American Political Science Review

Curation Bubbles

--Manuscript Draft--

Manuscript Number:	APSR-D-23-00955R2
Full Title:	Curation Bubbles
Article Type:	Research Article
Section/Category:	American Government & Politics
Keywords:	partisanship; political information; Twitter; facebook; social media
Corresponding Author:	Jon Green, PhD Duke University Durham, NC UNITED STATES
Corresponding Author's Institution:	Duke University
Corresponding Author's Secondary Institution:	
Order of Authors:	Jon Green, PhD
	Stefan McCabe, PhD
	Sarah Shugars, PhD
	Hanyu Chwe, BA
	Luke Horgan, BA
	Shuyang Cao, BA
	David Lazer, PhD
Order of Authors Secondary Information:	
Abstract:	Information on social media is characterized by networked curation processes in which users select other users from whom to receive information, and those users in turn share information that promotes their identities and interests. We argue this allows for partisan "curation bubbles" of users who share and consume content with consistent appeal drawn from a variety of sources. Yet, research concerning the extent of filter bubbles, echo chambers, or other forms of politically segregated information consumption typically conceptualizes information's partisan valence at the source level as opposed to the story level. This can lead domain-level measures of audience partisanship to mischaracterize the partisan appeal of sources' constituent stories -- especially for sources estimated to be more moderate. Accounting for networked curation aligns theory and measurement of political information consumption on social media.
Suggested Reviewers:	Fabrizio Gilardi University of Zurich gilardi@ipz.uzh.ch
	Christopher Bail Duke University christopher.bail@duke.edu
	Leticia Bode Georgetown University lb871@georgetown.edu
	Jacob Montgomery Washington University, St. Louis jacob.montgomery@wustl.edu
	Betsy Sinclair Washington University, St. Louis bsinclai@wustl.edu

Opposed Reviewers:	
Additional Information:	
Question	Response
Please enter the total number of pages of your manuscript.	39
The following questions ask you about the ethics of the research you conducted for the manuscript you are submitting. **Unless otherwise noted, the information collected below will be shared with editors and reviewers, along with any appendix uploaded with the manuscript. Please anonymize the information as needed.** In 2020, APSA Council approved *Principles and Guidance for Human Subjects Research*, which outlines several principles that should guide all political science research. The Principles include General Principles such as respect for the autonomy of participants (a wider category than "human subjects" as defined in US federal regulations), openness about ethical issues, the researcher's non-delegable responsibility for their ethical practices, and the reasoned justification of deviations from the Principles in scholarly publications, as well as additional specific Principles. The *APSR* editorial team adheres to these Principles. When submitting research for publication in the *APSR* that draws on direct engagement with human participants in the research process (see the Principles for clarification), you will be asked to affirm that your research conforms to all Principles and to state where you discuss the relevant aspects of your research practices (whether in the text or an appendix). If claiming an exception to one of the Principles, you will be asked to provide a reasoned	No

justification for doing so. Did your research directly engage human participants in the research process?	
Did you and/or your co-authors encounter any ethical issues or perceptions of ethical issues unrelated to the study of human participants?	No
Was this research funded by any agencies, organizations, or institutions?	No
Transparency and openness considerations The following questions ask you about the transparency and openness of your research design. Unless otherwise noted, the information collected below will be shared with editors and reviewers. Please anonymize the information as needed. Did you generate or use qualitative or quantitative data for the research reported in this manuscript? Please select all that apply.	Yes, a quantitative dataset
Where do you describe how the *quantitative* data were generated or how you gained access to the data? Please check all that apply. as follow-up to "**Transparency and openness considerations**	In the main text

The following questions ask you about the transparency and openness of your research design. Unless otherwise noted, the information collected below will be shared with editors and reviewers. Please anonymize the information as needed. Did you generate or use qualitative or quantitative data for the research reported in this manuscript? Please select all that apply. "	
Do you agree to deposit in the *APSR* Dataverse at the conditional acceptance stage the quantitative datasets (and related documentation, interview guides, protocols, and scripts) and to have these materials published at the time the article or letter is published? as follow-up to "**Transparency and openness considerations** The following questions ask you about the transparency and openness of your research design. Unless otherwise noted, the information collected below will be shared with editors and reviewers. Please anonymize the information as needed. Did you generate or use qualitative or quantitative data for the research reported in this manuscript?	Yes, with limitations

Please select all that apply. "	
Where do you provide a reasoned justification of the epistemological, ethical, or practical constraints that prevent public, free access to the data? Please check all that apply. as follow-up to "Do you agree to deposit in the *APSR* Dataverse at the conditional acceptance stage the quantitative datasets (and related documentation, interview guides, protocols, and scripts) and to have these materials published at the time the article or letter is published? "	In an appendix
Please enter the manuscript word count (maximum 12,000 words for Articles, 4,000 words for Letters)	10748
Have you pre-registered your research design and/or analysis for this submission?	No
Author Comments:	
Manuscript Classifications:	Methodology; Research Methods; Public Opinion and Voting Behavior; Political Behavior

To the editors,

Thank you for conditionally accepting our manuscript, “Curation Bubbles.” In addition to our finalized materials here, we have submitted our replication materials to the APSR Dataverse. Our finalized manuscript includes the following changes based on your requests and our own review:

Editor-requested:

- Added notes to more fully describe dependent variables in Tables E.1 and H.1
- Added notes to Figure 6 and corresponding manuscript text indicating the supplementary tables in which the full model specification can be found
- Revised formatting for Figures 4 and F.1 to increase font size and overall clarity

Additional:

- Revised figure formatting to move captions outside of image files and into figure notes.
- Revised the following sentence for clarity (in the ‘Estimating Partisanship’ subsection of the ‘Data and Methods’ section):
    - Previous: "To avoid this issue, we add the absolute value of negative counts to all five political categories for the given URL, changing the formerly negative count to zero so as to keep resulting audience score constrained to the [-1, 1] range."
    - Current: "To avoid this issue, for URLs with any negative counts in any political categories, we add the largest absolute value of category-level negative counts to all categories. This coerces the minimum category-level count to zero and constrains the resulting audience score to the [-1,1] range."
- Removed the following sentence from the ‘Dataset 2: Facebook URLs’ subsection of the ‘Data and Methods’ section: “We expand this list by further collecting URLs from domains which appear in our Twitter dataset and are primarily associated with political content -- a classification process described in the next section."
    - This sentence should have been removed in the previous submission as our revised manuscript did not include this step in the analysis. We opted against adding extra Twitter-derived URLs in the Facebook analysis to allow for each platform-specific analysis to be fully independent.
- Removed the following sentence from footnote 5: “Shugars, et al (2021) shows that these estimates of partisanship, when aggregated to the county level, are strongly correlated with two-way Clinton vote share in 2016.”
    - This citation is now redundant as we show this correlation in Figure G.1
- Revised the following sentence for clarity (in the ‘Statistical and Substantive Evidence of Curation Bubbles’ subsection of the ‘Data and Methods’ section):
    - Previous: “allowing us to infer that the audience for this specific story does not reflect a data-generating process in which every *New York Times* story is sampled from an identical distribution.”

- Current: “allowing us to infer that the audience for this specific story does not reflect a data-generating process in which every *New York Times* story's audience is sampled from an identical distribution.”
- Corrected the number of hand-coders referenced in Appendix B (from eight to nine)
  - The original analysis was based on all nine / this correction to the manuscript does not alter any results.
  - Removing the one outlier referenced in Appendix B reduced the effective number from nine to eight rather than eight to seven (again, the correction reflects the original analysis, not a change to the analysis).
- Updated the regression results around the hand-coding.
  - When preparing the replication files, we discovered the regression models supporting our discussion of partisan appeal accidentally used rounded versions of some of the domain scores as input. When correcting this, the regression output and associated F-test changed slightly. These changes do not affect the interpretation of the results in any way.

Definitely let us know however we can be of assistance moving forward in the publication process!

All the best,
Jon Green
Stefan McCabe
Sarah Shugars
Hanyu Chwe
Luke Horgan
Shuyang Cao
David Lazer

Title: Curation Bubbles

Authors:
- Jon Green, Assistant Professor, Department of Political Science, Duke University, jon.green@duke.edu, ORCID: 0000-0002-6542-6745
- Stefan McCabe, Postdoctoral Associate, Institute for Data, Democracy & Politics, George Washington University, Stefan.mccabe@email.gwu.edu, ORCID: 0000-0002-7180-145X
- Sarah Shugars, Assistant Professor, Department of Communication, Rutgers University, sarah.shugars@rutgers.edu, ORCID: 0000-0002-1917-4313
- Hanyu Chwe, Ph.D. Student, Network Science Institute, Northeastern University, chwe.h@northeastern.edu, ORCID: 0000-0002-7745-8162
- Luke Horgan, Research Engineer, Network Science Institute, Northeastern University, horgan.l@northeastern.edu
- Shuyang Cao, Ph.D. Student, Department of Computer Science and Engineering, University of Michigan, caoshuy@umich.edu, ORCID: 0000-0002-3628-8291
- David Lazer, University Distinguished Professor, Network Science Institute, Northeastern University, d.lazer@northeastern.edu, ORCID: 0000-0002-7991-9110

Abstract: Information on social media is characterized by networked curation processes in which users select other users from whom to receive information, and those users in turn share information that promotes their identities and interests. We argue this allows for partisan "curation bubbles" of users who share and consume content with consistent appeal drawn from a variety of sources. Yet, research concerning the extent of filter bubbles, echo chambers, or other forms of politically segregated information consumption typically conceptualizes information's partisan valence at the source level as opposed to the story level. This can lead domain-level measures of audience partisanship to mischaracterize the partisan appeal of sources' constituent stories -- especially for sources estimated to be more moderate. Accounting for networked curation aligns theory and measurement of political information consumption on social media.

Acknowledgements: The authors would like to thank Taylor Carlson; Bruce Desmarais; Shannon McGregor; C. Daniel Myers; participants at the 2021 meetings of the American Political Science Association, Midwest Political Science Association, and Society for Political Methodology; participants at the Duke Behavior and Identities Workshop; and three anonymous reviewers for helpful feedback on earlier versions of this manuscript. The authors would also like to thank Ronald Robertson for making reference data available and John Harrington for research assistance.

Funding Statement: Stefam McCabe was supported by the John S. and James L. Knight Foundation through a grant to the Institute for Data, Democracy & Politics at the George Washington University. Shuyang Cao is supported by a Bloomberg Data Science Ph.D. Fellowship. David Lazer acknowledges support from the William & Flora Hewlett Foundation and the Volkswagen Foundation.

INTRODUCTION

The internet provides individuals with an essentially limitless amount of information and a high degree of choice in which elements of that information to consume. This has prompted concerns over whether the internet threatens societies' abilities to establish common bases of reliable information and, by extension, the sustainability of democracy in the 21st century. One side of this debate argues that the internet fuels political polarization, as users may choose to avoid conflicting viewpoints (Sunstein 2002, 2017) or may be algorithmically steered toward politically favorable content (Pariser 2011). Conversely, an extensive empirical literature on news consumption shows that incidental exposure to politically diverse sources is extremely common (Gentzkow and Shapiro 2011; Messing and Westwood 2014; Bakshy, Messing, and Adamic 2015; Guess 2021), and therefore suggests that concerns over "echo chambers" or "filter bubbles" are overblown.

We argue that networked curation processes lead information consumption on social media in particular to be more politically homogeneous than this empirical literature has thus far suggested. However, this is more a *reflection* of democracy than a *threat* to democracy – a product of individuals engaging with information, and each other, on their own terms – highlighting tradeoffs between cross-cutting exposure and active participation (Mutz 2006; Stroud 2011; Kreiss and McGregor 2023). Users on social media platforms curate the information they share with others and simultaneously receive curated streams of information tailored to their interests (Davis 2017). This involves "unbundling" discrete pieces of information from their parent sources and re-bundling them into user-level streams of content – transforming a hierarchical distribution of information (from sources to consumers) into a networked distribution of information (from users to users).

One implication of this process, which we develop in this paper, is that source-level estimates of audience partisanship may mistake heterogeneity for moderation. Users choose other users to follow based on their tendencies to share useful or otherwise appealing information (Barberá 2015), and those accounts will in turn selectively share information from a given source with their network ties (such as followers, friends, or group members) based on the extent to which that content serves *social* as well as *informational* functions (Marwick and boyd 2011; Epstein et al. 2023). That is, users share information not only to inform others, but also to perform their identities, advance their interests, and generate social returns (such as likes, retweets, or followers). Since individual stories are subject to these networked curation processes, source-level estimates do not reflect cases where individual stories from a given source are useful for different partisan audiences.

We test this implication of networked curation by comparing common measures of audience partisanship at the source and story levels. Virtually all the literature in this area aggregates partisan consumption to the level of the domain, or source, rather than examining the partisan audiences of individual stories (Robertson et al. 2018; Eady et al. 2019; Guess 2021; Peterson, Goel, and Iyengar 2021). Bakshy, Messing, and Adamic (2015) and González-Bailón et al. (2023) are notable exceptions, discussed below. Source-level aggregation implicitly assumes that every story from a given source is drawn from a consistent distribution of partisan appeal that attracts a stable ratio of Democratic to Republican users. By contrast, we find evidence of partisan *curation bubbles*, defined as sets of users who share and consume content with consistent appeal from a variety of sources.[1] When users in curation bubbles are able to identify and circulate congenial information from a variety of sources, individual stories may reach audiences atypical of the sources that produced them – introducing heterogeneity into the source's aggregate audience that doesn't necessarily reflect heterogeneity in the partisan valence of the information users consume.

We use two large-scale datasets to test for the presence and extent of partisan curation bubbles. First, we analyze sharing patterns on Twitter in 2017 and 2018 using a panel of over 1.6 million user accounts linked to a commercial voter file. We then examine sharing and exposure patterns on Facebook between 2017 and 2021 using data made available through Social Science One (King and Persily 2020). We consistently find evidence of partisan curation bubbles. The fact that we find substantively similar results in all three analyses (of sharing on Twitter and Facebook, and of exposure on Facebook) suggests a robust pattern.

1. Consistency here is with respect to a given dimension of information. A user can be in multiple curation bubbles along different dimensions, such as sports, politics, and music.

CURATION ON SOCIAL MEDIA

The online information ecosystem in the early 21st century is characterized by unbundling and abundance. An individual's news consumption near the end of the 20th century would typically be clustered in a small number of sources offering packages of information. As one could not read a story in a newspaper without buying at least a single copy of the whole newspaper, information search was largely a search for preferred packages, or sources, from which to habitually consume a variety of information. This could take the form of a subscription to a newspaper that covered news, opinion, sports, and culture – or an opinion magazine that offered a particular editorial direction. Contemporary information consumption presents a fundamentally different proposition as it is largely unbundled at the story level, such that it is practical for individuals to consume information a la carte from a wide range of sources. The task of information search is now less about identifying the most desirable sources and is instead about identifying content of interest from a functionally infinite set of options.

The internet offers individuals several different strategies to manage the task of wading through an ocean of information to identify what they want to see. Centralized aggregators such as search engines and news portals (Fischer, Jaidka, and Lelkes 2020; Robertson et al. 2023) are perhaps the most obvious and widely used, allowing individuals to input queries ranging from general ("political news") to specific ("2024 Nevada caucus results") and receive relevant information in return. Here, we focus on a different, commonly-used setting: feed-based social media, in which users follow accounts and posts from these accounts are aggregated into a flow of content. We take *curation* to be the processes through which people are matched with content that appeals to them. We consider curation to encompass both platform architecture, such as a ranking algorithm, and user choice within that architecture. Following Davis (2017), this includes both *consumptive curation*, or users' selection of accounts from which to receive information, and *productive curation*, or users' choice of what to share with others. Importantly, consumptive curation effectively delegates the search for relevant information to others. Rather than actively searching for specific information, and rather than choosing news sources to habitually consume information, users choose other users from which to habitually consume information and then scroll through whatever those users choose to post.

Online curation is analogous to prior accounts of the "two-step flow" of information from radio and print media to opinion leaders, and from opinion leaders to ordinary citizens (Lazarsfeld, Berelson, and Gaudet 1948; Katz and Lazarsfeld 1955). However, on social media this has the potential to happen with more structure and on a far larger scale. Rather than opinion leaders (perhaps haphazardly) recounting news they read earlier in the day to another individual, users on social media can immediately and directly share news with hundreds or thousands of other users at a time. Moreover, opinion leading relationships as envisioned by the Columbia school are largely formed incidentally, as a consequence of proximity within one's local community (e.g., Lazer et al. 2010; Minozzi et al. 2020). By contrast, social media allows users much more choice in who they form ties with and why, potentially including the provision of information.

These affordances of social media renewed longstanding concerns over how much choice in information consumption is too much. The ability to pick and choose individual accounts from which to receive political information carries the potential for users to select into politically homogeneous "echo chambers" (Sunstein 2002, 2017). The increased reliance on platforms that algorithmically filter, sort, and recommend content prompts parallel concerns over "filter bubbles" (Pariser 2011; Ribeiro et al. 2020) in which consuming partisan information begets exposure to more partisan information. These related concerns involve the same outcome: politically homogeneous information diets that, in theory, frustrate democratic societies' abilities to make collective decisions using common bases of reliable information.

Empirical research regarding the extent these potentially undesirable outcomes manifest is mixed (Prior 2013; Barberá 2020; Dahlgren 2021). This is in part because individuals' tendencies to engage in selective exposure within their information environments are not as straightforward as early theories regarding the concept predict -- in line with early skepticism (Freedman and Sears 1965, 1967). While some individuals do select pro-attitudinal sources (Stroud 2011), this does not necessarily mean they are actively avoiding counter-attitudinal information. Indeed, individuals are especially likely to seek (and subsequently share) pro-attitudinal information when they are exposed to counter-attitudinal information (Garrett 2009; Weeks et al. 2017). This dynamic is less obviously concerning, and can take place in the context of healthy deliberative exchange. In addition, people often rely on heuristics other than partisanship when deciding which information to consume, such as topical relevance (Kobayashi and Ikeda 2009; Mummolo 2016) or social endorsements (Messing and Westwood 2014). As a result, partisan segregation in aggregate news consumption online is typically found to be relatively low (Gentzkow and Shapiro 2010; Flaxman, Goel, and Rao 2016; Guess 2021).

This finding initially extended to social media. Early research on Facebook showed that since friendship ties formed for a variety of reasons – many of which were incidental to politics – Facebook users were frequently exposed to politically distant sources (Bakshy et al. 2012; Bakshy, Messing, and Adamic 2015). However, when extending this analysis on Facebook to Pages and Groups, which form for more specific reasons, González-Bailón et al. (2023) find stronger evidence of political segregation in information consumption – consistent with other work finding evidence of political homophily on social media (Conover et al. 2011; Conover et al. 2012).[2] Put simply, one may be friends with a politically distant acquaintance or relative on Facebook in spite of their politics but follow a Page because of its politics, which will have consequences for the diversity of information to which one is exposed. Moreover, pro-attitudinal information spreads more quickly, is consumed more frequently, and is received more approvingly within political communities on social media sites than counter-attitudinal information (Halberstam and Knight 2016; Garz, Sörensen, and Stone 2020). This imbalance is likely attributable to the political information users choose to share on social media sites. Sharing information with one's followers is inherently more public than consuming it oneself, and can be used to signal (or, from the opposite perspective, infer) political identities and commitments (Marwick and boyd 2011; Settle 2018). In the rare instances in which users share political information from opposing partisans, it is often accompanied by negative comments that indicate disagreement (Cinelli et al. 2021; Wojcieszak et al. 2022).

The Facebook Page, in the above example, is acting as a curator – an account that shares or reshares content. An account that posts a link to a story in the *New York Times* is identifying that content as worthy of attention. Consumers are accounts on social media that are exposed to content. Users have the ability to act as both a consumer and curator, though in practice the vast majority of productive curation is done by a small number of accounts (Grinberg et al. 2019; Wojcik and Hughes 2019; Hughes et al. 2021) who are more politically active offline (and exhibit more partisan extremity) than

2. See also work showing that politically engaged users occasionally sever ties on social media for political reasons (Bode 2016; Neely 2021).

users who do not post about politics themselves (Hughes 2019). These curators, in turn, take an active role in purposively identifying individual stories to share, and how to frame those stories for their followers (Park and Kaye 2018; Billard 2021). Importantly, curators do not necessarily share information solely for information's sake – the act of sharing specific information (as opposed to other information one could potentially share) is a means by which users can signal aspects of their identity that are important to them (e.g., Van Bavel et al. 2021; Osmundsen et al. 2021). Political information sharing on social media will therefore likely feature partisan curators – i.e. users who selectively share information that promotes their political in-group or detracts from political out-groups. These users are, in a sense, performing "hidden labor" for their preferred party, attempting to shape the character of online discourse by selectively sharing politically favorable information.[3]

The underlying logic and architecture of information sharing on social media is therefore likely to produce *curation bubbles*, or sets of users who share and consume content with consistent appeal from a variety of sources. We view curation bubbles as a general property of social media not limited to partisanship. For example, Taylor Swift fans will curate information related to Taylor Swift from a variety of sources. This will include atypical sources, such as ESPN, when ESPN publishes stories about Taylor Swift. Here, we are interested in partisan curation bubbles, or users who tend to share (and, through homophilous tie formation, see) politically consistent information from a variety of sources.

If and when politically neutral or distant sources publish individual stories that are useful for promoting partisan identities and interests, partisan users will share them with their followers (who are in turn likely to be co-partisans themselves), introducing heterogeneity into those sources' audience for their constituent stories. By extension, partisan curation bubbles are formed via co-partisan users sharing information favorable to their party. The breadth and variety of information available on the internet allows partisan users to easily find politically favorable information (Peterson and Iyengar 2021). This information can originate from a variety of sources, and indeed partisans tend to overestimate the extent to which mainstream outlets perceived as ideologically distinct offer substantively different coverage (Peterson and Kagalwala 2021). Furthermore, politically favorable information may be most useful for promoting one's party *precisely when* it is attributable to a source perceived to be politically

3. We thank an anonymous reviewer for raising this point.

FIGURE 1. Stylized examples.

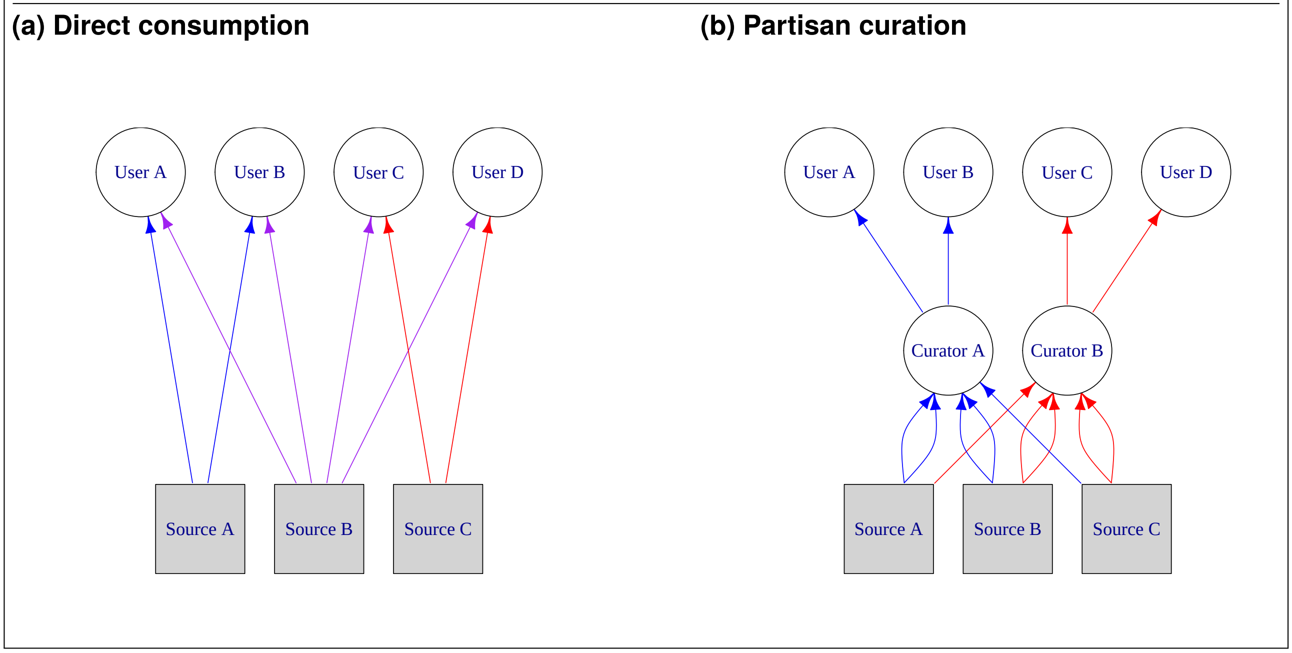

neutral or distant, as this can increase its credibility (Baum and Groeling 2009). Temporal variation in whether the news is broadly favorable to the political left or right can also introduce selective engagement with the news itself (Kim and Kim 2021), which would lead to variation in which partisan curation bubbles are circulating more or less raw information at any given time.

Figure 1 provides an illustrative characterization of the partisan curation process, in comparison to a process solely driven by users consuming information directly from sources. In both cases, there are three sources that, based on users' overall consumption behavior, appear to be left-leaning, neutral, and right-leaning, respectively. In Figure 1a, this is reflected by two left-leaning users consuming the left-leaning source, two right-leaning users consuming the right-leaning source, and all four users consuming the neutral source. In Figure 1b, there are two curators mediating these users' consumption. Curator A only shares blue stories with the two left-leaning users who follow them and Curator B only shares red stories with the two right-leaning users who follow them, irrespective of the sources that produced those stories. The pattern of consumption is quite integrated at the producer level; yet is completely segregated at the story level. Source B, in particular, appears neutral overall not by producing stories that all users consume, but by producing stories that are curated by either left-leaning *or* right-leaning users.

Partisan curation bubbles carry implications for how researchers understand the political valence of the information being shared on social media. Empirical researchers frequently quantify source-level partisan slant using estimates of the partisanship of news outlets' overall audiences (e.g., Robertson et al. 2023; Eady et al. 2019; Guess 2021; Garimella et al. 2021). These estimates typically represent a normalized ratio of how often URLs from the given domain were shared by Democrats compared to Republicans. For instance, a domain shared exclusively by Democrats would receive a score of -1, a domain shared exclusively by Republicans would receive a score of 1, and a domain shared by equal numbers of Democrats and Republicans would receive a score of 0. The major exceptions that construct scores at the URL as well as domain level are Bakshy, Messing, and Adamic (2015) and González-Bailón et al. (2023). The former evaluates exposure to cross cutting partisan content on Facebook; the latter examines segregation in news consumption on Facebook. Both sets of results are consistent with the possibility of partisan curation bubbles, but neither directly studies their presence.

The core assumption of this approach – common to many approaches for quantifying political valence on social media (e.g., Barberá 2015) – is that behavior reflects revealed preferences. This is agnostic to the substance of the content in question, in contrast with methods that infer the slant of a given story or source based on its text (Ho and Quinn 2008; Gentzkow and Shapiro 2010).[4] This can make reputations self-fulfilling. If, for example, Republican users avoid the *New York Times* because it is regarded as left-leaning, then the *New York Times* will garner left-leaning audience regardless of what the newspaper publishes (Peterson and Kagalwala 2021) – which will carry through to its location on the [-1,1] scale. Similarly, a score of zero doesn't mean that the domain is "neutral" in any sense deeper than that it was shared by Democrats and Republicans at equal rates. In other words, the average partisanship of sources' audiences is a relative measure of partisanship, not an absolute one (Robertson et al. 2018; Guess 2021).

As the stylized example in Figure 1 suggests, partisan curation bubbles have the potential to distort estimates of partisan appeal at the source level because they can introduce substantial heterogeneity at the story level. Importantly, this distortion is unlikely to be uniform – there will be more story-level heterogeneity in audience partisanship within sources that carry more moderate overall estimates. The reason for this is mechanical as well as theoretical: there is only one way for individual stories to

4. Though even in these settings, it is not obvious that such estimates can be taken as ground truth measures of "bias," per se. Gentzkow and Shapiro (2010) argue that newspapers may adapt the slant of the content they produce to match the political preferences of their audience.

aggregate to an extreme domain score (circulation among consistently partisan audiences), but there are two ways to produce a moderate score. A moderate domain can produce stories that are consistently circulated by both Democrats and Republicans at relatively even rates, or it can produce stories that are disproportionately circulated by *either* Democrats *or* Republicans. The moderate domain-level average will only reflect the individual stories that produced it in the former case. However, as partisan curators selectively share stories that are socially useful for them, the latter will frequently occur (we expand on this point in Appendix A).

Finally, we note that our theoretical framework is agnostic as to the potential role of social media platform's recommendation algorithms. While algorithmic curation is undoubtedly important for determining which information users see, algorithms themselves do not inevitably lead to the consumption of politically homogeneous information. Algorithms can optimize on a variety of criteria, some politically salient and others not (Fischer, Jaidka, and Lelkes 2020; Bandy and Diakopoulos 2021), and different platforms may make different design choices that could (either intentionally or incidentally) encourage or discourage exposure to counter-attitudinal information (Garrett and Resnick 2011). For example, algorithm-based recommendations from centralized news aggregators such as MSN or Google may be more likely to direct users toward large, mainstream sources than they are to direct users toward niche, ideological sources (Guess 2021). The best evidence in this area on social media in particular comes from a platform-wide experiment on Twitter, which found that its algorithmic timeline led users to be exposed to more political content than users who remained on chronological timelines (Huszár et al. 2022). However, that same study found inconsistent effects with respect to whether the amplification of political content was disproportionately in favor of left- or right-leaning content. While we are unable to isolate the potential contributions of platforms' algorithms to our empirical findings, we view it as exceedingly unlikely for individual stories to circulate among politically atypical audiences in the absence of users intentionally curating those stories for their social ties.

Hypotheses

Our theoretical framework carries a set of empirical implications that we test in this paper.

First, because networked curation occurs at the story level, we expect to observe audience heterogeneity within sources.

> **Hypothesis 1a:** Productive curation. The partisan composition of *sharing* behavior will exhibit story-level heterogeneity within sources.
>
> **Hypothesis 1b:** Consumptive curation. The partisan composition of *viewing* behavior will exhibit story-level heterogeneity within sources.

We further expect that this heterogeneity corresponds with substantive differences in the latent partisan appeal of the information being circulated and that it is not, in expectation, due to idiosyncrasies such as "hate-sharing" or noise.

> **Hypothesis 2:** Partisan audience scores estimated at the story level will reflect the substantive partisan appeal of those stories.

Finally, we expect systematic variation in the extent to which partisan curation bubbles pose a challenge to interpreting and using source-level estimates of audience partisanship. Specifically, moderate domain-level estimates of audience partisanship are more likely to mis-characterize the partisan audience for any given story. Relatively more extreme source-level estimates, by contrast, will more frequently reflect the audiences for each individual story.

> **Hypothesis 3:** Moderate source-level scores will more frequently mis-characterize the partisan appeal of their constituent stories.

We illustrate these general points in Table 1, which displays the stories (URLs) from the *Wall Street Journal* with the 10 most Republican audience scores and the 10 most Democratic audience scores. Domain-level scores often identify the *Wall Street Journal* as "neutral" (Gentzkow and Shapiro 2010; Bakshy, Messing, and Adamic 2015), and it carries a relatively centrist domain-level audience score of -0.34 in our Twitter data. However, we see that this score, if applied to every story produced by the *Wall Street Journal*, fails to adequately describe its cross-cutting content. The stories disproportionately circulated by Republicans are largely conservative opinion pieces from the editorial page. In contrast, none of the stories disproportionately circulated by Democrats are opinion pieces; they are mainstream reporting with content that is good news for Democrats and bad for Republicans. In other words, the *Wall Street Journal* does not have a consistent moderate audience; it produces individual stories with differential partisan appeal that reach different partisan audiences. This is the dynamic we will explore and test through our curation bubbles framework.

TABLE 1. Top headlines from the *Wall Street Journal*, minimum 250 shares.

Democratic-Leaning Stories	Score	Republican-Leaning Stories	Score
'Boss, I Miss You So Much': The Awkward Exile of Michael Cohen	-0.71	Who Paid for the 'Trump Dossier'?	0.54
GOP Activist Who Sought Clinton Emails Cited Trump Campaign Officials	-0.70	When Justice Is Partial	0.52
Hacking Probe Spurs Back-and-Forth on Eve of Georgia Governor's Race	-0.70	The Memo and the Mueller Probe	0.51
Melania Trump's Military Flights Before Her Move to Washington Cost More Than $675,000	-0.70	U.S. Is World's Most Competitive Economy for First Time in a Decade	0.51
The Varied—and Global—Threats Confronting Democracy	-0.69	The FBI's Trump 'Insurance'	0.51
Trump Organization Tied to Deal to Keep Stormy Daniels Quiet	-0.69	U.S. Consumer Confidence Hits 14-Year High	0.50
U.S. Eyes Michael Flynn's Links to Russia	-0.68	Brennan and the 2016 Spy Scandal	0.49
ACLU Will No Longer Defend Hate Groups Protesting With Firearms	-0.68	The Scandal That Matters	0.49
How the House GOP Health Plan Compares to the ACA	-0.68	Lifting the Steele Curtain	0.46
U.S. Military's Space in Trump Tower Costs $130,000 a Month	-0.68	Mueller's Fruit of the Poisonous Tree	0.46

Note: On the left, the 10 headlines with the most left-leaning audience scores; on the right, the 10 headlines with the most right-leaning audience scores. Audience scores drawn from Twitter sample.

DATA AND METHODS

We use data from Twitter and Facebook to examine both domain and story-level partisan curation bubbles in the U.S. This cross-platform comparison allows for validation of our key result, but comes with challenges. People use different social media platforms for different reasons (Evans et al. 2017), meaning that we would expect variation in engagement and exposure across the two platforms. However, both platforms are of scientific interest, with Twitter being particularly influential among journalists and Facebook being the social media platform the general public most frequently uses for news consumption overall (McGregor and Molyneux 2018; Molyneux and McGregor 2021; Jurkowitz and Gottfried 2022). While these platforms' user bases differ in size and composition, leading us to expect variation in the precise stories and domains which circulate on each site, users on both platforms engage in similar styles of networked curation.

Perhaps the bigger challenge to cross-platform analysis is methodological. Data on Twitter and Facebook are collected and structured differently, requiring slightly different approaches for estimating partisanship, as discussed in detail below. These differences, however, also come with opportunities. Our Twitter data contains information for individual users with fine-grained measures of their likely partisan affiliation. Our Facebook data does not contain such individual-level data, but it does include clicks, reactions, and views in addition to shares. Each platform therefore allows us to test phenomena that the other does not.

In total, our Twitter data consists of 405,531 unique URLs shared on that platform between January 1, 2017 and December 31, 2018. These URLs originated from a total of 8,378 domains. Our Facebook data consists of 218,395 unique URLs from 908 domains shared between January 1, 2017 and February 28, 2021. We analyze less content on Facebook because we focus on domains and URLs less affected by privacy-preserving noise. Each of these datasets and their partisanship measures are described in detail below.

Dataset 1: Twitter users with matched voter data

For this study, we collected tweets from a panel of Twitter users matched to U.S. voting records. Taking a user-focused approach to data collection allows us to identify a consistent population over time and to bring in user-level demographic information, including measures of party affiliation. A pilot version of this dataset was described in Grinberg et al. (2019) and more descriptives are provided in Hughes et al. (2021) and Shugars et al. (2021). For purposes of this analysis, two details from those papers are relevant: first, although slightly more white and female than the population of American Twitter users, our panel is otherwise generally representative of Twitter users (Hughes et al. 2021); second, our vendor for voter data (TargetSmart) provides a modeled estimate of party identification that correlates well with aggregate electoral results, allowing us to avoid the vagaries of interpreting party registration between states (Shugars et al. 2021).[5]

Panel users were identified in 2017. Starting with 290 million profiles retrieved from Twitter's 10% Decahose sample, we searched for profiles in which the Twitter names (display name or handle) and locations matched entries in the voter file that were unique at the city level (or state level, if the Twitter profile does not list city). We successfully matched 1.6 million accounts corresponding to registered U.S. voters. Because some users may go inactive, this represents an upper bound on our population size. Once identified, we retroactively collected panelists' past tweets dating back to 2010. Since 2017, we have regularly collected all new, publicly posted panelists' tweets.

In this analysis, we analyze URLs shared or retweeted by our panelists between January 1, 2017 and December 31, 2018. We include retweets in our analysis because this is the primary means of sharing content authored by others. While it is certainly the case that, on the margins, some sharing (and by extension in our Facebook data, clicking, reacting, and viewing) behavior is done with disapproval of the underlying content, past work has has found that on Twitter this sort of "hate sharing" is concentrated in quote tweets (Wojcieszak et al. 2022), and so we exclude URLs shared through this mechanism from our analysis.[6] We restrict our focus to URLs shared a minimum of 10 times, giving us an initial set of 1,404,035 unique URLs originating from 82,293 domains. After excluding URLs not likely to contain political content (see discussion below) and domains with fewer than 1,000 total shares, we focus our analysis on a subset of 369,675 politically relevant URLs originating from 718 domains.

Dataset 2: Facebook URLs

To analyze partisan circulation of content on Facebook, we use the FORT URLs Shares dataset which is available to researchers through a collaboration with Social Science One (King and Persily 2020;

5. Our copy of the voter file, and associated estimates of partisanship, are from 2017. These estimates of partisanship have been used elsewhere in the literature (Broockman and Kalla 2023).

6. Quote tweets are also relatively uncommon; on average for every one quote tweet a user sends, they send roughly one user-authored tweet, two replies, and three retweets (Shugars et al. 2021).

Messing et al. 2021). The dataset counts the number of people who viewed, clicked, reacted to, or shared any given URL. A URL must have received at least 100 public shares to be included in the dataset, but data from private individuals is included in the released counts. The dataset fulfills differential privacy guarantees by adding Gaussian noise to all counts (Messing et al. 2021). Because the Gaussian function is constant across URLs, the signal-to-noise ratio is highest for high-engagement URLs.

Specifically, we begin by collecting all URLs from domains shared on Facebook at least 1 million times in the U.S. between January 1, 2017 and February 28, 2021. This initial sweep yields 5,545,381 URLs from 1,132 domains. We then impose three filtering processes. First, to avoid irregularities introduced by the added statistical noise (Buntain et al. 2023), we only consider URLs that have been shared at least 1,000 times, viewed 10,000 times, clicked 5,000 times, and reacted to 5,000 times. Second, we remove all URLs not classified as political. Finally, we remove domains with fewer than 10 unique URLs. These three filtering processes trim the Facebook dataset to 214,995 unique political URLs from 780 domains.

Classifying political content

Since we are interested in the curation of political content, we filter both datasets to URLs we classify as political. To do this, for every URL we retrieve the title and "blurb" – the short text which is displayed for a URL on social media. For Facebook, this information is directly available through the FORT URLs Shares dataset. For Twitter, we scrape this information.[7] For both platforms, we classify each URL as related to politics or not politics using a convolutional neural network and word vectors initialized with the GloVe pre-trained embedding (Pennington, Socher, and Manning 2014). The final classifier is trained on New York Times, Wikipedia, and Facebook data and achieves a precision of 99% and a recall of 92%.[8]

7. We scraped each URL in the dataset and extracted the text present in the `og:description` HTML tag. (The Open Graph protocol defines a number of HTML tags a publisher can add to their website to enable platforms like Facebook and Twitter to easily discover information about a web page's contents). This is the same process by which Facebook extracted the blurbs present in their dataset (Messing et al. 2021).

8. We used a 98% training-test split; precision and recall are reported for the held-out test set (48,000 articles).

Estimating Partisanship

We estimate a URL's partisanship as the average partisanship of interactions with that content. This means slightly different things on different platforms, though we have conducted our analysis so as to make the platforms as closely comparable as possible. We elaborate on relevant measurement considerations here.

Our Twitter panel is matched to a commercial voter file that includes a reliable modeled estimate of each user's likelihood of identifying as a Democrat on a 0-100 scale.[9] Where possible, we use this numeric representation of likely Democratic identification rather than trichotomizing the measure to partisan categories in order to preserve as much information about model uncertainty as possible. For the purposes of capturing the partisanship of sharing behavior in a manner comparable to traditional approaches, we implement a linear transformation of this score by subtracting it from 50 and then dividing by 50 to put this modeled estimate on $[-1, 1]$ scale running from most Democratic to least Democratic. This is preferable to relying solely on party registration, which is not collected in all states.

Our Facebook data does not include a direct measure of user partisanship, but does aggregate interaction counts into five categories of political ideology: -2, -1, 0, +1, and +2, from very liberal to very conservative. These labels are included with the FORT URL Shares dataset and are estimated based on the political pages a user follows, similar to Barberá et al.'s tweetscores (Barberá et al. 2015; Messing et al. 2021). While typically interpreted as ideology, this measure anchors Democratic politicians on one side and Republican politicians on the other. For consistency with our Twitter data we therefore refer to this measure as partisanship.

Added Gaussian noise in the Social Science One data can make this calculation difficult: popular but hyperpartisan URLs in our Facebook dataset are shown as having negative share counts among out-partisans. Constructing audience scores with these negative counts could lead URLs to fall outside of the $[-1, 1]$ range when normalizing. To avoid this issue, for URLs with any negative counts in any political categories, we add the largest absolute value of category-level negative counts to all categories. This coerces the minimum category-level count to zero and constrains the resulting audience score

9. In Appendix G show that this modeled estimate of partisanship is strongly correlated with county-level election returns.

to the $[-1, 1]$ range. This allows us to calculate political scores in a straightforward manner without substantively altering our methodological approach or eventual results.

For both platforms, we then calculate a URL's partisan audience score as the average partisanship of its interactions. For Twitter, this means assigning each sharing event the modeled partisanship of the user who shared the URL, and then averaging those scores. For Facebook, we construct this average based on the counts of shares in the five partisanship categories, weighting interactions by the partisan values of -2 to +2, before normalizing to the common $[-1, 1]$ scale. Consistent with past work finding that Twitter has a more left-leaning user base than Facebook (Wojcik and Hughes 2019), the total average audience score on Twitter is -0.39, while on Facebook the analogous score (based on sharing behavior) is -0.12.

We use sharing behavior on both Twitter and Facebook as our primary estimate of content partisanship on these platforms. In order to preserve information regarding individuals who share the same content multiple times, we use the total number of *shares* rather than the unique number of *sharers* which has been used in past work (Bakshy, Messing, and Adamic 2015; Robertson et al. 2018).[10]

In addition to this cross-platform comparison, we calculate three more estimates of content partisanship on Facebook using measures of views, clicks, and reactions. While only available for our Facebook dataset, these measures establish the robustness of our main findings and give additional insight into the multi-faceted curation process of social media. Views refers to the number of times a piece of content appeared within a user's feed; clicks captures the consumption choices of what users click on once they are exposed; and finally, reactions indicate users' public responses of "like", "love", "haha", "wow", "sorry", or "anger."

Statistical and Substantive Evidence of Curation Bubbles

We use a number of strategies to test the empirical implications of our curation bubbles framework. Our tests of H1a and H1b begin with comparing URL-and domain-level audience scores on each platform.

10. The two measures used by Bakshy et al and Robertson et al differ slightly as well, where measures used by Robertson represent sources as slightly further towards the political right. The former's domain-level scores are averages of URL-level ratios of Democratic to Republican sharers; the latter is a ratio of Democratic to Republican sharers at the domain level. We present comparisons to these earlier measures in Appendix D.

If story-level partisan composition follows source-level composition, these distributions will be similar. However, if story-level partisan composition is heterogeneous as we hypothesize, these distributions will differ. Specifically, the distributions of URL-level scores should exhibit more extremity than the distributions of domain-level scores.

On Facebook we are able to test both productive (H1a) and consumption (H1b) curation. While the sharing data of Twitter only allows for testing H1a, the user-level data from this platform allows for further tests by user partisanship. For example, we can compare the average domain-level audience score of URLs Democrats share with the average URL-level audience score of URLs Democrats share. This allows us to further test H1a by examining possible partisan drivers of our results.

Differences between URL and domain-level distributions are important because we expect that scores estimated at the story level will reflect the substantive partisan appeal of those stories (H2). To test this, we had a team of hand coders evaluate the partisan appeal of a sample of 1,000 news stories drawn from our Twitter data. These hand coders, a collection of graduate students and postdocs, were asked to evaluate the appeal of selected stories to Democrats (-1) or to Republicans (1), or both equally (0). The full instructions are included in Appendix B. We sampled URLs for coding with probability proportional to the absolute deviation between the URL-based audience score and domain-based audience score; that is, we oversampled stories in partisan curation bubbles. Each coder evaluated 500 stories (Krippendorf's $\alpha = 0.673$), and we averaged the results to produce a hand-coded score of partisan appeal on the same scale as the URL-based audience score.

Finally, we expect systematic variation in how well domain estimates capture the substantive partisan appeal of their constituent stories (H3). Specifically, we expect that moderate domain scores are more likely to mis-characterize the partisan appeal of stories (URLs) from that domain. We test this by estimating the extent to which we can statistically distinguish URL-level audience scores from their parent domains' audience scores. These tests also serve as an important robustness check for H1a and H1b, demonstrating that differences in distributions are not merely due to partisan variation in the volume of URLs associated with different types of sources.

More formally, we test the extent to which individual URLs have partisan audiences that are statistically and substantively distinguishable from the aggregate audience of their associated domain.

This assumes that, in the absence of partisan curation bubbles, URL-level estimates of audience composition would be sampled from a normal distribution centered at the domain-level audience score[11]. We can then test, for each constituent story, whether its observed audience score is statistically distinguishable from a story generated under this null hypothesis.

One concern with this approach is that, given the large share volume for many stories, we could reject the null hypothesis despite trivial substantive differences. We therefore conduct a further test for *substantive* significance by widening our uncertainty intervals by 0.1 in each direction, analogous to the use of two one-sided tests in equivalence testing (Rainey 2014). Testing for differences of at least 0.1 on the [-1,1] scale (i.e. 5% of the full possible range) corresponds to perceptible differences in audience composition on the left and the right. For example, the difference in Twitter-based domain scores between `breitbart.com` and `nationalreview.com` is 0.11 and the difference between `thenation.com` and `theatlantic.com` is 0.09.

By way of example, consider individual URLs associated with the *New York Times* on Twitter (top left panel of Figure 4). The mean partisanship (and therefore the domain score) for `nytimes.com` is -0.59 and the standard deviation is 0.59. Under the null hypothesis of no partisan curation bubbles, we would expect most URLs to have audience scores that fall within an interval characterized by the standard error of the mean, with the width depending on our chosen confidence level. For the *New York Times*, the interval will be centered on -0.59, with the width depending on our chosen confidence level, the domain-level standard deviation of shares (0.59) and on the square root of the number of URL shares. One such URL, an editorial criticizing a court decision on voter-registration policies,[12] was shared 75 times and has a URL score of -0.72. This point estimate is to the left of the domain-wide average, but is not statistically distinguishable from it at the 99% confidence level. as this score falls within the interval $-0.59 \pm 2.57\frac{0.59}{\sqrt{75}} = [-0.76, -0.41]$. On the other hand, an editorial denouncing the shooting of Representative Steve Scalise[13] (175 shares, URL score 0.25) falls well outside of the interval, $-0.59 \pm 2.57\frac{0.59}{\sqrt{175}} = [-0.70, -0.47]$, allowing us to infer that the audience for this specific

11. There are two notable approximations here: first, we are assuming that the number of shares for the URL is relatively large; second, we are assuming that the URL score is not bounded to the interval $(-1, 1)$. The former is only a concern for infrequently shared URLs on Twitter and would not be a problem on Facebook, due to our stricter inclusion criteria. In Appendix F we show, using a non-parametric bootstrap estimator, that our results are not sensitive to this approximation.

12. https://www.nytimes.com/2018/10/19/opinion/sunday/north-dakota-addresses-voting-id.html

13. https://www.nytimes.com/2017/06/14/opinion/steve-scalise-congress-shot-alexandria-virginia.html

story does not reflect a data-generating process in which every *New York Times* story's audience is sampled from an identical distribution.

One concern with this approach is that, given the large share volume for many stories, we could reject the null hypothesis of no difference despite trivial substantive differences. To address this, we conduct a further test for *substantive* significance by widening the confidence intervals by 0.1 in each direction, analogous to the use of two one-sided tests in equivalence testing (Rainey 2014). Testing for differences of at least 0.1 on the [-1,1] scale (i.e. 5% of the full possible range) corresponds to perceptible differences in audience composition on the left and the right. For example, the difference in Twitter-based domain scores between `breitbart.com` and `nationalreview.com` is 0.11 and the difference between `thenation.com` and `theatlantic.com` is 0.09.

For each domain, we then calculate the proportion of constituent URLs that have audience scores statistically and substantively distinct from the domain-level average. For example, for the New York Times, only 5% of stories are substantially different from the domain score of -0.59. Per H3, we expect this proportion to be higher for more moderate domains (e.g., those with a score close to 0), and lower for more extreme domains (e.g., those closer to -1 or 1).

RESULTS

We evaluate: whether story-level heterogeneity in partisan appeal is reflected in sharing behavior (H1a) and viewing behavior (H1b); whether partisan appeal is recognizable at the story level (H2); and whether this heterogeneity leads moderate source-level estimates to mischaracterize the partisan appeal of their constituent stories more frequently than extreme source-level estimates (H3).

Testing for productive and consumptive curation

We first show evidence of productive curation (H1a) by plotting the distributions of domain- and URL-level audience scores based on sharing behavior on Twitter and Facebook in Figure 2. Here, we focus on domains within the top quartile by number of political URLs, though in Appendix I we show that these results are consistent across a variety of thresholds. In both cases, the distributions of URL and domain score differ from one another. The differences are statistically significant under a Kolomogorov-Smirnov test (Facebook: $D = 0.12$; $p = 0.0101$; Twitter: $D = 0.23$; $p = 0.0003$). When visually comparing the distributions between platforms, it may seem counter-intuitive that the Twitter distributions have a larger test statistic. The KS test is comparing the largest absolute deviation between the empirical cumulative distribution functions rather than the whole shape of the distribution. If you use a distance measure that compares the full distribution—Wasserstein distance—the larger substantive difference between the distributions on Facebook is apparent ($D = 0.18$ for Facebook; $D = 0.08$ for Twitter). Here, what is relevant is that in both cases the Kolmogorov-Smirnov test rejects the null hypothesis of no differences between the domain- and URL-level distributions. As we show below, the relative visual similarity between the domain- and URL-level distributions on Twitter masks substantial heterogeneity across domains.

In our Twitter sample, there is considerably more political news sharing on the political left than on the right (the platform-wide audience score is -0.39), which is at least partially attributable to the partisan composition of Twitter's user base and our sample of Twitter users (Hughes et al. 2021), while Facebook is more balanced (its corresponding score is -0.14). However, on both platforms – and especially on Facebook – the distributions of audience scores at the URL level exhibit more extremity than they do at the domain level – providing preliminary support for H1a.

To better understand productive curation in the context of a platform's user base, we further examine variation in domain- and URL-level audience scores by users' modeled partisanship (Figure 3). For ease of visualization, here we trichotomize our user-level measure of modeled partisanship into likely Democrats ($p(Dem) > .65$), unlikely Democrats ($p(Dem) < .35$), and users for whom likely partisanship is uncertain ($.35 \leq p(Dem) \leq .65$). The figure shows that domain-level scores mischaracterize the sharing profiles of a significant number of users in all three groups[14], but especially those who are unlikely to be Democrats. Under a domain measure, 28% of unlikely Democrats' political information sharing is, on average, to the right of 0; and 1.5% is to the right of 0.5. Using a URL measure, these respective percentages are 39.1% and 7.3%. This indicates that there are a substantial number of unlikely Democrats who tend to share political information from sources with

14. KS test statistics for likely Democrats $D = 0.30$; $p < 10^{-16}$, uncertain: $D = 0.11$; $p < 10^{-16}$, unlikely Democrats: $D = 0.15$; $p < 10^{-16}$).

FIGURE 2. Distribution of URL and domain-level scores based on sharing behavior for Twitter and Facebook.

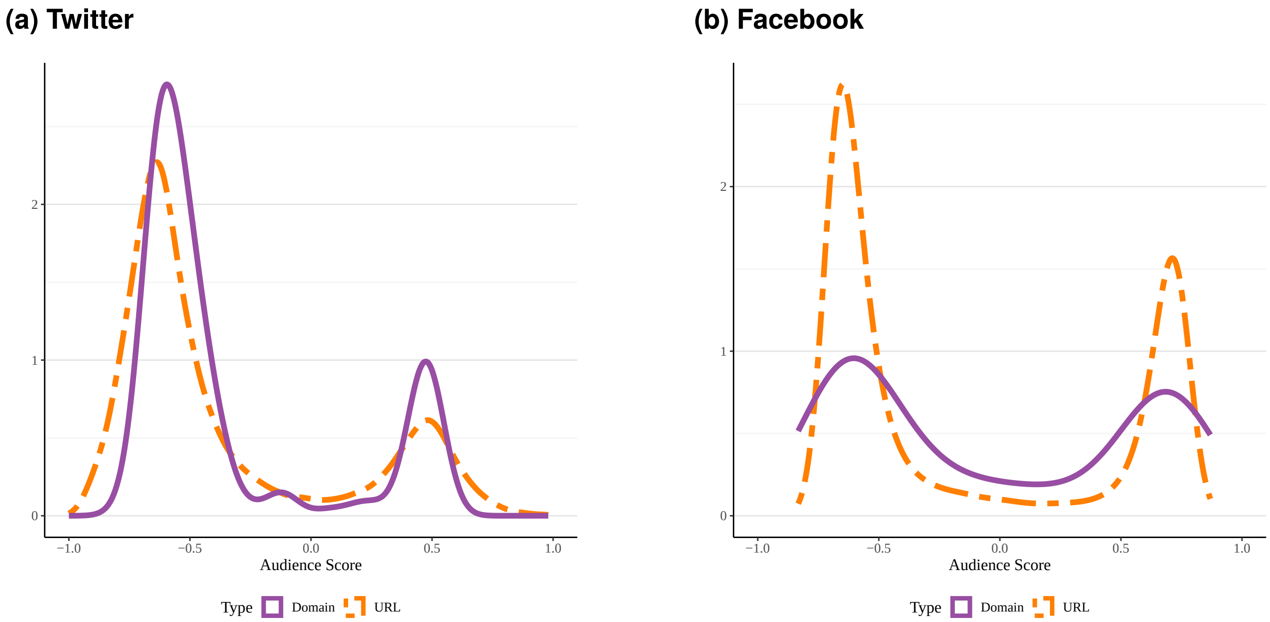

Note: Data limited to domains in the top quartile of unique political URLs.

generally left-leaning audiences when specific stories from those sources are disproportionately shared by right-leaning users, further supporting H1a.

We illustrate this dynamic in Figure 4, which plots individual URLs' audience scores and share volume on both Twitter and Facebook for the *New York Times,* the *Wall Street Journal*, Mediaite, Fox News, RT, and *Reason*. These cases are illustrative in that they vary in size and overall audience partisanship. The domain-level audience score for each is shown with a vertical dashed line; individual stories with audience scores statistically and substantively distinguishable from the domain score at the 99% level are shown with greater opacity, and in blue (more Democratic than expected under the null) or red (less Democratic than expected under the null), relative to those that are within this uncertainty interval. Stories that could be statistically significantly distinguished from the source under a null hypothesis of no difference, but whose difference did not meet our threshold of substantive significance, are shown in yellow. This figure shows that, even for sources with more extreme overall audience scores and a relatively lower share of stories in partisan curation bubbles, atypical partisan audiences do often find specific information from those sources to circulate at high volume, consistent with H1a.

This figure also previews the dynamic we will systematically test in H3. For sources with neutral

FIGURE 3. Twitter URL- and domain-level partisanship of shared political URLs, by modeled partisanship of users.

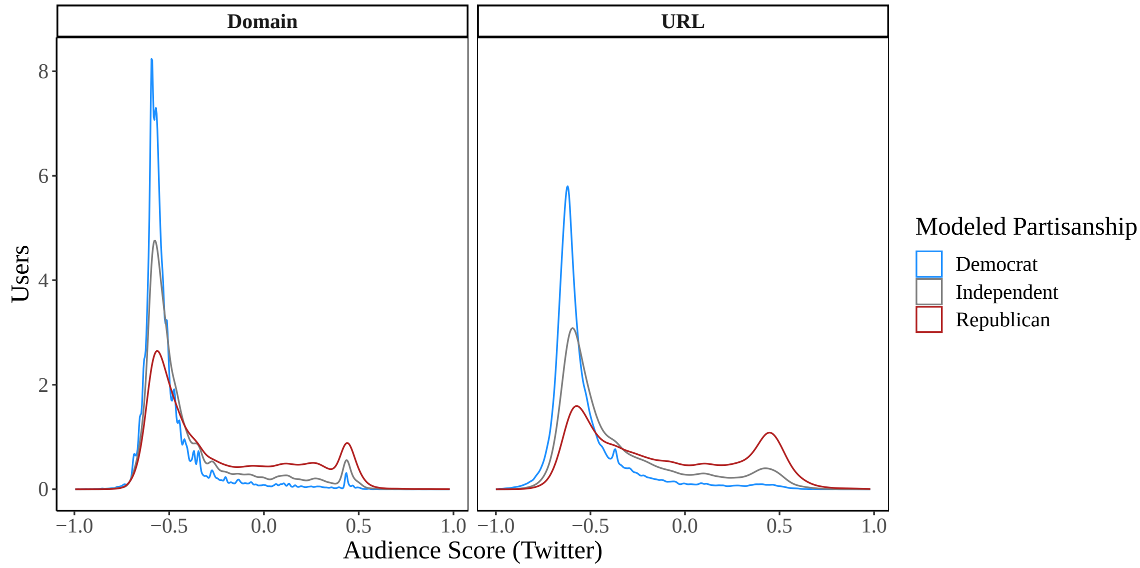

Note: Only users who shared at least five politics-related URLs are included.

domain scores, story-level fluctuations between different partisan audiences are in some cases the norm – *especially* for the stories that circulate at high volume. The *Wall Street Journal*'s well-known divide between its "hard news" and editorial content, discussed above, is further apparent here – as is Mediaite's idiosyncratic audience.[15] While we further test H3 across all sources, these results provide preliminary evidence that domain-level scores near zero cannot be straightforwardly interpreted as indicating reliably neutral content. Moreover, it is precisely the domains with the most neutral audience scores that exhibit the most within-domain partisan heterogeneity. These domains are not garnering neutral audience scores solely by producing content that is consistently shared by Democrats and Republicans at equal rates; they often produce content that is alternately shared by either Democrats or Republicans disproportionately.

We first test whether the heterogeneity we observe in productive curation (H1a) extends to consumptive curation (H1b). Using Facebook's FORT URLs dataset, we recalculate domain and URL-level scores using the consumptive measures of clicks, reactions, and views, along with the productive measure of shares for comparison. The results are shown in Figure 5. While we find that

15. Mediaite is a self-consciously bi-partisan outlet that primarily posts video clips of pundits and politicians from across the political spectrum commenting on current events.

FIGURE 4. URL scores by share volume for selected domains on Twitter and Facebook.

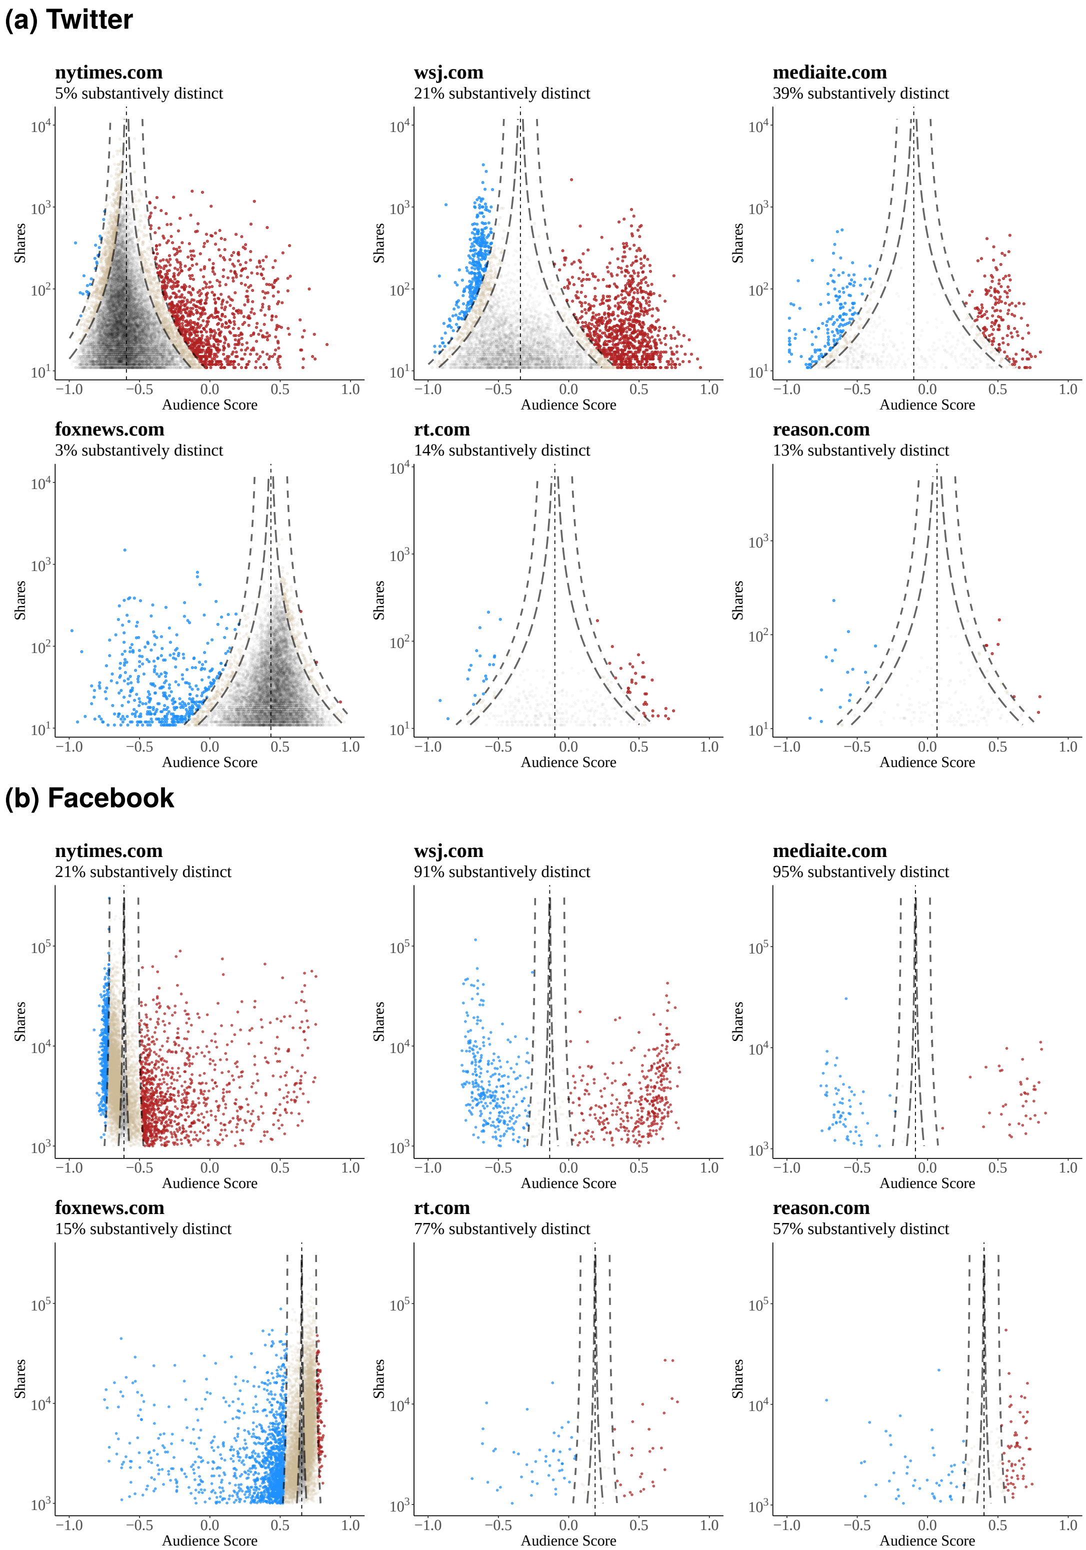

Note: Points represent URLs, colors represent relationship to domain-level score. Grey points are not statistically distinguishable from the domain-level average, yellow points are statistically but not substantively (> 0.1) distinguishable, blue points are substantively more left-leaning, and red points are substantively more right-leaning. Total proportion of URLs substantively distinct from domain-level average shown in facet subtitles.

FIGURE 5. Distributions of URL-level audience scores by engagement type (Facebook).

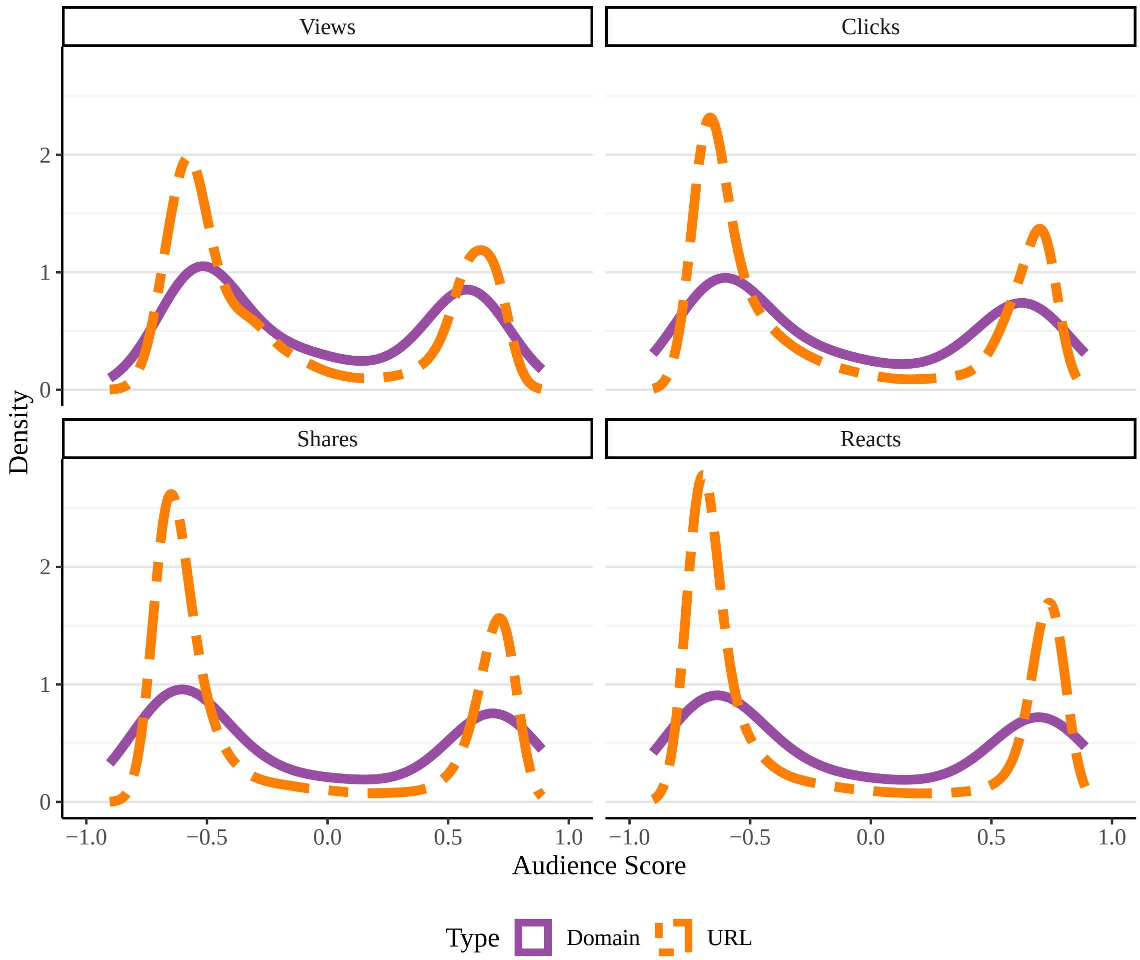

Note: Data limited to domains in the top quartile of unique political URLs.

public-facing behaviors (shares and reactions) exhibit more extremity than private behaviors (clicks and especially views), we again find that the domain-based approach consistently understates the partisan extremity of engagement,regardless of how engagement is measured. This supports H1b, showing that story-level heterogeneity in sharing behavior (productive curation) is carried through to story-level heterogeneity in viewing behavior (consumptive curation).

Substantive differences in partisan appeal

We argue that this heterogeneity is attributable to substantive differences in story-level partisan appeal (H2). To test this, we use the subset of 1,000 URLs for which we have hand coded partisan appeal.

FIGURE 6. Hand-coded partisan appeal against Twitter-based audience scores.

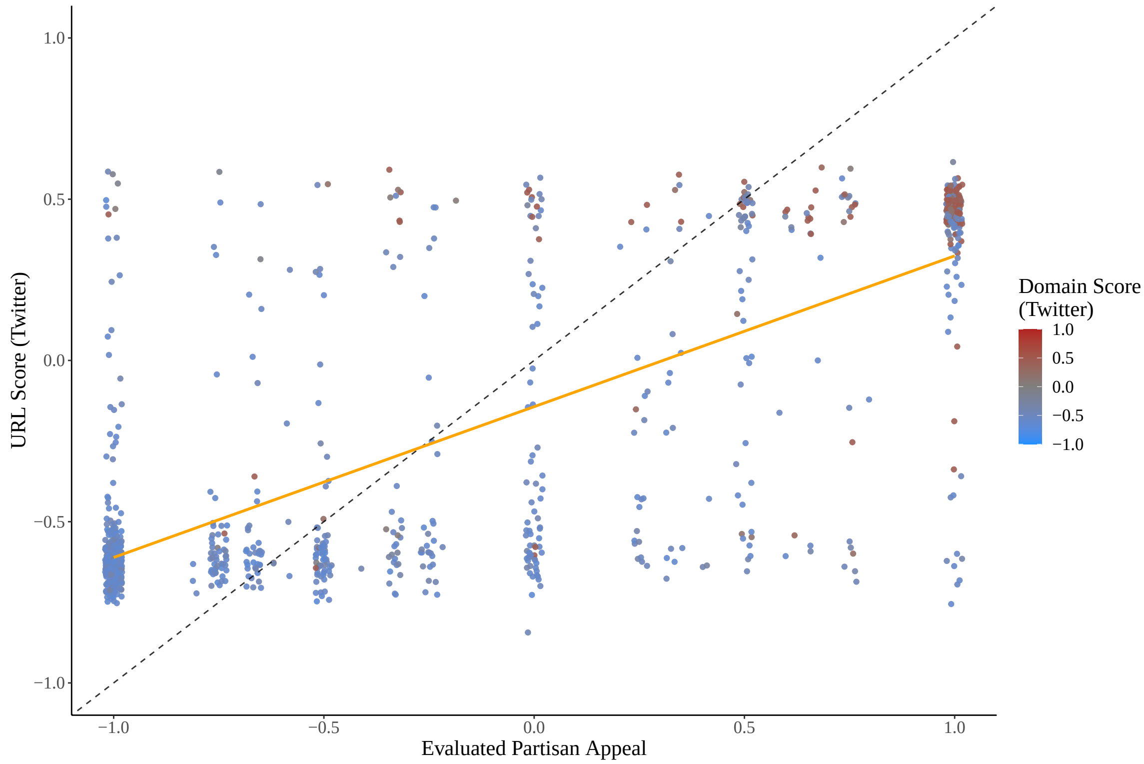

Note: Dashed black line is the identity line ($y = x$). Solid orange line is line of best fit. Articles with large URL score-domain score discrepancies were oversampled for hand coding. See Table E.1 for related regression results.

Specifically, we test whether variation in humans' assessments of whether Democrats or Republicans would view a story more favorably is better explained by URL-level or domain-level audience scores. We find there is a strong correlation ($r = 0.75$) between the URL audience scores and the evaluated partisan appeal, and that this relationship is much stronger at the URL level than at the domain level ($r = 0.55$). This is further illustrated in Figure 6, which plots human-evaluated partisan appeal against URL-level audience scores, with domain-level audience scores reflected in the color gradient. When partisan appeal is regressed against URL and domain scores together (Table E.1), an F-test supports the inclusion of domain scores as improving model fit ($F = 41.01$), but the substantive improvement is minimal, shifting the adjusted R^2 from 0.554 to 0.571. Put simply, we find support for H2: variation in URL-based audience scores reflect variation in the substantive partisan appeal of a story, and this is not a product of source cues.

Differential implications for domain-level estimates

Finally, we test systematic differences in the extent to which curation bubbles distort estimates of partisan audiences across the [-1,1] scale at the source level (H3). Figure 7 shows, for our Twitter data, the proportion of each domain's stories that significantly differ from the domain-level audience score under a null hypothesis of differences less than 0.1 (i.e. the proportion of stories that are *substantively* distinct from the domain-level average) by the domain-level audience score, with illustrative sources labeled. We supplement this analysis in Figure 8, showing the same dynamics across both productive and consumptive measures using our Facebook data.

While Facebook's larger user base and higher volume narrows confidence intervals such that larger proportions of stories are statistically distinguishable from their domain-level averages in general, both platforms show a clear trend. While every domain at least occasionally produces stories that circulate among atypical partisan audiences, this is significantly more common as domains' audience scores approach zero (see OLS regressions in Appendix H). That is, the more neutral the domain-level audience score, the more frequently that domain's constituent stories have partisan audience scores that are substantively different than the domain's audience as a whole. This supports H3, showing that moderate domain scores tend to mischaracterize the partisan appeal of their constituent stories at the highest rates.

Results for specific domains, shown on the plot, reflect qualitative understandings of those domains' audiences as well. For example, hyper-partisan outlets such as Daily Kos and the Huffington Post on the left, and Breitbart and Fox News on the right, have extreme domain scores and fewer stories that circulate among atypical partisan audiences. By contrast, the *Wall Street Journal*'s domain-level audience score of -0.34 is frequently ill-suited to describe individual stories the newspaper publishes, as previously indicated in Table 1. Furthermore, the domains with audience scores that least frequently capture the partisan appeal of their constituent stories are those with audience scores near zero, such as Mediaite or the *New York Post*. It is also worth noting that domains with audience scores near zero and relatively less within-domain heterogeneity are often outlets that are ideological in ways that do not neatly reflect partisanship in the United States, such as the Russian state-sponsored RT.

FIGURE 7. Proportion of URLs substantively distinct from domain by domain-level audience score for Twitter.

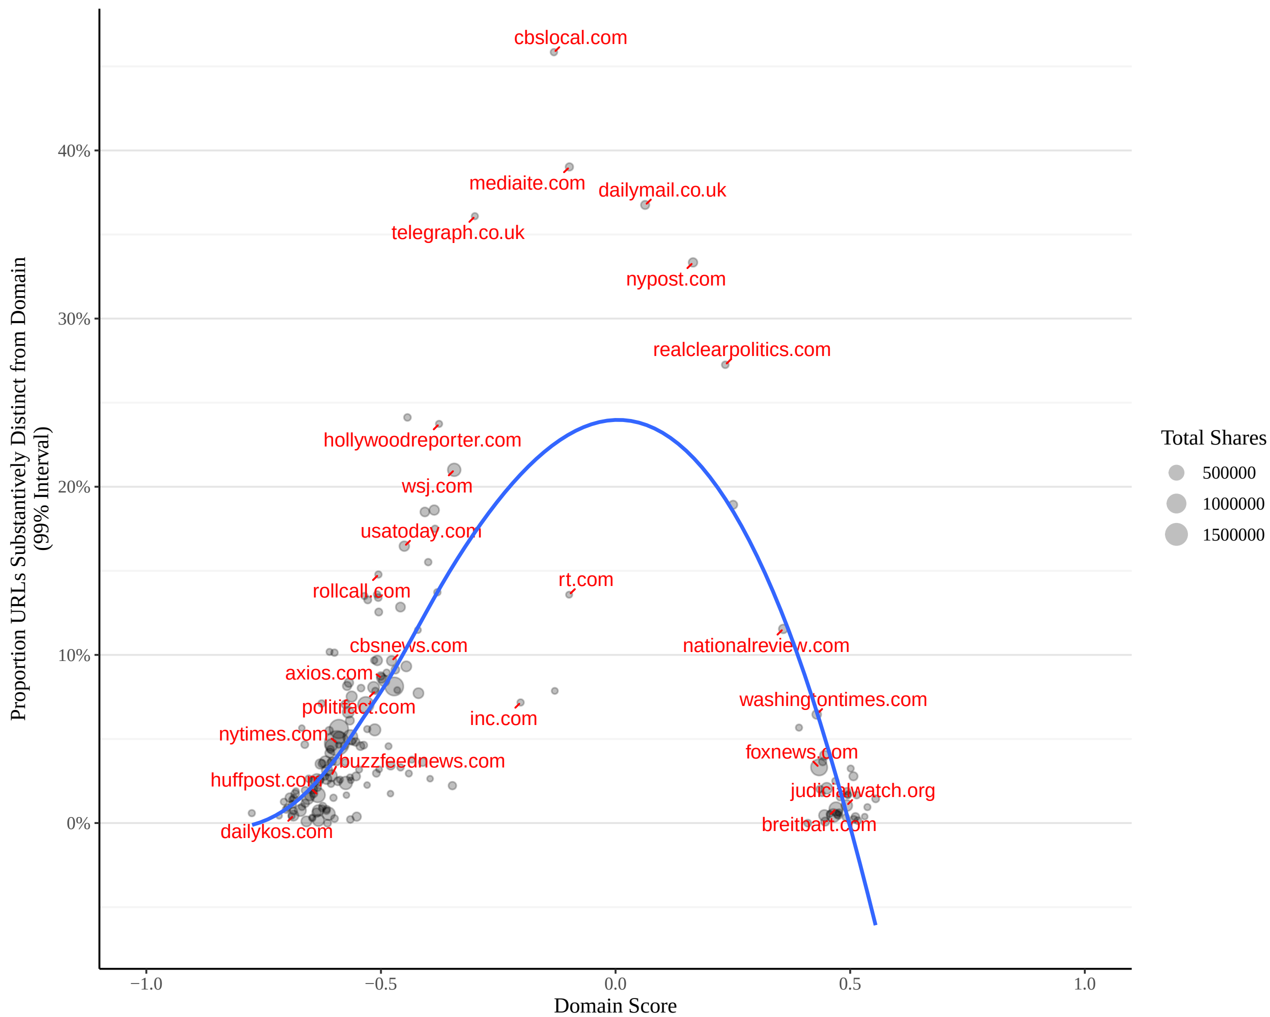

Note: Solid line indicates loess curve of best fit. Data limited to domains in the top quartile of unique political URLs.

DISCUSSION

We find evidence of partisan curation bubbles across our analyses, as users share and consume information with consistent partisan appeal from a variety of sources. These partisan curation bubbles frequently lead to story-level heterogeneity within sources for both productive (H1a) and consumptive (H1b) curation. This audience heterogeneity likely reflects a heterogeneity of within-source partisan appeal, as audience scores estimated at the story level do reflect the partisan appeal of content (H2). Furthermore, we find systematic variation in this heterogeneity—with more moderate estimates

FIGURE 8. Proportion of URLs substantively distinct from domain for different Facebook engagement types.

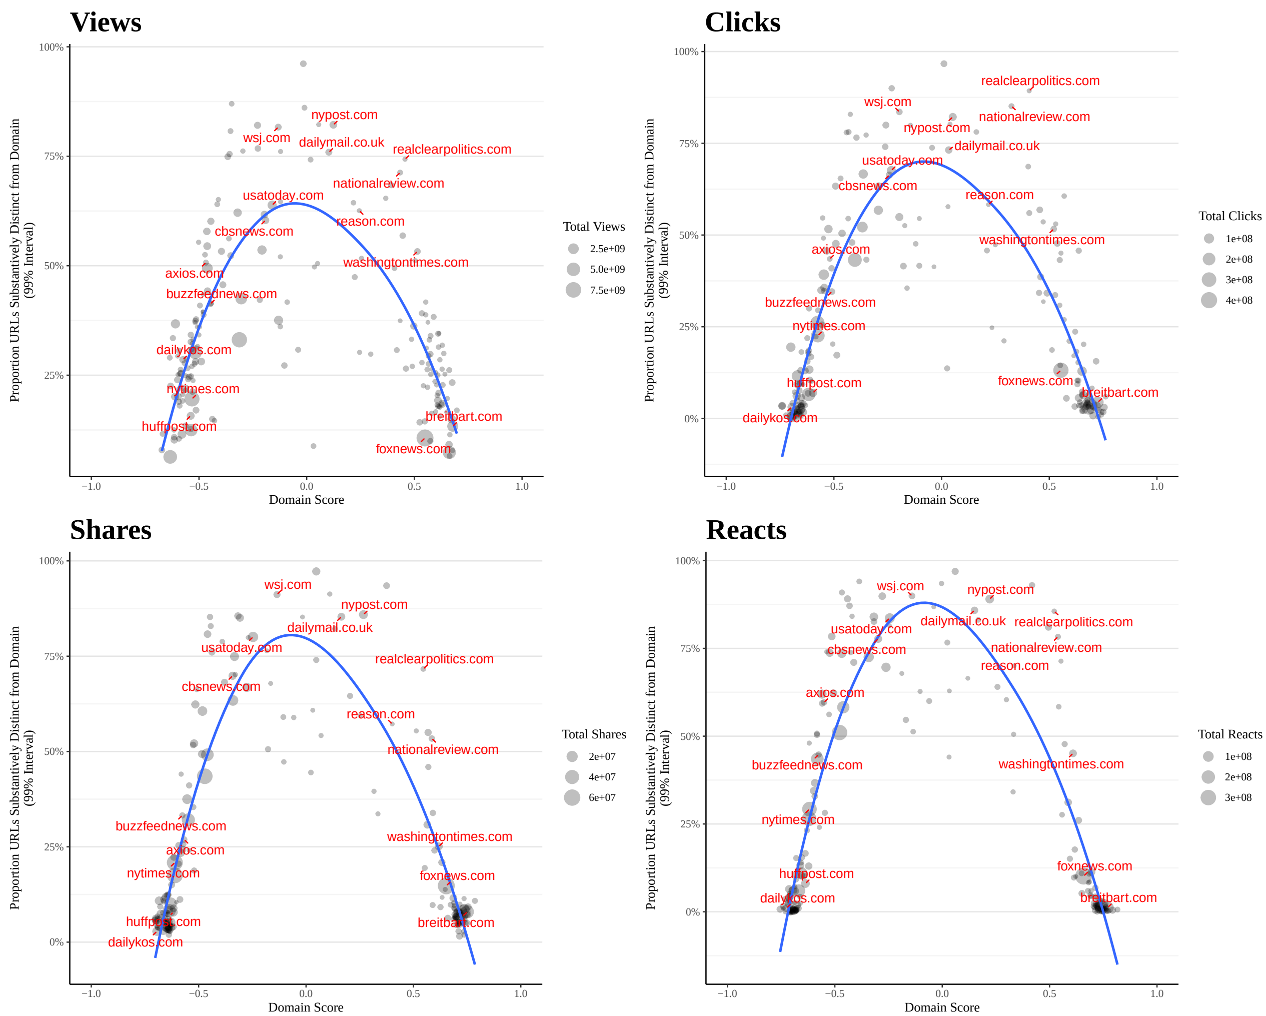

Note: Solid line indicates loess curve of best fit.

of domain-level audience partisanship more frequently mis-characterizing the partisan valence of individual stories (H3). This suggests that relatively moderate domain-level scores are often the result of different stories circulating among different partisan audiences, rather than every story reaching a consistently balanced audience.

It is likely that elements of these curation processes predate the internet and social media. For example, opinion leaders who subscribed to a given newspaper may have tended to read and talk about particular stories that matched their prior political preferences. However, observing this process prior to the internet would have required an impossible scale of instrumentation. There is a sense in which the internet has merely made this process visible. For example, recent work examining user behavior within Google Search indicates that even though users' search results do not systematically vary by partisanship, their choice of which search results to click on does (Robertson et al. 2023). The theoretical mechanisms underlying these findings could be extended in further work – such as by experimentally manipulating the pairing of politically (in)congruent stories with politically (in)congruent sources to directly test the extent to which users are willing to share politically favorable information from ideologically distant sources.

However, we argue that the internet has dramatically changed the structure of supply and demand to make networked processes of curation far more important. The fundamental logic of the internet is competition for attention at a granular story level, and the ratio of available information to human attention has increased by many orders of magnitude. This is in contrast to pre-internet competition at the outlet level, with consumers choosing stations to watch and newspapers to subscribe to. The networked curation processes of social media allows individuals to delegate the task of navigating a functionally infinite amount of information to other users who regularly share information that appeals to their identities and interests. One of the natural results of this process are partisan curation bubbles.

Purely with respect to measurement, our findings suggest that source-level measures of audience partisanship should be used with caution as they risk overestimating the partisan diversity of information consumption. All but the most extreme sources have a meaningful amount of partisan heterogeneity at the story level, and for some sources this is the rule rather than the exception. There will be times when source-level aggregation is theoretically warranted or practically necessary. In settings outside of social media where information consumption is not characterized by networked curation, source- and story-level estimates may generate similar results. Our point is to emphasize that source-level aggregation is a measurement choice that must be considered on a case-by-case basis.

More broadly, these findings shed new light on the macro structure of information consumption on feed-based social media. While we empirically demonstrate that information consumption on these platforms is more politically homogeneous than prior empirical accounts, we view the networked curation processes that produce these results as a feature of democratic participation as much as others might view the resulting polarized consumption as a normative concern. Social media is, at its core, social, allowing users to *use* information to perform their identities and advance their interests in the context of democratic participation. To the extent to which these identities and interests diverge – particularly among the most politically engaged, who are the most likely to perform opinion-leading functions on social media (Hughes 2019) – so too will the information that circulates among different audiences. While much of the literature takes polarized information consumption as distressing for democracy, it is not obvious that this, in and of itself, is a problem to solve (Kreiss and McGregor 2023). In this sense, these findings underscore the longstanding tradeoffs between exposure to opposing views and democratic participation (Mutz 2006; Stroud 2011) – with different sites at which individuals express themselves and exchange their views being better suited for one or the other.

While the analyses here focus on audience partisanship, our theoretical framework problematizes source-level analyses of information consumption on social media more generally. For example, with respect to the study of political misinformation, preliminary evidence indicates that users interested in promoting false or misleading narratives often strategically repurpose factually true information from reliable sources in order to do so (Goel et al. 2024). Domain-level measures of political information cannot detect this behavior, but it naturally follows from individuals engaging with and *using* information on social media to perform their identities and advance their interests. Furthermore, individuals who report low levels of trust in mainstream sources on surveys may base these evaluations more on the sources' reputations than their specific interactions with information those sources produce (Peterson and Kagalwala 2021), and likely still recognize that such sources are perceived as credible by others (see also Pennycook and Rand 2019). Despite their stated distrust in mainstream sources overall, these individuals may nevertheless find specific information from these sources useful when it suits their purposes (Baum and Groeling 2009). Accounting for networked curation is crucial for aligning theory and measurement on large-scale platforms where such affordances are available.

Finally, it is important to note that socio-technical systems are elastic, and that different design choices may lead to different outcomes (Bail 2021). For example, emerging social media, such as TikTok, have de-emphasized the curation influence of followed accounts to rely more directly on the estimated relevance of specific pieces of content. The variation of platform features and affordances suggests a promising line of future research in examining the dynamics and democratic outcomes of networked curation across different platforms. We believe that content choice for information consumers is permanently expanded relative to the 20th century. While curation bubbles may pop, the process of networked curation connecting people to content they want to see from a set of vast choices is a permanent feature of the information landscape.

DATA AVAILABILITY STATEMENT

The raw data underlying this article cannot be shared due to privacy concerns arising from matching data to administrative records, data use agreements, and platforms' terms of service. Research documentation (including all code) and secondary Twitter data that preserves user anonymity are openly available at the American Political Science Review Dataverse: https://doi.org/10.7910/DVN/1ONKDX.

ACKNOWLEDGMENTS

The authors would like to thank Taylor Carlson; Bruce Desmarais; Shannon McGregor; C. Daniel Myers; participants at the 2021 meetings of the American Political Science Association, Midwest Political Science Association, and Society for Political Methodology; participants at the Duke Behavior and Identities Workshop; and three anonymous reviewers for helpful feedback on earlier versions of this manuscript. The authors would also like to thank Ronald Robertson for making reference data available and John Harrington for research assistance.

FUNDING STATEMENT

Stefan McCabe was supported by the John S. and James L. Knight Foundation through a grant to the Institute for Data, Democracy & Politics at the George Washington University. Shuyang Cao is supported by a Bloomberg Data Science Ph.D. Fellowship. David Lazer acknowledges support from the William & Flora Hewlett Foundation and the Volkswagen Foundation.

CONFLICT OF INTEREST

The authors declare no ethical issues or conflicts of interest in this research.

ETHICAL STANDARDS

Facebook Data in this study were obtained from Meta, as part of Facebook Open Research & Transparency (FORT), an initiative to facilitate the study of social media's impact on society. Researchers seeking permission to use the FORT Platform must 1) apply to become an approved partner and 2) sign the Research Data Agreement (RDA), a publicly available legal agreement. The RDA prohibits sharing Facebook Data with any third-party. Researchers may request access to Facebook Data at: https://socialscience.one/rfps.

Collection of Twitter data and linkage to administrative records was approved by the Institutional Review Board at Northeastern University (#17-12-13).

REFERENCES

Bail, Christopher. 2021. *Breaking the Social Media Prism: How to Make our Platforms Less Polarizing*. Princeton, NJ: Princeton University Press.

Bakshy, Eytan, Solomon Messing, and Lada Adamic. 2015. "Exposure to ideologically diverse news and opinion on Facebook." *Science* 348 (6239): 1130–1132. https://doi.org/10.1126/science.aaa1160.

Bakshy, Eytan, Itamar Rosenn, Cameron Marlow, and Lada Adamic. 2012. "The Role of Social Networks in Information Diffusion," 519–528. International Conference on the World Wide Web. Lyon, FRA. https://doi.org/10.1145/2187836.2187907.

Bandy, Jack, and Nicholas Diakopoulos. 2021. "More Accounts, Fewer Links: How Algorithmic Curation Impacts Media Exposure in Twitter Timelines." In *ACM on Human-Computer Interaction.* CSCW1. New York, NY. https://osf.io/p59wd.

Barberá, Pablo. 2015. “Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data.” *Political Analysis* 23 (1): 76–91.

———. 2020. “Social Media, Echo Chambers, and Political Polarization.” In *Social Media and Democracy: The State of the Field (eds. Nate Persily and Josh Tucker).* Cambridge, UK: Cambridge University Press.

Barberá, Pablo, John Jost, Jonathan Nagler, Joshua Tucker, and Richard Bonneau. 2015. “Tweeting from Left to Right: Is Online Political Communication More Than an Echo Chamber?” *Psychological Review* 26 (10): 1531–1542.

Baum, Matthew, and Tim Groeling. 2009. “Shot by the Messenger: Partisan Cues and Public Opinion Regarding National Security and War.” *Political Behavior* 31 (2): 157–186.

Billard, Thomas. 2021. “Deciding What’s (Sharable) News: Social Movement Organizations as Curating Actors in the Political Information System.” *Communication Monographs* 89 (3): 1–22.

Bode, Leticia. 2016. “Pruning the news feed: Unfriending and unfollowing political content on social media.” *Research & Politics* 3 (3). https://doi.org/10.1177/2053168016661873.

Broockman, David, and Joshua Kalla. 2023. *Selective Exposure and Partisan Echo Chambers In Television News Consumption: Evidence from Linked Viewership, Administrative, and Survey Data.* Preprint. Open Science Framework, April 14, 2023. Accessed May 2, 2023. https://doi.org/10.31219/osf.io/b54sx. https://osf.io/b54sx.

Buntain, Cody, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker. 2023. “A Summary Statistic for Political Ideology of Web Domains Using Differentially Private Engagement Data with Applications to Social Science One,” 17:72–83. Proceedings of the International AAAI Conference on Web and Social Media. Palo Alto, CA. Accessed June 26, 2023. https://ojs.aaai.org/index.php/ICWSM/article/view/22127.

Cinelli, Matteo, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. “The echo chamber effect on social media.” *Proceedings of the National Academy of Sciences* 118 (9).

Conover, Michael, Bruno Goncalves, Alessandro Flammini, and Filippo Menczer. 2012. “Partisan Asymmetries in Online Political Activity.” *EPJ Data Science* 1 (6).

Conover, Michael, Jacob Ratkiewics, Matthew Francisco, Bruno Goncalves, Filippo Menczer, and Alessandro Flammini. 2011. "Political Polarization on Twitter." *International AAAI Conference on Weblogs and Social Media* (Barcelona, ESP).

Dahlgren, Peter. 2021. "A critical review of filter bubbles and a comparison with selective exposure." *Nordicom Review* 42 (1): 15–33.

Davis, Jenny. 2017. "Curation: A Theoretical Treatment." *Information, Communication, & Society* 5:770–783.

Eady, Gregory, Jonathan Nagler, Andrew Guess, Jan Zilinsky, and Joshua Tucker. 2019. "How Many People Live in Political Bubbles on Social Media? Evidence From Linked Survey and Twitter Data." *SAGE Open* (SMaPP Global Special Issue-Original Research): 1–21.

Epstein, Ziv, Nathaniel Sirlin, Gordon Pennycook, and David Rand. 2023. "The social media context interferes with truth discernment." *Science Advances* 9 (9): abo6169.

Evans, Sandra K., Katy E. Pearce, Jessica Vitak, and Jeffrey W. Treem. 2017. "Explicating Affordances: A Conceptual Framework for Understanding Affordances in Communication Research." *Journal of Computer-Mediated Communication* 22 (1): 35–52. https://doi.org/10.1111/jcc4.12180.

Fischer, Sean, Kokil Jaidka, and Yphtach Lelkes. 2020. "Auditing local news presence on Google News." *Nature Human Behavior* 4:1236–1244.

Flaxman, Seth, Sharad Goel, and Justin Rao. 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly* 80 (51): 298–320.

Freedman, Jonathan, and David Sears. 1965. "Selective Exposure." *Advances in Experimental Social Psychology* 2:57–97.

———. 1967. "Selective Exposure to Information: A Critical Review." *Public Opinion Quarterly* 31:194–213.

Garimella, Kiran, Tim Smith, Rebecca Weiss, and Robert West. 2021. "Political Polarization in Online News Consumption." https://arxiv.org/abs/2104.06481.

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

Garrett, R. Kelly, and Paul Resnick. 2011. "Resisting Political Fragmentation on the Internet." *Daedalus* 140 (4): 108–120.

Garz, Marcel, Jil Sörensen, and Daniel F. Stone. 2020. "Partisan selective engagement: Evidence from Facebook." *Journal of Economic Behavior & Organization* 177 (September): 91–108. Accessed April 30, 2023. https://doi.org/10.1016/j.jebo.2020.06.016. https://linkinghub.elsevier.com/retrieve/pii/S0167268120302079.

Gentzkow, Matthew, and Jesse Shapiro. 2010. "What Drives Media Slant? Evidence from U.S. Daily Newspapers." *Econometrica* 78 (1): 35–71.

———. 2011. "Ideological Segregation Online and Offline." *The Quarterly Journal of Economics* 126 (4): 1799–1839. https://doi.org/10.1093/qje/qjr044.

Goel, Pranav, Jon Green, David Lazer, and Philip Resnik. 2024. *Misinformation is more than "fake news": Using co-sharing to identify use of mainstream news for promoting misinformation narratives.* https://www.researchsquare.com/article/rs-4427772/v1.

González-Bailón, Sandra, David Lazer, Pablo Barberá, Meiqing Zhang, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, et al. 2023. "Asymmetric ideological segregation in exposure to political news on Facebook." *Science* 381 (6656): 392–398.

Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. "Fake news on Twitter during the 2016 U.S. presidential election." *Science* 363 (6425): 374–378.

Guess, Andrew. 2021. "(Almost) Everything in Moderation: New Evidence on Americans' Online Media Diets." *American Journal of Political Science* 65 (4): 1007–1022. https://doi.org/10.1111/ajps.12589. https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12589.

Halberstam, Yosh, and Brian Knight. 2016. "Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter." *Journal of Public Economics* 143:73–88.

Ho, Daniel, and Kevin Quinn. 2008. "Measuring Explicit Political Positions of Media." *Quarterly Journal of Political Science* 3:353–377.

Hughes, Adam. 2019. *A small group of prolific users account for a majority of political tweets sent by U.S. adults.* Accessed October 23, 2019. https://www.pewresearch.org/short-reads/2019/10/23/a-small-group-of-prolific-users-account-for-a-majority-of-political-tweets-sent-by-u-s-adults/.

Hughes, Adam, Stefan McCabe, William Hobbs, Emma Remy, Sono Shah, and David Lazer. 2021. "Using administrative records and survey data to construct samples of Tweeters and Tweets." *Public Opinion Quarterly* 85 (S1): 323–346.

Huszár, Ferenc, Sofia Ira Ktena, Conor O'Brien, Luca Belli, Andrew Schlaikjer, and Moritz Hardt. 2022. "Algorithmic amplification of politics on Twitter." *Proceedings of the National Academy of Sciences* 119 (1): e2025334119. https://doi.org/10.1073/pnas.2025334119. https://www.pnas.org/doi/abs/10.1073/pnas.2025334119.

Jurkowitz, Mark, and Jeffrey Gottfried. 2022. "Twitter is the go-to social media site for U.S. journalists, but not for the public," June 27, 2022. https://www.pewresearch.org/short-reads/2022/06/27/twitter-is-the-go-to-social-media-site-for-u-s-journalists-but-not-for-the-public/.

Katz, Elihu, and Paul Lazarsfeld. 1955. *Personal Influence: The Part Played by People in the Flow of Mass Communications.* London, UK: Routledge.

Kim, Jin Woo, and Eunji Kim. 2021. "Temporal Selective Exposure: How Partisans Choose When to Follow Politics." *Political Behavior* 43 (4): 1663–1683. https://doi.org/10.1007/s11109-021-09690-1.

King, Gary, and Nathaniel Persily. 2020. "A New Model for Industry–Academic Partnerships." *PS: Political Science & Politics* 53, no. 4 (October): 703–709. https://doi.org/10.1017/S1049096519001021.

Kobayashi, Tetsuro, and Ken'ichi Ikeda. 2009. "Selective Exposure in Political Web Browsing." *Information, Communication, and Society* 12 (6): 929–953.

Kreiss, Daniel, and Shannon McGregor. 2023. "A review and provocation: On polarization and platforms." *New Media & Society* 26 (1): 556–579.

Lazarsfeld, Paul, Bernard Berelson, and Hazel Gaudet. 1948. *The People's Choice: How the Voter Makes up His Mind in a Presidential Campaign.* New York NY: Columbia University Press.

Lazer, David, Brian Rubineau, Carol Chetkovich, Nancy Katz, and Michael Neblo. 2010. "The Coevolution of Networks and Political Attitudes." *Political Communication* 27 (3): 248–274. https://doi.org/10.1080/10584609.2010.500187.

Marwick, Alice E., and danah boyd. 2011. "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience." *New Media & Society* 13 (1): 114–133.

McGregor, Shannon, and Logan Molyneux. 2018. “Twitter’s influence on news judgment: An experiment among journalists.” *Journalism* 21 (5): 597–613.

Messing, Solomon, Christina DeGregorio, Bennett Hillenbrand, Gary King, Saurav Mahanti, Zagreb Mukerjee, Chaya Nayak, Nathaniel Persily, Bogdan State, and Arjun Wilkins. 2021. *Facebook Privacy-Protected Full URLs Data Set.* Accessed January 3, 2022. https://dataverse.harvard.edu/file.xhtml?fileId=5212847&version=8.0.

Messing, Solomon, and Sean Westwood. 2014. “Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online.” *Communication Research* 41 (8): 1042–1063.

Minozzi, William, Hyunjin Song, David M. J. Lazer, Michael A. Neblo, and Katherine Ognyanova. 2020. “The Incidental Pundit: Who Talks Politics with Whom, and Why?” *American Journal of Political Science* 64 (1): 135–151. https://doi.org/10.1111/ajps.12469.

Molyneux, Logan, and Shannon McGregor. 2021. “Legitimating a platform: evidence of journalists’ role in transferring authority to Twitter.” *Information, Communication, & Society* 25 (11): 1577–1595.

Mummolo, Jonathan. 2016. “News from the other side: How topic relevance limits the prevalence of partisan selective exposure.” *Journal of Politics* 78 (3): 763–773.

Mutz, Diana. 2006. *Hearing the Other Side: Deliberative Versus Participatory Democracy.* Cambridge, UK: Cambridge University Press.

Neely, Stephen R. 2021. “Politically Motivated Avoidance in Social Networks: A Study of Facebook and the 2020 Presidential Election.” *Social Media + Society* 7 (4): 20563051211055438.

Osmundsen, Mathias, Alexander Bor, Peter Bjerregaard Vahlstrup, Anja Bechmann, and Michael Bang Petersen. 2021. “Partisan Polarization Is the Primary Psychological Motivation behind Political Fake News Sharing on Twitter.” *American Political Science Review,* https://doi.org/10.1017/S0003055421000290.

Pariser, Eli. 2011. *The Filter Bubble: What the Internet is Hiding from You.* London, UK: Penguin UK.

Park, Chang Sup, and Barbara Kaye. 2018. “News Engagement on Social Media and Democratic Citizenship: Direct and Moderating Roles of Curatorial News Use in Political Involvement.” *Journalism & Mass Communication Quarterly* 95 (4): 1103–1127.

Pennington, Jeffrey, Richard Socher, and Christopher Manning. 2014. "GloVe: Global Vectors for Word Representation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing,* 1532–1543. EMNLP 2014. Stroudsburg, PA: Association for Computational Linguistics. https://doi.org/10.3115/v1/D14-1162.

Pennycook, Gordon, and David Rand. 2019. "Fighting misinformation on social media using crowdsourced judgments of news source quality." *Proceedings of the National Academy of Sciences* 116 (7): 2521–2526. https://doi.org/10.1073/pnas.1806781116. https://www.pnas.org/doi/10.1073/pnas.1806781116.

Peterson, Erik, Sharad Goel, and Shanto Iyengar. 2021. "Partisan selective exposure in online news consumption: evidence from the 2016 presidential campaign." *Political Science Research and Methods* 9 (2): 242–258.

Peterson, Erik, and Shanto Iyengar. 2021. "Partisan Gaps in Political Information and Information-Seeking Behavior: Motivated Reasoning or Cheerleading?" *American Journal of Political Science* 65 (1): 133–147.

Peterson, Erik, and Ali Kagalwala. 2021. "When Unfamiliarity Breeds Contempt: How Partisan Selective Exposure Sustains Oppositional Media Hostility." *American Journal of Political Science* 115 (2): 585–598.

Prior, Markus. 2013. "Media and Political Polarization." *Annual Review of Political Science* 16:101–127.

Rainey, Carlisle. 2014. "Arguing for a Negligible Effect." *American Journal of Political Science* 58 (4): 1083–1091. https://doi.org/10.1111/ajps.12102.

Ribeiro, Manoel Horta, Raphael Ottoni, Robert West, Virgílio A. F. Almeida, and Wagner Meira. 2020. "Auditing Radicalization Pathways on YouTube," 131–141. Conference on Fairness, Accountability, and Transparency. Barcelona, ESP. https://doi.org/10.1145/3351095.3372879. https://doi.org/10.1145/3351095.3372879.

Robertson, Ronald E., Jon Green, Damian J. Ruck, Katherine Ognyanova, Christo Wilson, and David Lazer. 2023. "Users Choose to Engage with More Partisan News Than They Are Exposed to on Google Search." *Nature* (May 24, 2023). Accessed June 5, 2023. https://doi.org/10.1038/s41586-023-06078-5. https://www.nature.com/articles/s41586-023-06078-5.

Robertson, Ronald E., Shan Jiang, Kenneth Joseph, Lisa Friedland, and David Lazer. 2018. "Auditing Partisan Audience Bias within Google Search." In *ACM on Human-Computer Interaction.* New York, NY. https://doi.org/10.1145/3274417.

Settle, Jamie. 2018. *Frenemies: How Social Media Polarizes America.* Cambridge, UK: Cambridge University Press.

Shugars, Sarah, Adina Gitomer, Stefan McCabe, Ryan J. Gallagher, Kenneth Joseph, Nir Grinberg, Larissa Doroshenko, Brooke Foucault Welles, and David Lazer. 2021. "Pandemics, Protests, and Publics: Demographic Activity and Engagement on Twitter in 2020." *Journal of Quantitative Description: Digital Media* 1. https://doi.org/10.51685/jqd.2021.002.

Stroud, Naomi Jomini. 2011. *Niche News: The Politics of News Choice.* Oxford, UK: Oxford University Press.

Sunstein, Cass. 2002. *Republic.com.* Princeton, NJ: Princeton University Press.

———. 2017. *#Republic.* Princeton, NJ: Princeton University Press.

Van Bavel, Jay, Elizabeth Harris, Philip Pärnamets, Steve Rathje, Kimberly Doell, and Joshua Tucker. 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.

Weeks, Brian, Daniel Lane, Dam Hee Kim, Slgi Lee, and Nojin Kwak. 2017. "Incidental Exposure, Selective Exposure, and Political Information Sharing: Integrating Online Exposure Patterns and Expression on Social Media." *Journal of Computer-Mediated Communication* 22 (6): 363–379.

Wojcieszak, Magdalena, Andreu Casas, Xudong Yu, Jonathan Nagler, and Joshua A. Tucker. 2022. "Most users do not follow political elites on Twitter; those who do show overwhelming preferences for ideological congruity." *Science Advances* 8, no. 39 (September 30, 2022): eabn9418. Accessed October 6, 2022. https://doi.org/10.1126/sciadv.abn9418. https://www.science.org/doi/10.1126/sciadv.abn9418.

Wojcik, Stefan, and Adam Hughes. 2019. "Sizing Up Twitter Users." Pew Research, April 24, 2019. Accessed May 19, 2023. https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/.

APPENDIX A: STYLIZED EXAMPLE

We illustrate our theoretical framework with a stylized example, considering a hypothetical news source that produces fifteen stories, each of which are shared fifteen times (for 225 sharing events in total). Given this number of shares, there are 16 possible scores for each individual story, ranging from -1 if all fifteen shares are from Democratic users to 1 if all shares are from Republican users. The domain-level score could be calculated from any one of roughly 115 million combinations (with replacement). Across all these possible sharing patterns there is exactly one way for a domain to receive a partisanship score of -1 or 1 -- if all 15 shares for all 15 stories come from Democrats or Republicans, respectively. There are more than five million combinations of sharing patterns that could produce a domain score of zero.

As a consequence, there is a greater potential for variation in story-level audience scores that produce identically moderate domain-level audience scores than there is in story-level audience scores that produce more extreme domain-level scores. This is shown in Figure A1, which plots each of these potential sharing patterns with respect to the domain-level audience score we would observe (x-axis) against the variance in story-level audience score we would observe (y-axis). As the audience score moves toward either extreme, there are fewer combinations of sharing patterns that can produce it, and therefore the potential story-level variance decreases. This is instructive for thinking about the data generating process for the domain-level audience scores we observe. Sources with audience scores closer to zero, by definition, have more potential for the mean to obscure important story-level variation.

FIGURE A.1. Simulated Domain Audience Score Estimate and Variance

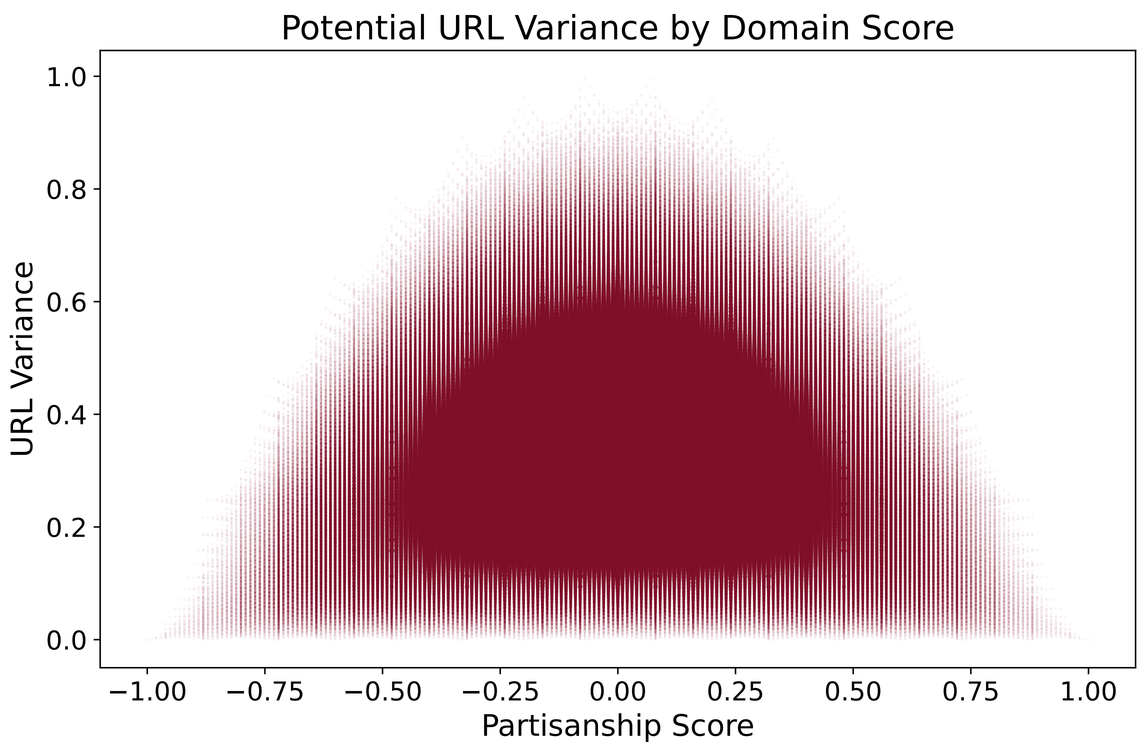

APPENDIX B: CODING PARTISAN APPEAL OF STORIES

To test the partisan appeal of stories, we selected a sample of URLs for coding. Nine uncompensated volunteer coders—graduate students and postdocs—were presented with the following instructions:

> In this survey you will be asked to code news stories based on whether or not you think they would appeal to Democrats or Republicans. These stories are all a few years old, dating from 2017 and 2018, so as a reminder about the political situation then:
>
> - The President of the United States was a **Republican**, **Donald Trump**.
> - The Senate had a **Republican** majority, led by Majority Leader **Mitch McConell**.
> - The House of Representatives had a **Republican** majority, led by Speaker of the House **Paul Ryan**.
>
> In each item, you will be shown a headline and (if available) the extra text that would have been displayed if the story was shared on Facebook or Twitter. If you need more information to make a decision, you can click on the URL and it will open in a new browser tab.

An example coding item is presented below:

> Do you think this article is more likely to appeal to Democrats or Republicans, or both equally?
>
> Headline:
>
> **Trump administration tells EPA to cut climate page from website: sources**
>
> Snippet:
>
> U.S. President Donald Trump's administration has instructed the Environmental Protection Agency to remove the climate change page from its website, two agency employees told Reuters, the latest move by the newly minted leadership to erase ex-President Barack Obama's climate...
>
> If you need more context, click here to open the URL.

We sampled 1,000 URLs for coding, with each person coding 500. Coders had four options: "Appeals to Democrats" (-1), "Appeals to both equally" (0), "Appeals to Republicans" (1), and "Don't Know" (NA).

To avoid source cues dominating the task, we sampled URLs with a probability weighted by the absolute value of the difference between its (Twitter-based) URL score and domain score. This was a relatively difficult coding task; as a result the intercoder reliability was fairly low (Krippendorf's α = 0.602). One coder exhibited notably lower correlations with other coders; after dropping this set of codings (down to eight coders), reliability improved to acceptable levels (Krippendorf's α = 0.673). We averaged the eight coders' decisions to produce human-coded partisan appeal scores.

APPENDIX C: ADDITIONAL DETAILS ON POLITICAL CLASSIFIER

TABLE C.1. Example *Wall Street Journal* headlines and probabilities of being political.

Score	Headline
0	Google and Salesforce Ink Cloud, Apps Deal
0.1	Alexa, Stop Making Life Miserable for Anyone With a Similar Name!
0.2	Verizon, a Century-Old Company, Sets Its Sights on Millennials
0.3	Catholic Church Considers Married Priests to Ease Amazon Clergy Shortage
0.4	SC Johnson CEO Gives $150 Million to Cornell's Business College
0.5	Beyond Soda: How and Why Your Beverage Options Are Exploding
0.6	Wall Street Made Charles Murphy Successful and Rich, but Happiness Eluded Him
0.7	How to Keep Facebook From Oversharing Your Info
0.8	Just Say No to Push Notifications
0.9	Appeals Court Rules Against Trump on Canceling DACA Protections

Note: Only articles with a score of greater than 0.9 were included in the analysis. Scores are floored to the nearest 0.1 increment for clarity.

TABLE C.2. Percent of URLs with blurbs collected and percent classified as political, select domains.

Domain	% with blurb	% political
axios.com	90%	69%
breitbart.com	100%	82%
buzzfeednews.com	99%	63%
cbsnews.com	100%	39%
dailykos.com	99%	83%
dailymail.co.uk	100%	54%
foxnews.com	96%	75%
hollywoodreporter.com	100%	18%
huffpost.com	99%	73%
inc.com	98%	10%
mediaite.com	93%	79%
nationalreview.com	98%	76%
nypost.com	100%	53%
nytimes.com	99%	48%
politifact.com	91%	57%
realclearpolitics.com	98%	91%
reason.com	60%	43%
redstate.com	97%	64%
rollcall.com	96%	92%
rt.com	97%	74%
usatoday.com	98%	59%
washingtontimes.com	99%	92%
wsj.com	99%	54%

APPENDIX D: COMPARISON TO EARLIER METHODS

FIGURE D.1. Domain score comparison to Robertson et al (2018).

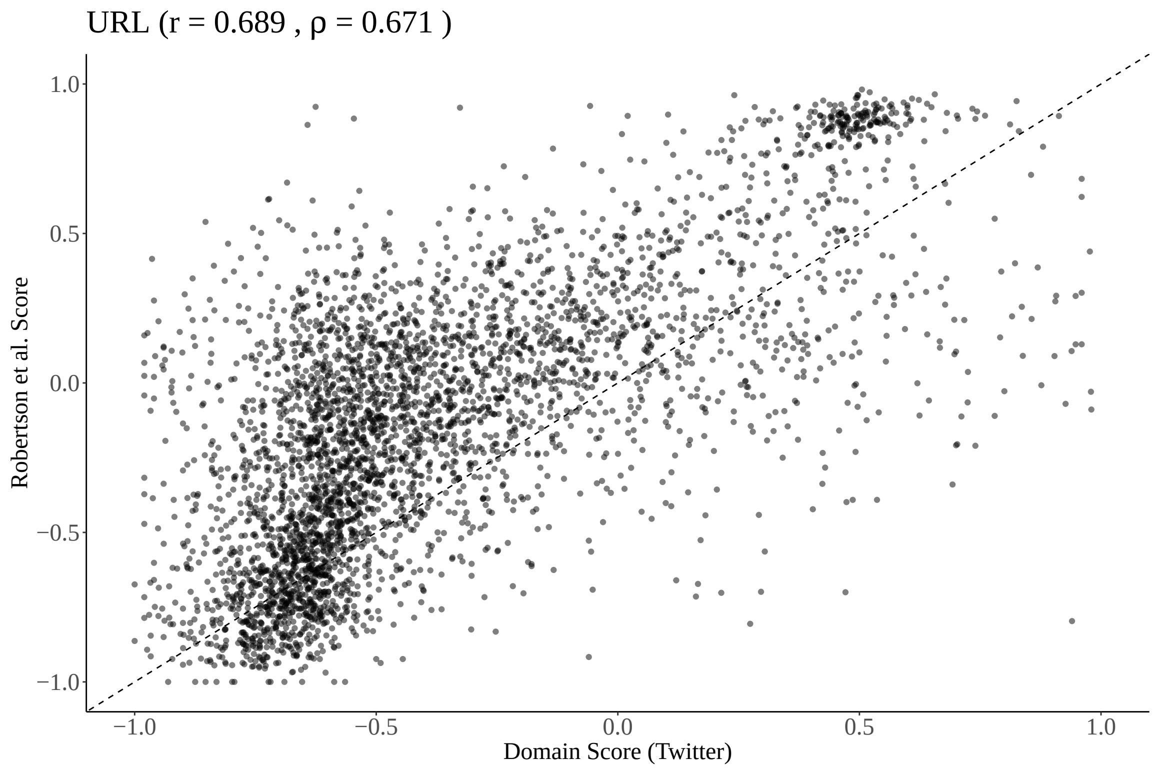

Note: Domains with 0% political URLs excluded.

FIGURE D.2. Domain score comparison to Bakshy et al (2015).

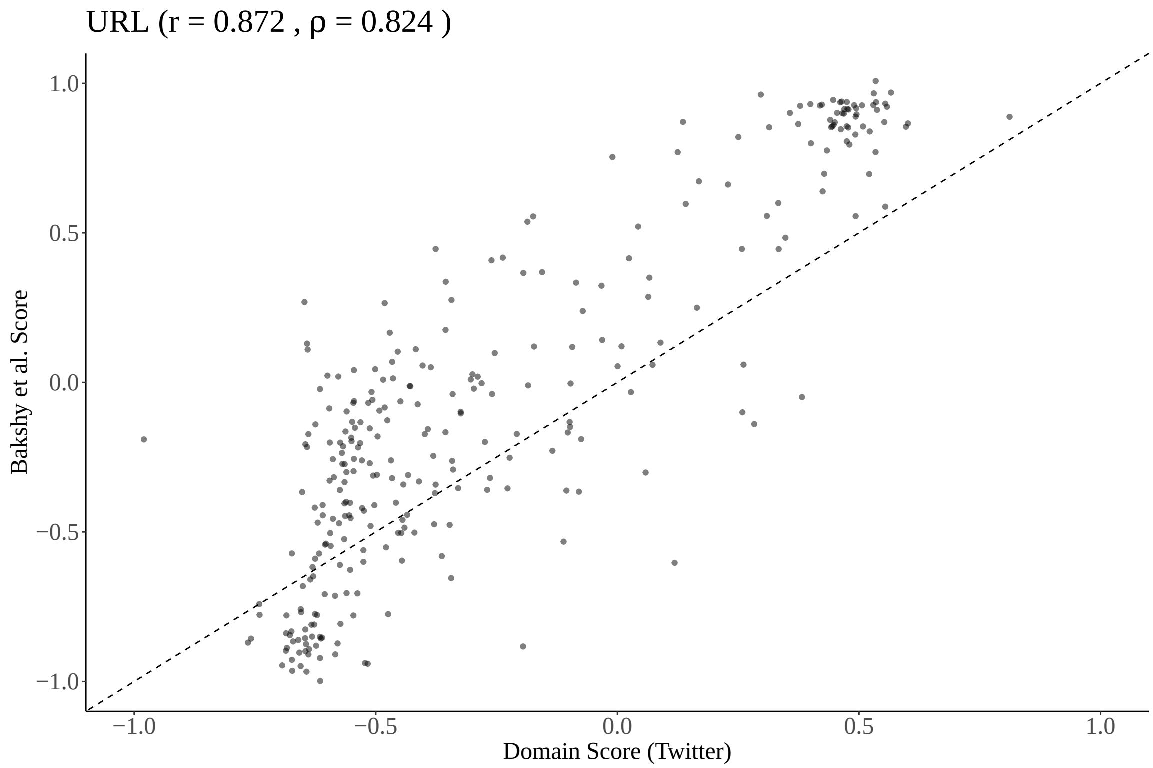

Note: Domains with 0% political URLs excluded.

APPENDIX E: REGRESSION RESULTS FOR HAND-CODED URLS

TABLE E.1. Models predicting human-evaluated partisan appeal.

	Full	URL only	Domain only
(Intercept)	0.154	0.058	0.201
	(0.024)	(0.019)	(0.030)
Domain score	0.362		1.192
	(0.056)		(0.058)
URL score	1.026	1.185	
	(0.042)	(0.034)	
Num.Obs.	972	972	972
R2	0.572	0.554	0.305
R2 Adj.	0.571	0.554	0.304
AIC	1433.9	1472.2	1904.4
BIC	1453.4	1486.8	1919.1
Log.Lik.	−712.935	−733.079	−949.207
RMSE	0.50	0.51	0.64

Note: OLS regression models for URL-only, domain-only, and URL-plus-domain models. The variable being predicted is the human-evaluated partisan appeal of the URL; the predictor variables are the (Twitter-based) URL score and/or domain score.

APPENDIX F: SIMULATION APPROACH

In the main text, we have preferred a parametric approach to determining if a URL's audience can be statistically distinguished from the audience of the domain as a whole. This is scalable across platforms and it draws on social scientists' existing intuitions around significance testing. However, as mentioned in the text, it does use a few approximations—sufficient sample size and unbounded support—and, as such, could be inaccurate for some tests.

In this Appendix, we present an alternative simulation-based approach and apply it to the Twitter dataset. In this approach, we compare the observed audience score for each URL to the range of scores that we would expect to observe if this audience was randomly sampled from that of the domain as a whole. Specifically, for each URL, we construct its vector of partisan score-shares. For each domain, we then concatenate these into single vectors representing all of the sharing events across the whole domain. For each URL, we then simulate a range of hypothetical audience scores by sampling sharing events from the domain-level vector, with replacement, as many times as the URL was shared – repeating this sampling 1000 times. So for example, if a URL was shared by three individuals with a partisan score of 1, two with a score of 2, and five with a score of 3, we would construct a vector of $[1, 1, 1, 2, 2, 3, 3, 3, 3, 3]$. We would then take ten draws, with replacement, from the sharing events for the URL's parent domain, repeating 1000 times – providing us with a non-parametric uncertainty interval. If a URL's observed audience score falls outside of this interval – that is, if it is more distant from its domain's score than we would expect to observe by randomly sampling from that domain's broader audience – we infer that the URL is in a partisan curation bubble. For a given domain, this allows us to quantify how consistently its associated stories circulate among (a)typical partisan audiences.

We reproduce Figures 4 and 7 below, using this alternative approach. Our substantive conclusions are unchanged.

FIGURE F.1. URL scores by share volume for selected domains on Twitter, simulation approach.

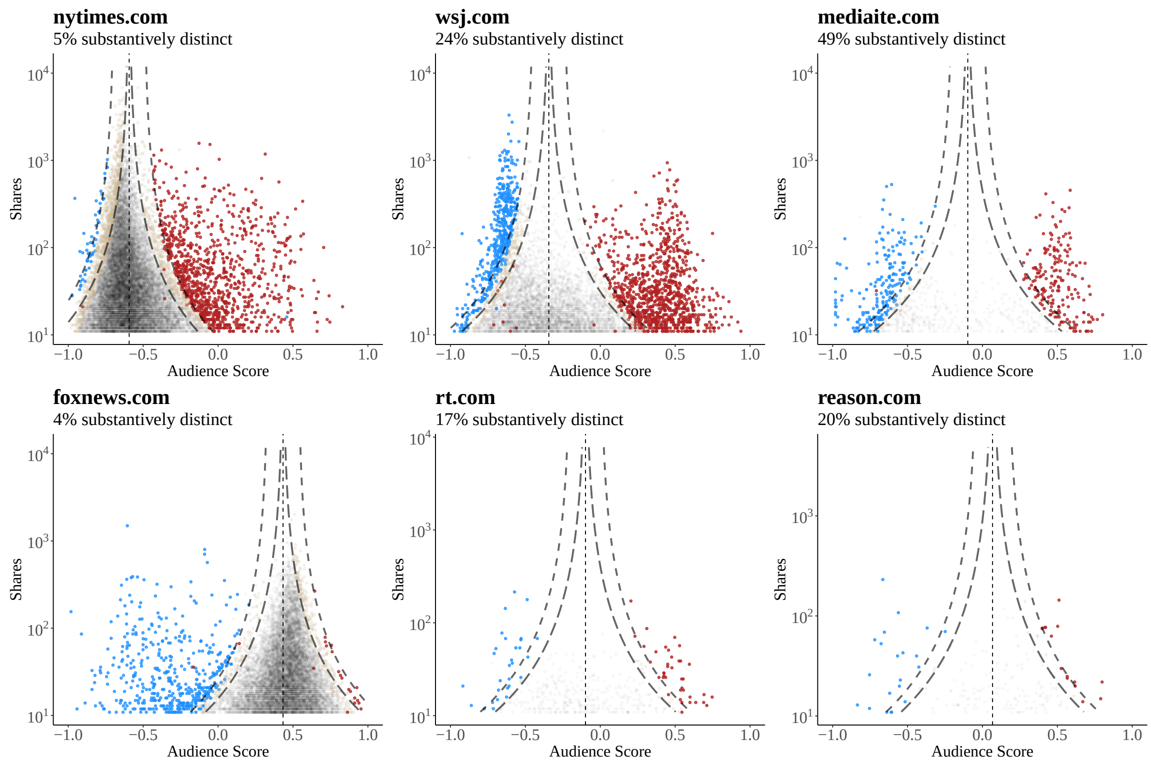

Note: Solid line indicates loess curve of best fit.

FIGURE F.2. Proportion of URLs substantively distinct from domain by domain-level audience score for Twitter, simulation approach.

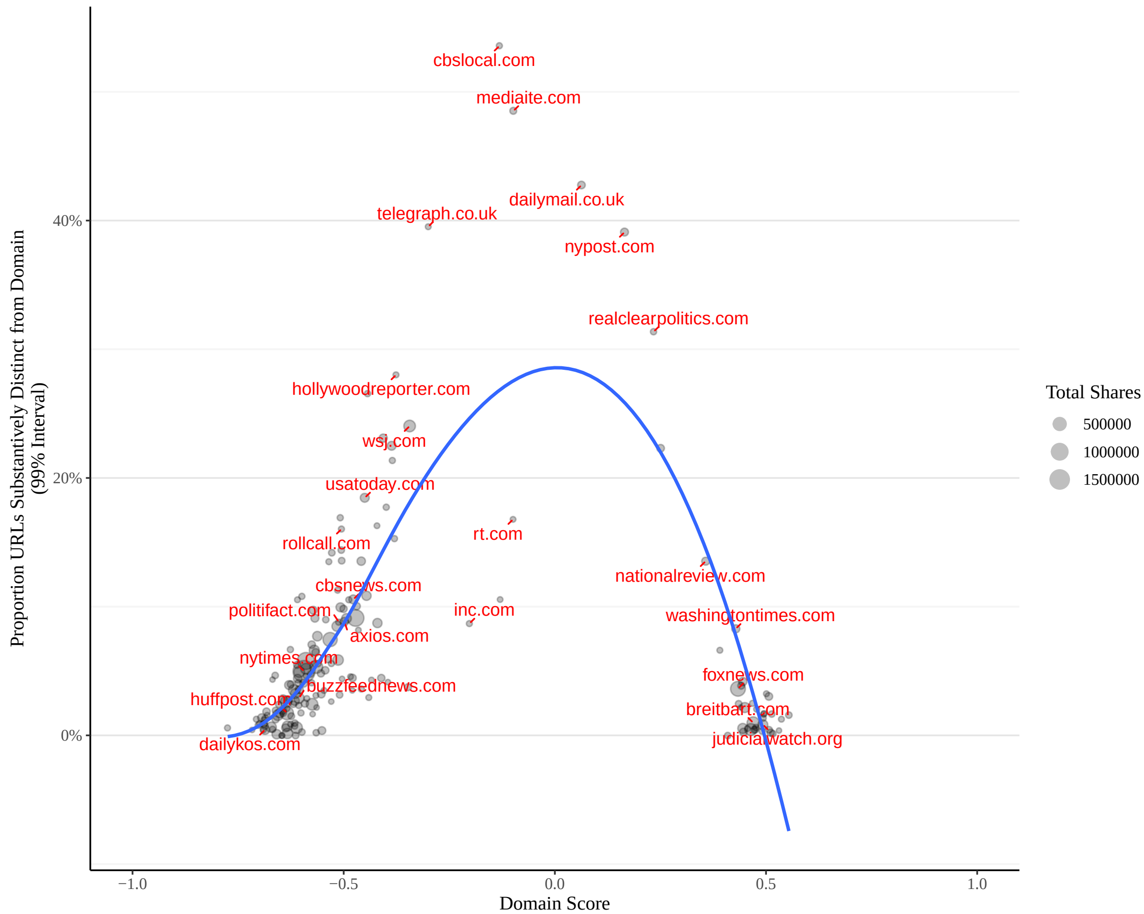

Note: Solid line indicates loess curve of best fit. Data limited to domains in the top quartile of unique political URLs.

APPENDIX G: COMPARING MODELED PARTISANSHIP AND PARTY REGISTRATION

In the main text, we present our (Twitter-based) analyses using individuals' modeled partisanship, the probability that the individual, if queried, would self-identify as a Democrat. This is the output of a proprietary machine learning model from 2017, trained on past vote history, partisan registration (where available), political contributions, demographics, geographic context, and so on. We have preferred to use this quantity, rather than party registration, because it provides a more granular estimate than party registration, and is more likely to be comparable across states.

We prefer to use these modeled estimates (as opposed to partisan registration) for two reasons. The first is that modeled partisanship is available for users in all states, while party registration is not available in many states (recall that some states, like Georgia, do not have party registration at all; using partisan registration to estimate audience scores would necessarily exclude sharing events from users who live in Georgia). The second is that using the numeric value representing the likelihood that a user identifies as a Democrat, rather than the categorical measure of Registered Democrat/Registered Republican/Neither, allows us to incorporate model uncertainty into our estimates. This reflects the intuition that a user with a modeled partisan score of 75 should not be counted the same as a user with a modeled partisan score of 95, even though both users may be registered to vote as Democrats.

We take two steps to show that this quantity is a useful measure of user partisanship. First, we take the full voter file and plot the county-level average propensity to identify as a Democrat against the county's two-way Democratic vote share in the 2016 presidential election. The results, which mirror the approach of Shugars et al. (2021), are shown in Figure G.1. At the county level, average modeled partisanship and two-way vote share are highly correlated ($r = 0.84$). The primary disagreements between the approaches are concentrated in predominantly White, rural counties in Appalachia with well-known differences between local and national political behavior; we have highlighted counties in Kentucky and West Virginia to illustrate this. The quantity being estimated in the proprietary model is the propensity to *identify* as a Democrat, not to *vote for* Democrats at the national level; these states, while often supporting Republicans for president, have long histories of local Democratic Party identification.

FIGURE G.1. Relationship between modeled partisanship and county-level vote choice.

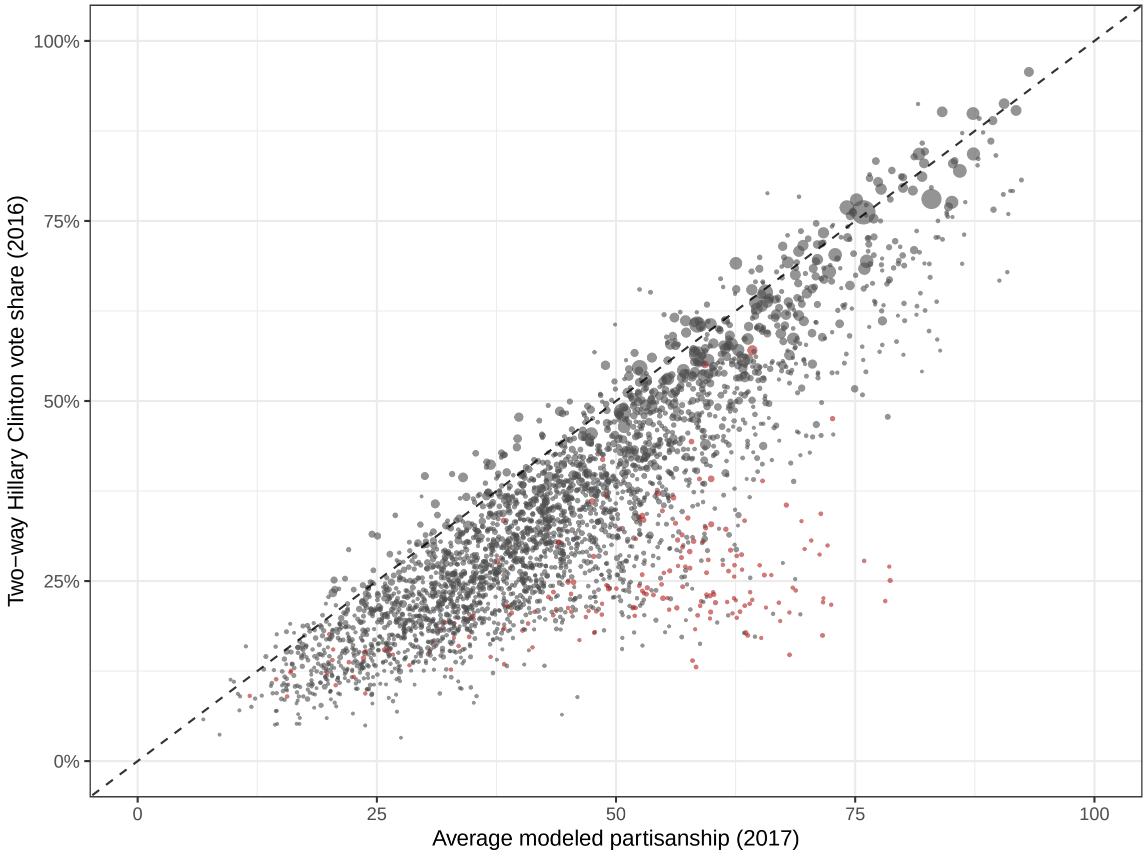

Note: Counties sized by number of voters. West Virginia and Kentucky highlighted in red. Alaska not included, as election results are not reported at the county level.

Second, we construct URL and domain scores using party registration instead of modeled partisanship. First, following Robertson et al. (2018), we do so using dichotomized party registration: a share from a registered Democrat counts as −1, a share from a registered Republican counts as +1, and independents do not contribute to the mean. The model-based and registration-based scores are highly correlated, as shown in Figure G.2; $r = 0.95$ at the URL level and $r = 0.88$ at the domain level. We also report coefficients regressing the continuous score used in the paper on the registration-based and an intercept term; for URLs, $\beta = 0.63$; for domains, $\beta = 0.61$.

To the extent to which these estimates differ, this is largely attributable to the exclusion of users who are registered to vote as something other than Democrats or Republicans (independents and those registered with minor parties) from the denominator. This tends to increase the extremity of audience scores estimated using the ratio of Democratic to Republican sharing events, relative to our estimates based on modeled probability of Democratic identification that include these users. When we include these users in a trichotomized measure of registration-based audience scores (Figure G.3), the raw correlations are nearly identical *and* the coefficients of the bivariate relationships are closer to 1.

FIGURE G.2. Comparison of model-based and dichotomized registration-based scores.

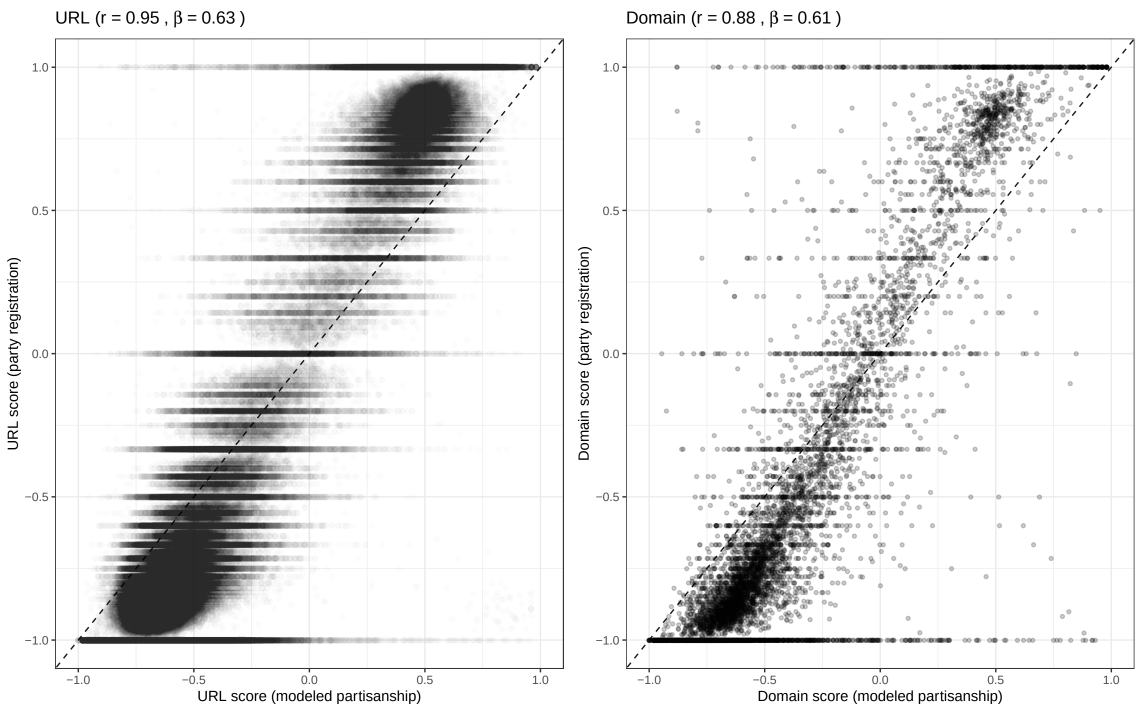

FIGURE G.3. Comparison of model-based and trichotomized registration-based scores.

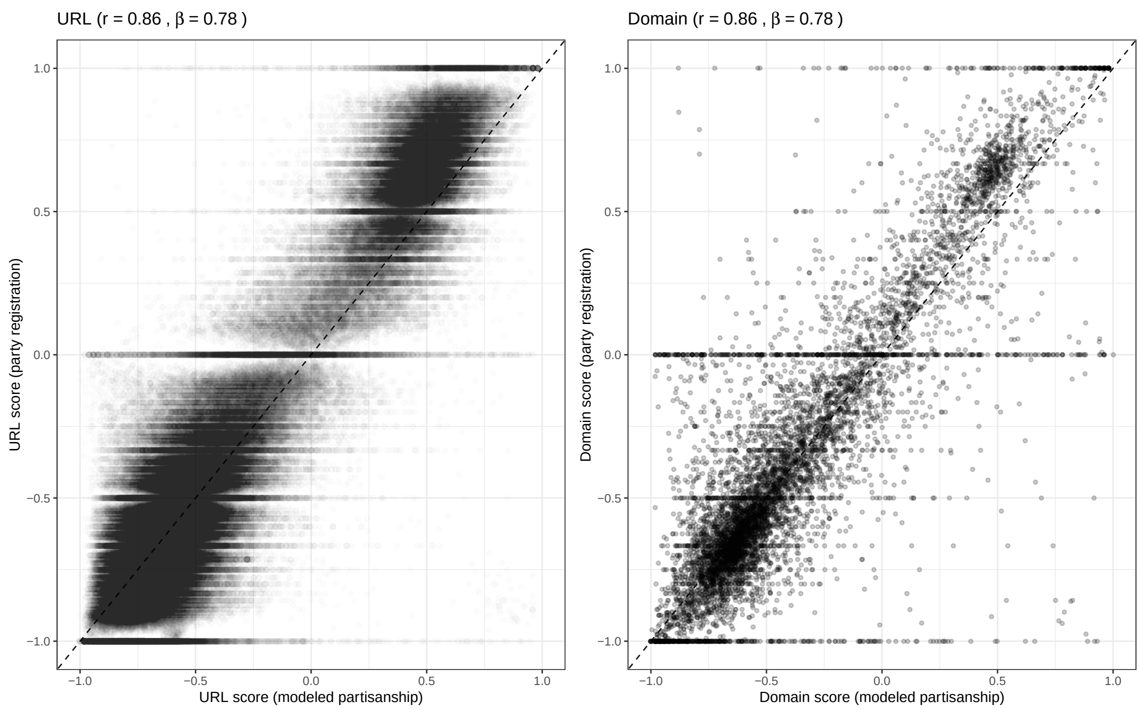

APPENDIX H: REGRESSION RESULTS FOR PROPORTION DISTINCT

TABLE H.1. Regression results for OLS models supporting Figures 7 and 8.

	FB Shares	FB View	FB Clicks	FB Reacts	TW Shares
Intercept	0.636***	0.561***	0.459***	0.748***	0.188***
	(0.080)	(0.076)	(0.081)	(0.086)	(0.035)
Domain Extremity	−0.972***	−0.663***	−0.870***	−1.014***	−0.401***
	(0.036)	(0.039)	(0.036)	(0.037)	(0.022)
Num. URLs	0.000	0.000*	0.000+	0.000	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	0.000
Num. Shares (logged)	0.030+	0.022	0.049**	0.020	0.024***
	(0.016)	(0.015)	(0.016)	(0.017)	(0.007)
Num.Obs.	780	780	780	780	717
R2	0.553	0.330	0.472	0.562	0.363
R2 Adj.	0.551	0.327	0.470	0.560	0.360
AIC	−252.7	−334.1	−231.8	−126.5	−1342.2
BIC	−229.4	−310.8	−208.5	−103.2	−1319.3
Log.Lik.	131.348	172.066	120.902	68.226	676.080
RMSE	0.20	0.19	0.21	0.22	0.09

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: OLS regression models. In each model, URL and domain scores have been constructed from the target metric (e.g., Twitter shares); the unit of analysis is the domain and the outcome variable is the proportion of URLs with scores substantively distinct from the domain score. Domain Extremity is the absolute value of the domain score for the action.

APPENDIX I: ALTERNATIVE THRESHOLDS

This appendix section varies the threshold of unique URLs per domain required for inclusion in Figures 2 and 5.

FIGURE I.1. Twitter distribution (no threshold)

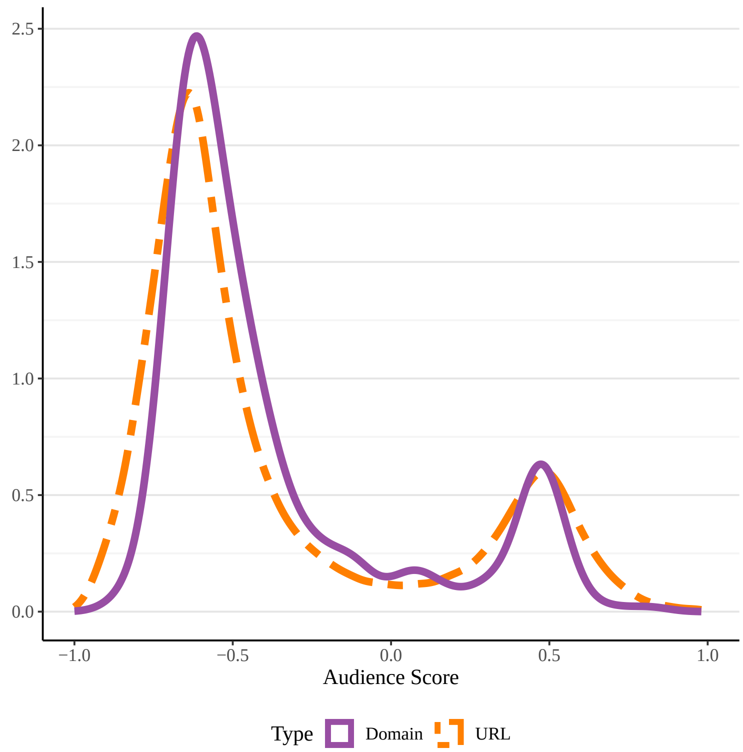

FIGURE I.2. Twitter distributions (median split)

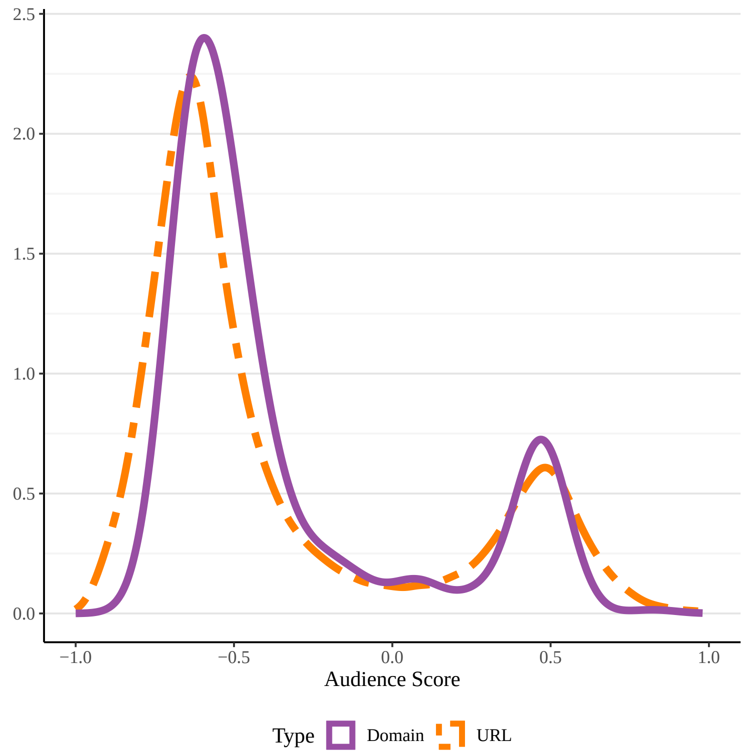

FIGURE I.3. Twitter distributions (high threshold)

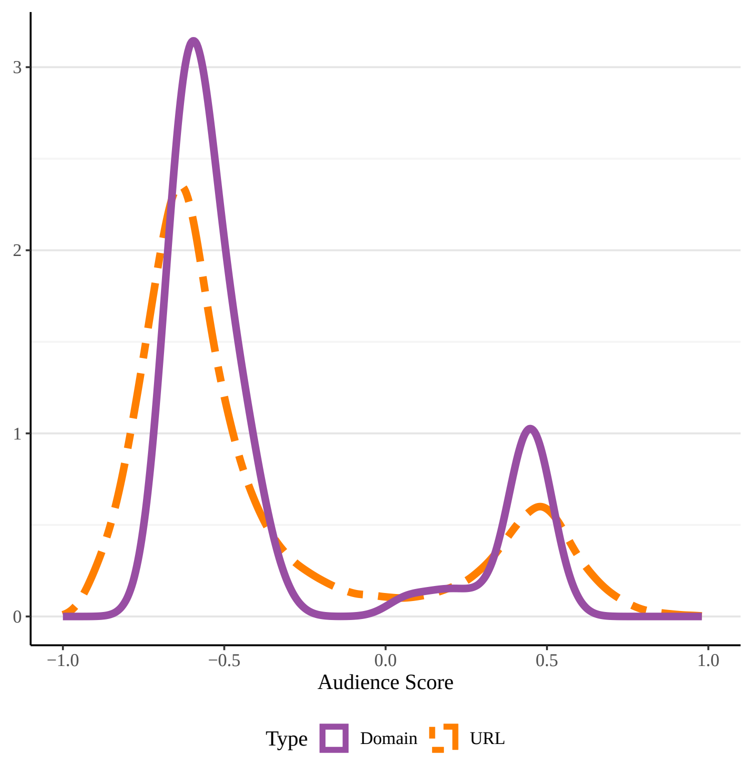

FIGURE I.4. Facebook distributions (no threshold)

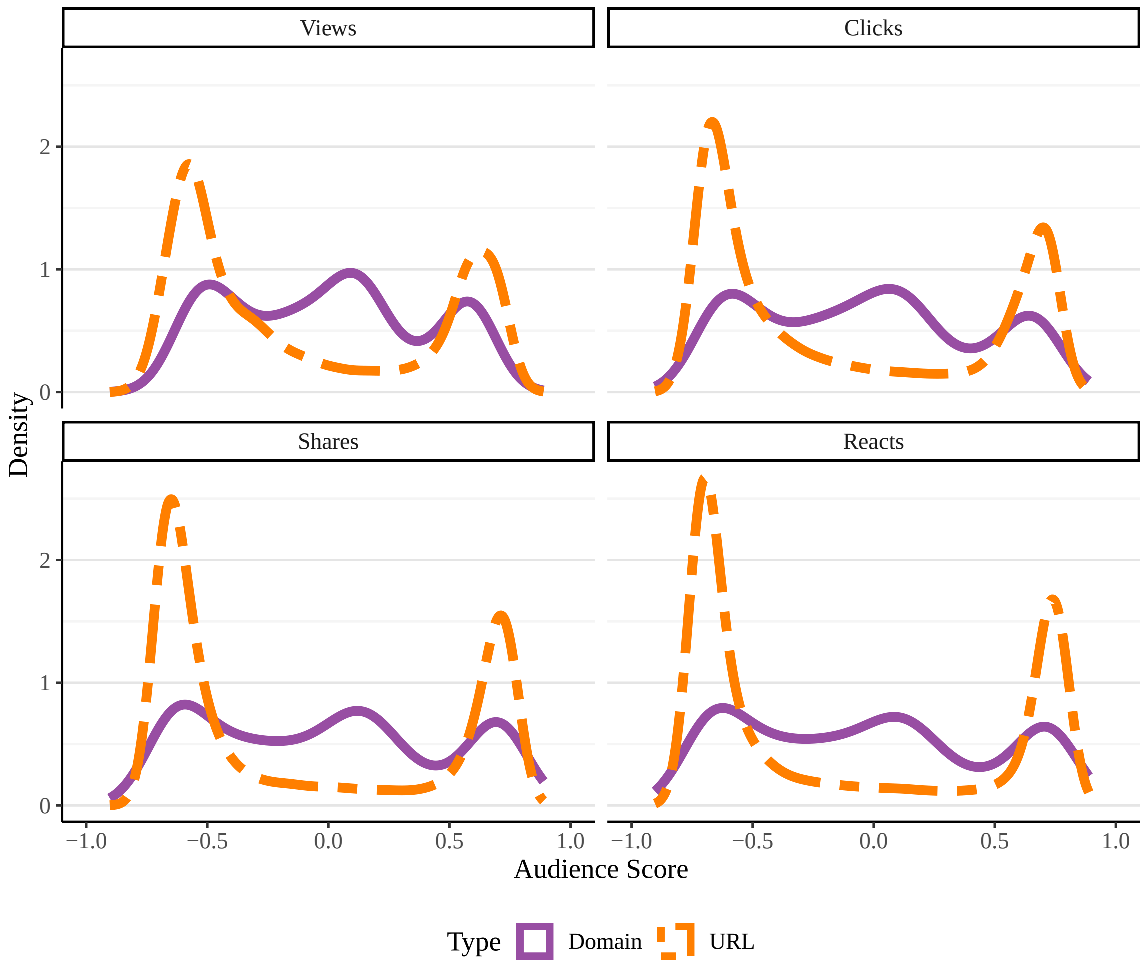

FIGURE I.5. Facebook distributions (median split)

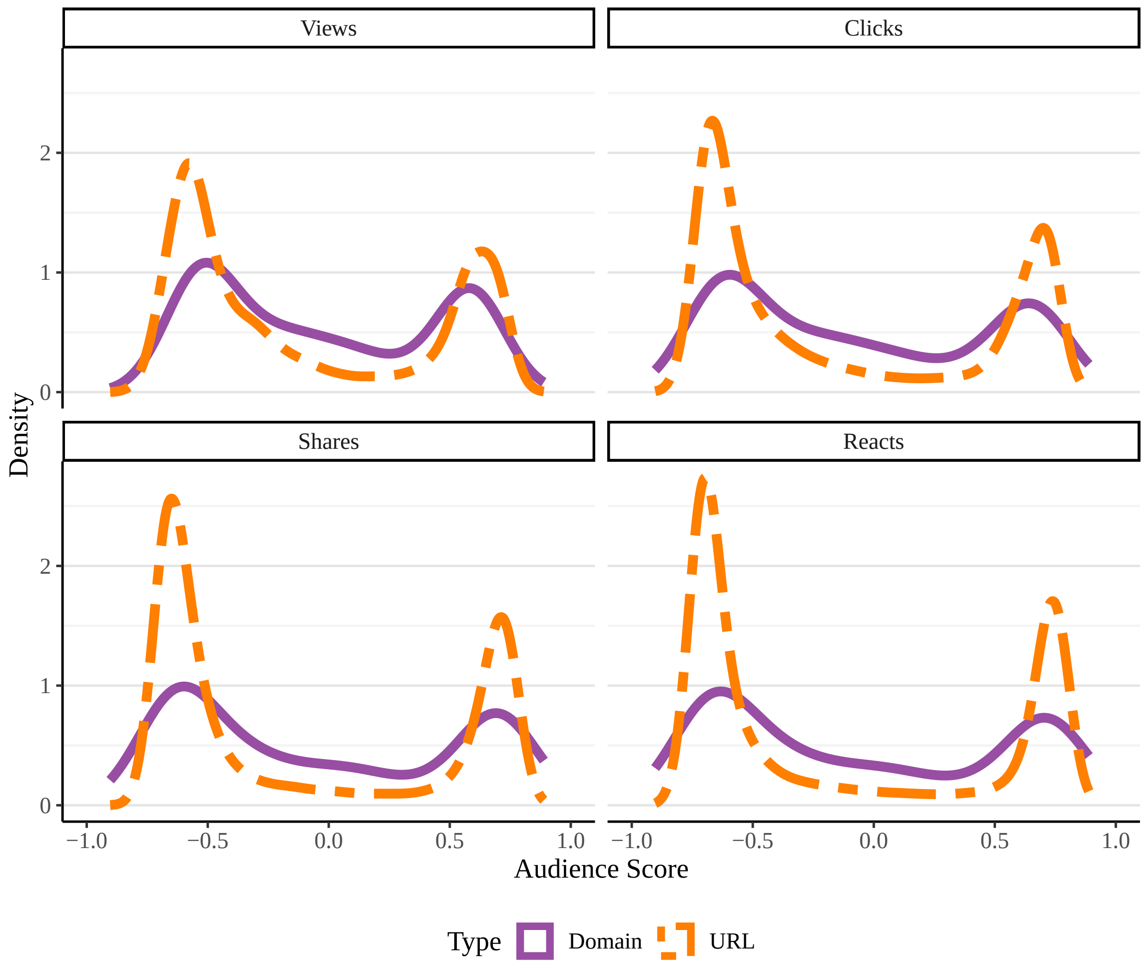

FIGURE I.6. Facebook distributions (high threshold)