

The advantage of the right in social media news sharing

Sandra González-Bailón ^{a,*}, Valeria d'Andrea ^b, Deen Freelon^c and Manlio De Domenico ^{b,d,*}

^aAnnenberg School for Communication, University of Pennsylvania, Philadelphia, PA 19104, USA

^bComplex Multilayer Networks Lab, Center for Digital Society, Fondazione Bruno Kessler, Trento 38123, Italy

^cHussman School of Journalism and Media, University of North Carolina at Chapel Hill, Chapel Hill, NC 27514, USA

^dDepartment of Physics & Astronomy, University of Padua, Padova 35131, Italy

*To whom correspondence should be addressed: Email: sgonzalezbailon@asc.upenn.edu; manlio.dedomenico@unipd.it

Edited By: Karen E. Nelson

Abstract

We analyze social media activity during one of the largest protest mobilizations in US history to examine ideological asymmetries in the posting of news content. Using an unprecedented combination of four datasets (tracking offline protests, social media activity, web browsing, and the reliability of news sources), we show that there is no evidence of unreliable sources having any prominent visibility during the protest period, but we do identify asymmetries in the ideological slant of the sources shared on social media, with a clear bias towards right-leaning domains. These results support the “amplification of the right” thesis, which points to the structural conditions (social and technological) that lead to higher visibility of content with a partisan bent towards the right. Our findings provide evidence that right-leaning sources gain more visibility on social media and reveal that ideological asymmetries manifest themselves even in the context of movements with progressive goals.

Keywords: news sharing, misinformation, online networks, political mobilization, computational social science

Significance Statement:

Existing research suggests that left- and right-wing activists use different media to achieve their political goals: the former operate on social media through hashtag activism, and the latter partner with partisan outlets. However, legacy and digital media are not parallel universes. Sharing mainstream news in social media offers one prominent conduit for content spillover across channels. We analyze news sharing during a historically massive racial justice mobilization and show that misinformation posed no challenge to the coverage of these events. However, links to outlets with a partisan bent towards the right were shared more frequently, which suggests that right-leaning outlets have higher reach even within the confines of online activist networks built to enact change and oppose dominant ideologies.

Introduction

On 2020 June 6, people across the United States joined one of the largest mobilizations in the country's protest history (1). The immediate trigger was the killing of George Floyd while in police custody, but the scale of the mobilizations reflected years of organizing by a decentralized political movement seeking to end police brutality and advocate for criminal justice reform (2). The movement, initiated in 2014 around the #BlackLivesMatter (BLM) hashtag, has become a global symbol of social justice, and a prominent example of a new form of digital activism that defies old forms of organizing. Online networks, and the flows of information they allow, are the backbone of this type of mobilization: these networks help activists create alternative spaces in which to articulate discourses that are excluded for mainstream media (3). Of course, these tools are also available to other types of activism, including less progressive forces. Recent scholarship has started to pay attention to the asymmetries that characterize different forms of digital activism (4), but there are still many unknowns about how different actors use online tools as

part of their repertoire, or how delimited different publics are online (5).

One idea that is receiving increasing empirical support is that the political left and the political right use media in different ways. Online media on the left and the right face different incentive structures and constraints that lead to different architectures and susceptibility to propaganda (6). This phenomenon, known as ideological asymmetry, manifests in a variety of ways. Adherents of the left, the argument goes, tend to consume more mainstream media (7), trust fact-checkers more (8), and curate more diverse personalized information environments (9) than right-wingers. The left engages heavily in hashtag activism—the use of social media hashtags to brand a political cause both on- and offline—but the lack of research on right-wing hashtag activism makes ideological comparisons in this area difficult (4). People on the right tend to trust mainstream media less (10), spread more false news (6, 11), and tolerate the spreading of misinformation by politicians more than the left (12). Right-wing distrust of both the mainstream media and the set of internet platforms known

Competing Interest: The authors declare no competing interest.

Received: March 30, 2022. **Accepted:** July 27, 2022

© The Author(s) 2022. Published by Oxford University Press on behalf of the National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial reproduction and distribution of the work, in any medium, provided the original work is not altered or transformed in any way, and that the work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

as “Big Tech” has led followers to opt in to the “right-wing media ecosystem” of more congenial outlets (7) as well as “alt-tech” social media that offer more permissive terms of use (13).

While existing work on ideological asymmetry tells us much about how the left and right differ in terms of their distinctive styles of online political engagement, less is known about what happens when the two sets of tactics collide. Social media are sites of political contestation, and individuals on opposite sides of the same issue often clash directly and indirectly for attention, resources, and ideological converts. A key question here is: When left and right clash online, who is more successful in spreading their message? On the one hand, proponents of what might be called the “advantage of the right” thesis point to the structural conditions that allow better message production and dissemination—conditions that take the form of money and free time (14) or, in the digital realm, algorithmic amplification, which gives right-leaning content more visibility (15); and they also point to the greater audience susceptibility and engagement with moralized content (16). On the other hand, those predicting greater prevalence and circulation of left-wing messages can point to the fact that left-leaning users outnumber the right on Twitter (17, 18), and that left-wing hashtag activism campaigns such as #BLM and #MeToo have gained massive success (2, 3). These two realities co-exist online. For example, in the early days of the BLM movement (2014 to 2015), left-wing, anti-police brutality voices far outnumbered the pro-police right on Twitter (2). But on the issue of mass shootings, messages supporting gun rights on Twitter are more numerous than those advocating gun control (19). So far, little if any research has investigated the question of how such ideological contests shape the visibility of political messages on social media. Here, we address this question in the context of the BLM movement.

Our analyses consider three outstanding questions regarding ideological asymmetries in contested online spaces: First, whose messages reach the most people? Second, which ideological side consumes the most low-quality content (i.e. misinformation)? And third, how common in the conversation is low-quality content in terms of prevalence and consumption? Our results support the “amplification of the right” thesis, finding that content with a right-leaning partisan slant is viewed and shared substantially more than left-leaning content. This gives right-leaning outlets an advantage to the extent that they gain higher visibility and levels of engagement, both crucial in the attention economy that characterizes digital media. Unreliable content, on the other hand, is very infrequent and low in visibility overall: only a very small community of Twitter users share unreliable sources. But to the extent that right-leaning sources are, on the aggregate, more visible, our results also suggest that activist networks are seeded with messages that are not necessarily aligned with their framing of events, limiting the impact that online activism can have on public discourse.

Data

Our data come from four different sources: Twitter’s Application Programming Interface (API); records of offline protest events; web tracking data; and reliability ratings for news websites. We use Twitter data to reconstruct social media activity around the 2020 June 6 mobilizations, and protest event data to build a benchmark to compare online activity with offline actions. Figure 1 shows a comparison of the timeline and location of protest events according to the Crowd Counting Consortium (20) with the timeline and location of the tweets, we analyze (see Figure S2 in the

“Supplementary Material” section for an alternative source of protest events data). There are obvious parallels both in terms of volume (the peak in online activity coincides with the peak of protest events on June 6) and in terms of geographic distribution in the contiguous United States (see the “Supplementary Material” section for more details on the data and alternative counts of offline protest events). To differentiate different types of Twitter users, we classify accounts in three groups: media accounts, bots, and human accounts. This classification relies on bot detection techniques (see the “Materials and methods” and the “Supplementary Material” sections for technical details). Building on prior work (21, 22), we used the label “media” as a shorthand to refer to accounts with bot-like behavior that are also verified by Twitter. These include the accounts of public figures, journalists, or news organizations. We labelled as “bots” automated accounts that are not verified by Twitter, and the rest are classified as human accounts. Figure 1C shows that most messages are generated by human accounts, but a large fraction of the total volume is generated by unverified bots. Less than 1% of all accounts fall in the “media” category. As we show in the Supplementary Material, this categorization also reveals other expected differences (for instance, verified accounts post more reliable content, see Figure S15).

We parsed all the tweets to extract URLs, when present, and we identified their registered domain. This yielded a list of $N = 2,176$ unique domains. We matched these domains with web browsing data collected during the same period from a representative panel of the US population (see the “Materials and methods” section). This web panel allowed us to obtain measures of audience reach (i.e. the fraction of the online population accessing the domains) and the partisanship composition of these audiences (i.e. the fraction of users accessing the domain that self-identify as Republican or Democrat). We use party identification to compute a variable of ideological slant, derived from the partisan sorting of the audiences accessing those domains. By looking at the composition of audiences along partisan divides, we are differentiating domains that are favored by different types of people. Our assumption is that audience composition tells us something about the slant of the coverage provided by the domains (i.e. news sources). This is an assumption that we share with prior research (e.g. 23, 24–26) but that may yield different classifications compared to content-based or editorial measures of ideological bias (e.g. those provided by AdFontes or AllSides). With this measure of audience composition, we compute ideological scores that we assign to domains and to the Twitter users posting those domains. We also matched the domains with reliability scores that rate the credibility of news and information websites. Sites receive a trust score on a 0 to 100 scale based on criteria related to the credibility of the information published and the transparency of the sources (see the “Materials and methods” and the “Supplementary Material” sections for more details).

Results

Most URLs shared during the protest mobilizations lead to “News/Information” domains. As Figure 2A shows, this is by far the most popular category (see Table S1 for examples of top domains within these groups). Within the news category, local news prevail (panel A inset). Panel B plots the distribution of the party scores for all domains. Most of these sites have audiences that include roughly the same number of Democrats and Republicans. Panel C shows a moderate correlation between audience reach on the web and number of tweets pointing to the domains (see Figures S5 to S8 in the “Supplementary Material” section for more details).

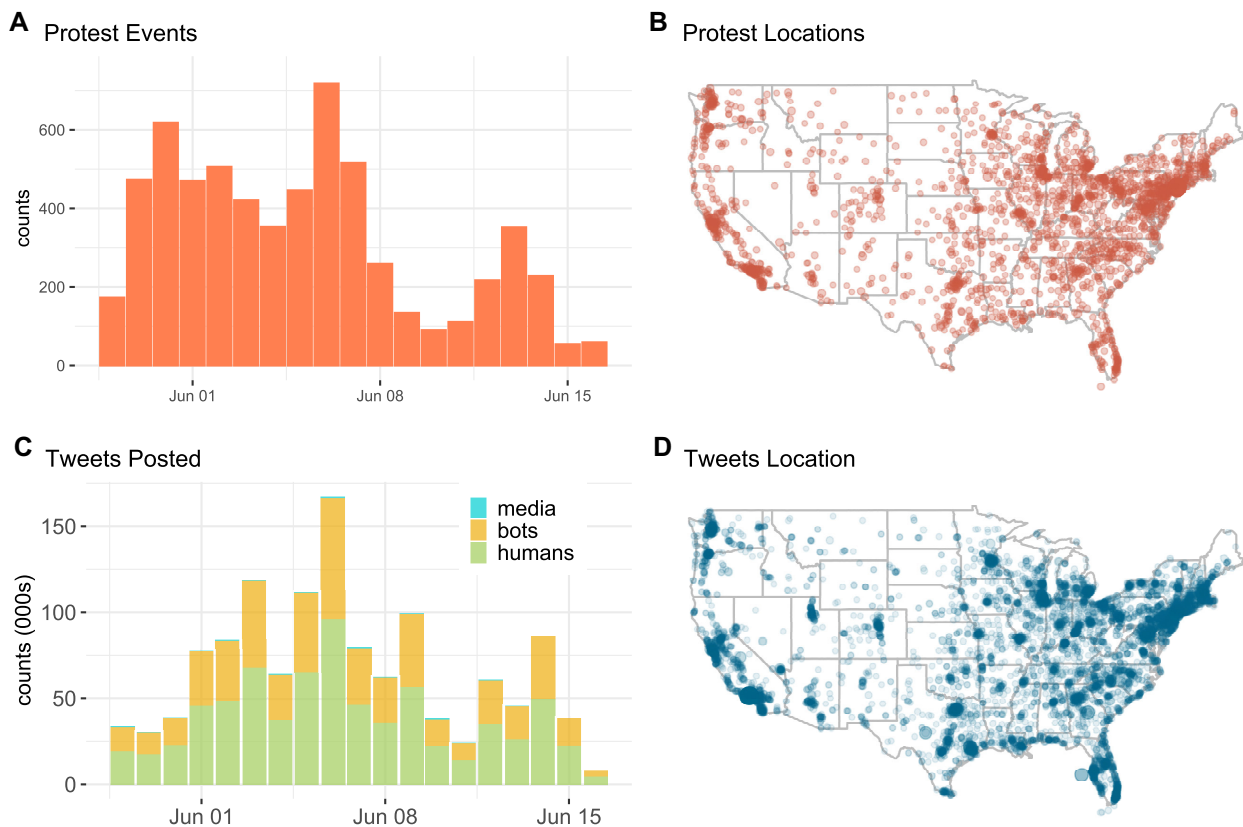


Fig. 1. Temporal and spatial distribution of protest events and tweets. The upper row shows the number (A) and location (B) of protests organized in the contiguous United States. The lower row shows the count (C) and location (D) of the tweets in our data set. Overall, Twitter activity reflects a similar temporal and spatial distribution to offline protest events, with June 6 being the day of greatest activity. Most of the tweets are generated by human accounts (~58% of all accounts), but unverified bots (~42%) generate a very large fraction of the total volume. Verified media accounts (<1%) generate a very small fraction of messages. See the “Materials and methods” and the “Supplementary Material” sections for more details on data sources and classifications.

on the reach and ideology distributions of these domains). Once we control for the audience reach of the websites and their classification in the “News/Information” category, domains that lean Republican are more frequent both in terms of total counts and unique counts of URLs shared (panel D). This pattern remains even when we exclude bot activity from the counts (see Figure S3 and Tables S3 and S4 in the “Supplementary Material” section). These results suggest that right-leaning outlets (in a partisan sense) are more visible, which is a surprising finding given the liberal bias of the Twitter user base and the specific stream of protest-related information we analyze.

Figure 3 unpacks the ideological slant of the three more frequent categories within the “News/Information” group: local news, general news, and politics (or partisan) domains (see Figures S4 and S7 in the “Supplementary Material” section for the top 30 domains within these categories). The figure reveals a shift towards the right of the distribution for party scores as we move from local news, to general news, to political domains. There are no strong differences in ideological scores for the three types of accounts (media, bots, and humans), but verified media accounts tend to share political URLs that lean more clearly towards the right (see Figures S13 to S15 in the “Supplementary Material” section). This ideological asymmetry suggests that the stream of information around BLM hashtags (and the offline protest events) was punctuated by messages systematically drawing attention to content from right-leaning domains. Even if sharing URLs does

not necessarily amount to endorsement, retweeting activity amplifies sources of information that (our results show) have a clear ideological leaning.

To further illustrate these dynamics of amplification, Figure 4 shifts the focus of attention from domains as the unit of analysis to the users posting those domains. We built the retweet (RT) network and calculated centrality measures (see Table S2 and Figure S16 in the “Supplementary Material” section for descriptive statistics). Panels A and B show the results of regression models where the dependent variable is number of RTs received (or incoming strength centrality, log-transformed). Panel A shows that verified media accounts are the most central accounts [consistent with prior research (21)], but also that posting tweets that contain URLs is associated with an increase in centrality in the RT network. We assign ideological scores to users based on the domains they share. The most central amongst this subset of users (panel B) are those posting a higher number of messages containing URLs; controlling for this, and for the number of followers, users that post URLs to Republican-leaning domains also receive a higher number of RTs (see Tables S5 and S6 in the “Supplementary Material” section for full regression results). These results suggest again that right-leaning domains have an advantage in terms of visibility and engagement: messages pointing to these sources of information receive more traction in the diffusion of content.

Panel 4C shows the RT network collapsed to its communities (the network has a high modularity score, $Q = 0.88$; see the

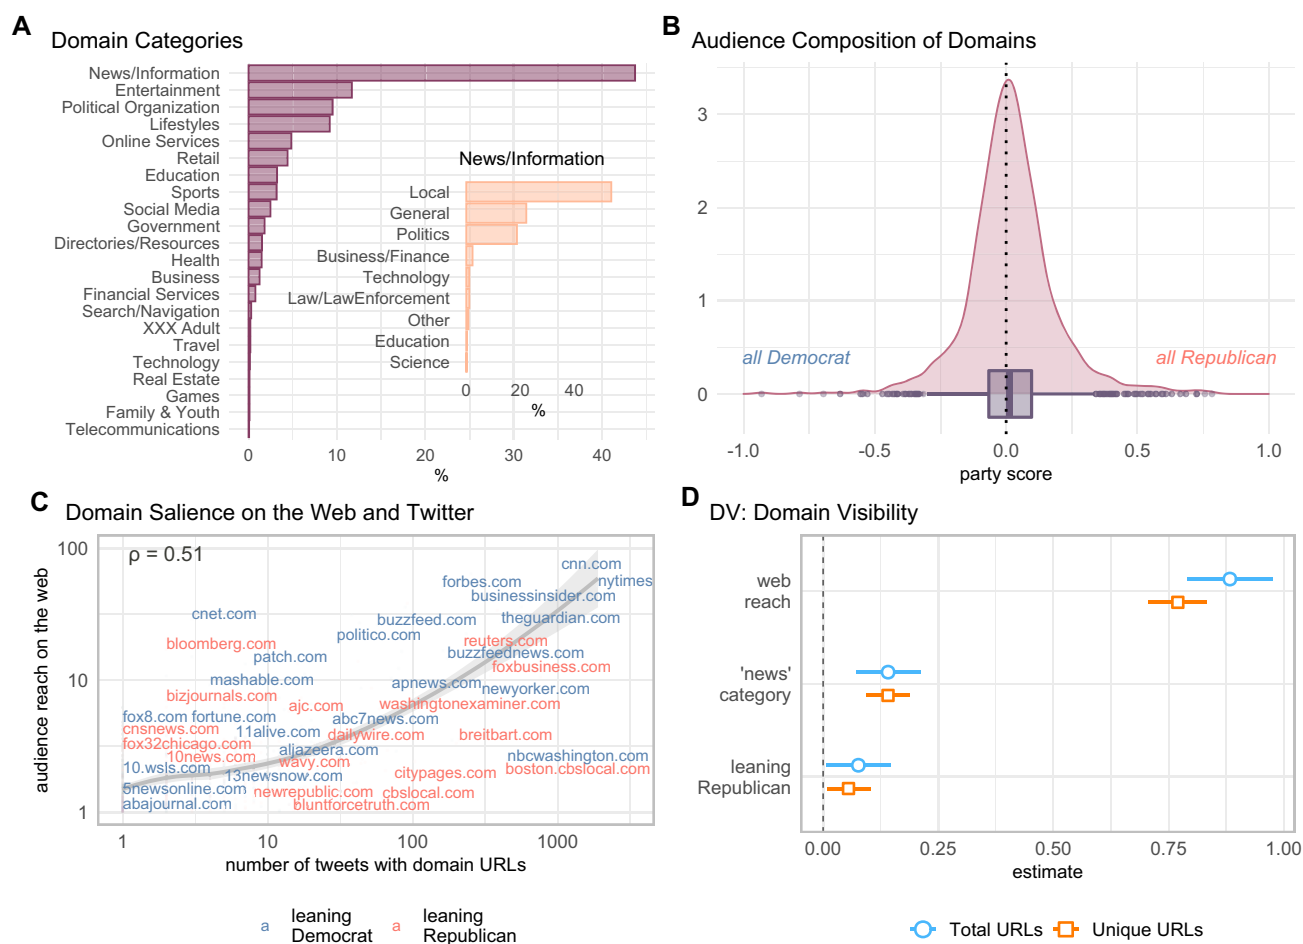


Fig. 2. Domains shared on Twitter during the protests. Most of the URLs go to news/information domains (panel A); within this category local news prevail (A inset). We assign ideological scores to these domains based on their audience composition in terms of party affiliation (panel B, see the “Materials and methods” section for details on calculations). The score equals -1 when a domain is visited exclusively by Democrats and 1 when it is visited exclusively by Republicans (0 means that Republicans and Democrats are equally likely to visit the domain). Panel C looks at the association between the audience reach of these domains on the web (i.e. the fraction of the US online population accessing the domains during this period) and the number of tweets that contain URLs to those domains (note that only a few labels are shown to improve legibility). Domains in blue have a favorability score below the mean (e.g. their audiences lean Democrat) and domains in red have a favorability score equal or above the median (e.g. they lean Republican). Panel D shows the results of linear models predicting domain visibility (measured as total URLs shared and unique URLs shared, both log-transformed, 95% CI). Web audience reach is the most important predictor of visibility on Twitter, and URLs pointing to “News/Information” domains are also more salient than non-news URLs. Controlling for these two variables, Republican-leaning URLs appear in more tweets, both in terms of total counts and unique counts (see the “Supplementary Material” section for regression tables and other specifications).

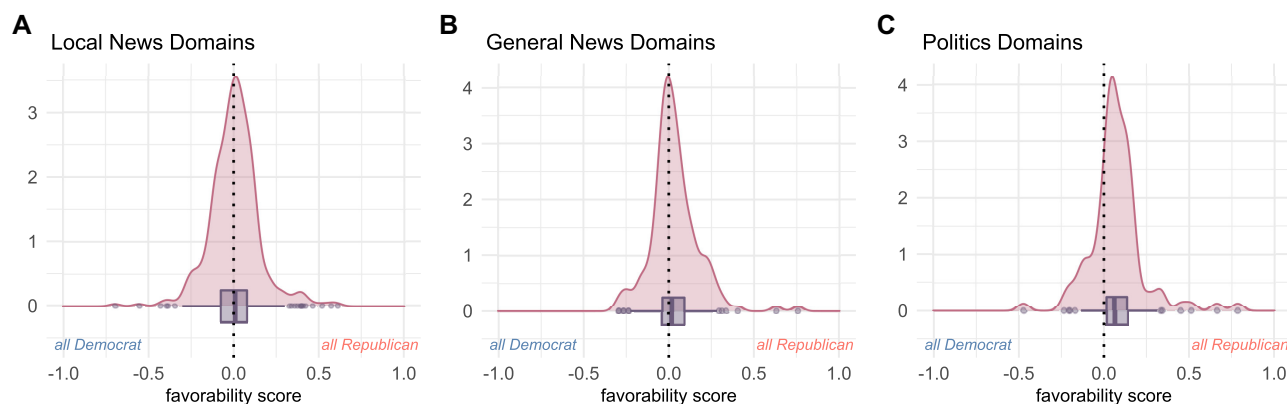


Fig. 3. Ideology distributions by domain subcategory. The panels in this figure show ideology distributions for the three most frequent categories within the “News/Information” group: (A) “local news”; (B) “general news”; and (C) “politics” (see Figure S4 in the “Supplementary Material” section for the top 30 domains within these subcategories). The “general news” and “political news” domains shared during the protests have a clear right-leaning slant. The shift to the right tail of the ideological distributions is particularly clear for political domains.

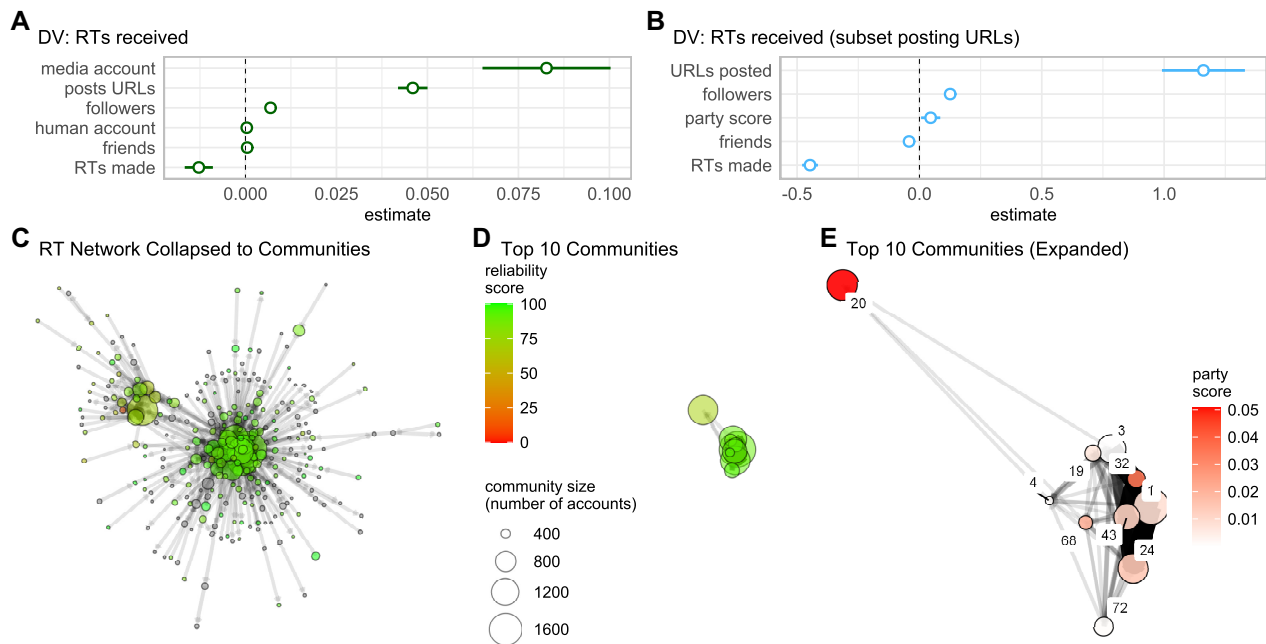


Fig. 4. Ideology and reliability in the RT network. Panels A and B show the results of linear models predicting the number of RTs received by user accounts (log-transformed, 99% CI). Accounts posting URLs have a higher centrality in the RT network. Accounts posting URLs to Republican-leaning domains receive more RTs (see the “Supplementary Material” section for regression tables and other specifications). Panel C shows the RT network collapsed to the ~280 communities identified by a random-walk algorithm (33). The network is very modular ($Q = 0.88$). Each community represents a group of accounts that RT each other more frequently; color encodes the reliability score assigned to each community based on the URLs shared. Most communities share URLs to reliable content, but the cluster on the top left has lower average reliability scores. Panel D extracts the top 10 communities in terms of size (i.e. number of user accounts classified in each), and panel E expands this subset to illuminate their average party scores (edge weight is proportional to number of RTs across communities). The two largest communities are #1 ($N \sim 1900$) and #20 ($N \sim 1400$). Users in these two communities do not RT each other. The most central user in community #1 is an American attorney who specializes in civil rights; the most popular news domain in this community is “cnn.com”. The most central user in community #20 is a conservative talk radio host; the most popular news domain is “foxbusiness.com”. The structural hole, or sparsity of RTs, separating these two communities (and the clusters around them) is suggestive of a divide in the diffusion of #BLM messages in two distinct sets of users; however, these two groups do not map onto the two halves of the ideological continuum, as indicated by the average party scores, all above the 0 line (and, therefore, with a right-leaning slant).

“Materials and methods” section for details on our community detection method). Each node represents a community, and the edges capture RTs among users in these communities. Node color encodes the average reliability score of the URLs shared within each community (see also Figures S9 and S10 in the “Supplementary Material” section for more details on reliability scores). As the plot shows, most content rates high in the reliability scale (60 or higher). There is a cluster of communities in the top left that have lower average reliability, and one small community in this cluster clearly sharing unreliable sources (in red) but the users in these communities amount to a very small fraction of all users [the secluded nature of the users interacting with less reliable sources is also consistent with prior research (6)]. Panel D extracts the top 10 communities, in terms of size, and panel E expands this subgraph to show additional detail. The RT network reveals a clear division separating two sets of users: the most popular user on one side of the divide is a conservative talk radio host; the most popular user on the other side of the divide is a civil rights attorney. The sparsity of RTs separating these users and their clusters suggests that #BLM messages diffused in two distinct sets of communities with different stances and attitudes. And yet these two groups do not map onto the two halves of the ideological continuum: the average party scores are all above the 0 line, reflecting the higher visibility that, overall, right-leaning domains have even within the confines of networks formed to advance social change.

Discussion

Social media have allowed users to create a public sphere where progressive movements (among others) can frame political causes in their own terms. These spaces, however, also host opposing voices, including those of conservative actors and mainstream media gatekeepers. Here, we document the clash of perspectives that arose on Twitter around the BLM protests in 2020. Critically, we address the questions of who produces news coverage and how audiences respond to it. Social media users consume much more content than they produce (27), which allows mainstream media and other content creators to have an influence on the platform.

Our analyses show that most of the news sources posted on social media as the massive street protests unfolded are produced by media with a right-leaning ideological slant (in partisan terms) and that this content generates more engagement in the form of retweet activity, thus increasing its reach. Right-leaning media are unlikely to portray the protest events in the movement’s terms, and their visibility in this stream of information is indicative of the ideological clash that characterize contested online spaces. Our results suggest that right-leaning domains do better (in terms of gaining visibility and engagement) than left-leaning domains. The right, in other words, has an advantage in the attention economy social media creates.

Even if some of this visibility arises from counter-attitudinal sharing (i.e. sharing content to criticize it), the boost in visibility is real, and is accelerated through social and algorithmic forms

of amplification. The social signals users create when engaging with content (to praise or criticize it) are picked up by the automated curation systems that determine which posts are seen first on users' feeds. Our analyses cannot parse out which of these mechanisms (intentional seeding versus counter-attitudinal sharing; social versus algorithmic amplification) are the most relevant in explaining the asymmetries we observe. But these asymmetries are the aggregated manifestation of social and technological mechanisms and, regardless of the motivation of users in sharing certain URLs or the specific parameters of curation algorithms, the result is still an asymmetric information environment where some coverage gains more traction and visibility.

The increased visibility of right-leaning content we identify can also result from the larger supply of content with a specific partisan slant: the users in our data are, in the end, picking content to share from a set of available sources. It is plausible that the asymmetry starts at the supply stage within the larger media environment, where news deserts are growing, and the gaps left by local newspapers is being filled by a network of websites created by conservative political groups (28) known as the right-wing media ecosystem (7). Given our focus on how Twitter was used during the protest events, our results can only illuminate how the asymmetry manifests on the platform, but not whether this is partly the downstream effect of an increasingly asymmetric news landscape with an unbalanced supply of news. Future research should try to document whether these asymmetries affect the supply of content, not just engagement with it.

Empirical results are always contingent on measurement. Future research should also try to replicate our findings on different social media platforms and using different measures of ideology. Our analyses use audience-based measures but alternative measures could be based, for instance, on voter registration records (e.g. 6) or content-related measures of media bias (e.g. 29). In the meantime, existing research already speaks to the generalizability of the patterns we identify. A recent research paper using an experimental approach, different measures of ideology, and a different timeframe shows results that are consistent with our main claim that the right has an advantage in social media (15). This work suggests there is an algorithmic amplification of the mainstream political right in six of seven countries analyzed. Even though the main focus of this work is on elected legislators from major political parties, the paper also contains results for news content in the United States. The results reported for news are less clear than the results reported for legislators, but the findings are still suggestive of amplification of right-leaning sources. Our approach to labeling the ideological slant of web domains is very different (again, we rely on the self-placement of audiences accessing those domains on the web, instead of content-based labels provided by AdFontes or AllSides); but our results showing the prevalence of right-leaning domains are consistent. This alignment suggests that the findings we report here are not contingent on our data or research design choices. If anything, our empirical context (the BLM protests) offers a more stringent test to document the advantage (i.e. the increased visibility) of right-leaning outlets.

Together, these findings show that mainstream media can shape events even in the context of activist networks. Their prominence dilutes the power of activism to frame events, and it lengthens the shadow of what we call the "advantage of the right": a disproportionate tendency of right-wing media voices to gain visibility in ideologically diverse social media spaces. The dominance of right-leaning voices on social media has been documented anecdotally in journalistic accounts (e.g. 30, 31) but there

is still a lack of systematic empirical evidence offering support to this claim. Here, we provide that type of evidence, and proof that ideological asymmetries manifest even in the context of movements with progressive goals.

In sum, our work shows that on one of the most prominent political issues of the 21st century, the perspective of right-leaning outlets dominated on Twitter. As discussed, this is surprising in some respects, especially given the well-documented population advantage of liberals over conservatives on the platform (17, 18). Yet it accords with studies of the right-wing media ecosystem, which has developed as an alternative to more centrist mainstream media and regularly attracts mass attention on controversial issues (7, 32). The prevalence of right-wing media content about BLM protests and protesters poses a challenge to the latter's attempts to set the media agenda and attract supporters. Future research should determine if the advantage of the right on display here also applies at other times and for other political issues; it should also aim to improve our measures and offer a more granular definition of what counts as misinformation or low-quality content (e.g. our domain-level measures mask heterogeneity in quality and bias at the news story level). Yet our results stand on their own as a demonstration of the prominence that right-leaning media have in the social media marketplace of ideas.

Materials and methods

Data

We collected social media data through Twitter's publicly available API by retrieving all messages that contained at least one relevant hashtag (see the "Supplementary Material" section for the list of keywords). In total, the data consists of $N \sim 52$ million tweets (all languages and locations). We restricted our main analyses to tweets written in English and produced from the United States for the time period 2020 May 28 to 2020 June 16 (20 days). This filtered dataset is formed by $N \sim 1.3$ million unique tweets (see Section 2 in the Supplementary Material, and Figures S11 and S12, for tests on how the filtering affects our measures of visibility). To benchmark online activity with the actions taking place on the streets across the country, we obtained protest event data from the Crowd Counting Consortium (20). The web-tracking data offering reach estimates and audience-based ideological scores comes from Comscore's Plan Metrix panel ($N \sim 12,000$), and it covers the same period as the Twitter activity data (2020 May to 2020 June). Following prior work (23, 24), we assign ideology scores to these domains using the audience-based measure $fav_p(d) = \frac{(R-D)}{(R+D)}$. These favorability scores equal -1, when a domain is visited exclusively by Democrats and 1, when it is visited exclusively by Republicans. The calculations exclude panelists that self-identify as "Independent". We used Comscore's classification of web domains when available (e.g. news/information, entertainment, etc.), and we manually checked to solve inconsistencies, errors, and missing labels (see the "Supplementary Material" section for more details). Finally, the reliability scores come from NewsGuard, a journalism and technology company that rates the transparency and credibility of news and information websites. We use the scores assigned to domains during the same period covered by the Twitter data (2020 May to 2020 June). These scores are provided at the domain (source) level, which may mask unreliable information published in specific news stories (see the "Supplementary Material" section for more details on the data, sources, and additional descriptive statistics).

Methods

We identify automated accounts using a bot classification technique trained and validated on publicly available datasets. Using 80% of the data for training and the remaining 20% for validation, the model achieves a classification accuracy of about 90%. When applied to an independent dataset to test out-of-domain performance, the classification accuracy decreases to 60%, which suggests the model can be generalized to new data but also that performance decreases with respect to training and validation sets (as is well known in the literature, see the “Supplementary Material” section for more details on the model and cross-validation checks). We build the weighted version of the retweet network and calculate centrality scores on the largest connected component (see Table S2 in the “Supplementary Material” section for descriptive statistics). We identify communities in the network using a random walk algorithm designed to identify dense subgraphs in sparse structures (33).

Models

The regression models have two main dependent variables (DVs): the number of URLs shared for every domain (total count and unique count); and centrality in the RT network for every user (number of RTs received). For domains, the main control variables are audience reach on the web and category (binary attribute identifying the domains classified as “news”). The ideological scores of the domains are the main variables of interest. For users, the main control variables are number of followers and friends, RTs made, and account type (media, bot, or human). The main variable of interest is whether the user posted URLs and, if so, of which ideological slant (see the “Supplementary Material” section for additional details and model specifications).

Acknowledgments

The authors would like to thank Oriol Artime and Riccardo Gallotti for useful discussions, and Anna Bertani for technical support during data collection.

Supplementary Material

Supplementary material is available at [PNAS Nexus](#) online.

Funding

This work was partly funded by NSF grants 1729412 and 2017655.

Author Contributions

S.G.B., V.A., D.F., and M.D.: research design; S.G.B., V.A., and M.D.: data collection and preprocessing; S.G.B., V.A., and M.D.: data analysis; and S.G.B., V.A., D.F., and M.D.: writing.

Data Availability

Data and code to replicate our results have been deposited in an Open Science Framework repository that can be accessed after filling in the following form: <https://forms.gle/aZ9L4ZHCnt9Pa2sEA>. Per Twitter policies, we cannot share the files with raw tweets, but we share the tweet IDs, which can be used to reconstruct the files used in the analyses (we include an index of these files). All other data are made available for replication purposes only (i.e. Comscore and NewsGuard require commercial license agreements).

References

- Buchanan L, Bui Q, Patel JK. 2020. Black lives matter may be the largest movement in U.S. history. *The New York Times*.
- Freelon D, McIlwain CD, Clark M. 2016. Beyond the hashtags: #Ferguson, #Blacklivesmatter, and the online struggle for offline justice. Center for Media & Social Impact. American University, Washington DC.
- Jackson SJ, Bailey M, Welles BF. 2020. #HashtagActivism. Networks of race and gender justice. Cambridge, MA: MIT Press.
- Freelon D, Marwick A, Kreiss D. 2020. False equivalencies: online activism from left to right. *Science*. 369:1197.
- Shugars S et al. 2021. Pandemics, protests, and publics: demographic activity and engagement on Twitter in 2020. *J Quant De-scr Digit Media*. 1:68.
- Grinberg N, Joseph K, Friedland L, Swire-Thompson B, Lazer D. 2019. Fake news on Twitter during the 2016 U.S. presidential election. *Science*. 363:374–378.
- Benkler Y, Faris R, Roberts H. 2018. Network propaganda: manipulation, disinformation, and radicalization in American politics. Oxford: Oxford University Press.
- Shin J, Thorson K. 2017. Partisan selective sharing: the biased diffusion of fact-checking messages on social media. *J Commun*. 67:233–255.
- Jost JT, van der Linden S, Panagopoulos C, Hardin CD. 2018. Ideological asymmetries in conformity, desire for shared reality, and the spread of misinformation. *Curr Opin Psychol*. 23: 77–83.
- Mourão RR, Thorson E, Chen W, Tham SM. 2018. Media repertoires and news trust during the early Trump administration. *Journal Stud*. 19:1945–1956.
- Guess A, Nagler J, Tucker J. 2019. Less than you think: prevalence and predictors of fake news dissemination on Facebook. *Sci Adv*. 5:eaau4586.
- De Keersmaecker J, Roets A. 2019. Is there an ideological asymmetry in the moral approval of spreading misinformation by politicians? *Pers Individ Dif*. 143:165–169.
- Rogers R. 2020. Deplatforming: following extreme internet celebrities to Telegram and alternative social media. *Eur J Commun*. 35:213–229.
- Schradie J. 2019. The revolution that wasn't: how digital activism favors conservatives. Cambridge, MA: Harvard University Press.
- Huszár F et al. 2022. Algorithmic amplification of politics on Twitter. *Proc Natl Acad Sci*. 119:e2025334119.
- Brady WJ, Wills JA, Burkart D, Jost JT, Van Bavel JJ. 2019. An ideological asymmetry in the diffusion of moralized content on social media among political leaders. *J Exp Psychol Gen*. 148:1802–1813.
- Freelon D. 2019. Tweeting left, right, & center: how users and attention are distributed across Twitter. Miami, FL: John S. & James L. Knight Foundation.
- Wojcik S, Hughes A. 2019. Sizing up Twitter users. Washington, DC: Pew Research Center.
- Zhang Y et al. 2019. Whose lives matter? Mass shootings and social media discourses of sympathy and policy, 2012–2014. *J Comput Mediat Commun*. 24:182–202.
- Pressman J, Chenoweth E, Crowd Counting Consortium. <https://sites.google.com/view/crowdcountingconsortium/view-download-the-data>. <https://sites.google.com/view/crowdcountingconsortium/view-download-the-data>.
- González-Bailón S, De Domenico M. 2021. Bots are less central than verified accounts during contentious political events. *Proc Natl Acad Sci*. 118:e2013443118.

22. Stella M, Ferrara E, De Domenico M. 2018. Bots increase exposure to negative and inflammatory content in online social systems. *Proc Natl Acad Sci.* 115:12435–12440.
23. Tyler M, Grimmer J, Iyengar S. 2022. Partisan enclaves and information bazaars: mapping selective exposure to online news. *J Polit.* 84:1057–1073.
24. Yang T, Majó-Vázquez S, Nielsen RK, González-Bailón S. 2020. Exposure to news grows less fragmented with an increase in mobile access. *Proc Natl Acad Sci.* 117: 28678–28683.
25. Gentzkow M, Shapiro JM. 2011. Ideological segregation online and offline. *Q J Econ.* 126:1799–1839.
26. Guess AM. 2021. (Almost) Everything in moderation: new evidence on Americans' online media diets. *Am J Pol Sci.* 65:1007–1022.
27. Antelmi A, Malandrino D, Scarano V. 2019. Characterizing the behavioral evolution of Twitter users and the truth behind the 90–9–1 rule. In *Companion Proceedings of The 2019 World Wide Web Conference*; Association for Computing Machinery, San Francisco, CA, pp 1035–1038.
28. Alba D, Nicas J. 2020. As local news dies, a pay-for-play network rises in its place. *The New York Times*.
29. Budak C, Goel S, Rao JM. 2016. Fair and balanced? Quantifying media bias through Crowdsourced Content Analysis. *Public Opin Q.* 80:250–271.
30. Thompson A. 2020. Why the right wing has a massive advantage on Facebook. *Politico*.
31. Roose K. 2020. What if Facebook is the real “silent majority”? *The New York Times*.
32. Pak C, Cotter K, DeCook J. 2020. Intermedia reliance and sustainability of emergent media: a large-scale analysis of American news outlets' external linking behaviors. *Int J Commun.* 14:3546–3568.
33. Pons P, Latapy M. 2006. Computing communities in large networks using random walks. *J Graph Algorithms Appl.* 10:191–218.