE = 0.0005$, 95% CI [-0.0017, 0.0001].

Interestingly, in August 2017 the United States had experienced one of most contentious events in its recent history. “Unite the Right” rally in Charlottesville, Virginia was an exemplar of a hyper-polarized event, resulting in a white supremacist killing one person and injuring 19 other people (Tien et al., 2020). Therefore, while contributing to a potentially inaccurate trend, high levels of affective polarization in August 2017 do make sense given the context.

Table S1. Correlation table between poll responses and lagged twitter language. Adjusted for multiple comparisons using the Holm method.

Lag	<i>r</i>	CI low	CI high	<i>t</i>	<i>df</i>	<i>p</i>
1	0.47	0.21	0.66	3.62	47	0.076
2	0.48	0.23	0.67	3.75	46	0.054
3	0.50	0.25	0.69	3.89	45	0.036
4	0.52	0.28	0.71	4.08	44	0.021
5	0.56	0.32	0.73	4.45	43	0.007
6	0.61	0.38	0.77	4.99	42	0.001
7	0.66	0.44	0.80	5.57	41	<.001
8	0.67	0.46	0.81	5.70	40	<.001
9	0.67	0.46	0.81	5.63	39	<.001
10	0.66	0.44	0.81	5.47	38	<.001
11	0.64	0.41	0.79	5.07	37	0.001
12	0.62	0.38	0.79	4.78	36	0.003
13	0.61	0.36	0.78	4.55	35	0.007
14	0.60	0.34	0.77	4.36	34	0.013
15	0.58	0.30	0.76	4.06	33	0.032
16	0.54	0.25	0.74	3.63	32	0.104

Table S2. List of known politically leaning subreddits, adapted from Soliman et al. (2019).

Subreddit	Political Leaning
BlueMidterm2018	Left
CaliforniaForSanders	Left
DepthHub	Left
Enough_Sanders_Spam	Left
esist	Left
FriendsofthePod	Left
GrassrootsSelect	Left
GreenParty	Left
justicedemocrats	Left
Keep_Track	Left
Kossacks_for_Sanders	Left
LateShow	Left
Maher	Left
occupywallstreet	Left
Political_Revolution	Left
PoliticalDiscussion	Left
progressive	Left
RussiaLago	Left
altright	Right
AskThe_Donald	Right
CBTS_Stream	Right
DarkEnlightenment	Right
DrainTheSwamp	Right
europeannationalism	Right
greatawakening	Right
hottiesfortrump	Right
kekistan	Right
Le_Pen	Right
Mr_Trump	Right
Physical_Removal	Right
redacted	Right
The_Congress	Right
The_Europe	Right
The_Farage	Right
tucker_carlson	Right
WhiteRights	Right

Figure. S1. Term frequency-inverse document frequency of the top 25 polarized words, by Twitter sample. These are the top polarized words in each sample.

Figure S2. Dendrogram of the hierarchical relationship in the hierarchical clustering analysis, based on 200 dimensions GloVe embeddings.

Figure. S3. Polarization score by population (American controls, Russian trolls, Iranian trolls, Venezuelan trolls) and polarization components (Issue and Affective). Points denote means; error bars denote 95% confidence intervals. All comparisons were matched on timeframe.

Figure. S4. Scatter plot of the average polarized subcomponent (Affective and Issue) by Twitter sample. Values on the Y-axis represent the average percent of polarized language in the month. Shaded areas around the regression line denote 95% CI. The size of the dots corresponds to the monthly sample size. Note that the Y-axis is fixed to 0-5, data points exceeding this limit are not shown in the figure; the regression lines take these observations into account.

Figure S5. Polarized language predicts retweets in political Russian trolls. The graph depicts the number of retweets predicted for a given tweet as a function of polarized language present in the tweet and type of troll. Bands reflect 95% CIs. For constant Y-axes, see Figure 4.

References

- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(28), 7313–7318.
- Kearney, M. W. (2019). *rreddit*. Github. Retrieved November 3, 2020, from <https://github.com/mkearney/rreddit>
- Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., & Booth, R. J. (2007). The Development and Psychometric Properties of LIWC2007 (LIWC2007 Manual). *Austin, TX: LIWC*.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E. P., & Ungar, L. H. (2013). Personality, gender, and age in the language of social media: the open-vocabulary approach. *PloS One*, 8(9), e73791.
- Soliman, A., Hafer, J., & Lemmerich, F. (2019). A Characterization of Political Communities on Reddit. *Proceedings of the 30th ACM Conference on Hypertext and Social Media*, 259–263.
- Tien, J. H., Eisenberg, M. C., Cherng, S. T., & Porter, M. A. (2020). Online reactions to the 2017 “Unite the right” rally in Charlottesville: measuring polarization in Twitter networks using media followership. *Applied Network Science*, 5(1), 10.