

THE PINCHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

Frictions in News Consumption: Evidence from Social Media
Luca Braghieri, Ro'ee Levy, and Hannah Trachtman

Discussion Paper No. 9-2026

Article-Level Slant and Polarization of News Consumption on Social Media*

Luca Braghieri, Sarah Eichmeyer, Ro’ee Levy,
Markus Mobius, Jacob Steinhardt, and Ruiqi Zhong

July 2026

Abstract

We fine-tune a large language model on expert ratings to construct a content-based, article-level measure of news slant, applied to the near universe of articles (~ 1 million) published online by the top ~ 100 U.S. outlets in 2019. 70% of the variation in slant occurs within rather than across outlets, leaving wide scope for within-outlet ideological selection. Consistent with this, polarization in news consumption on Facebook—the gap in average slant between articles Republicans and Democrats consume—equals 89% of the distance between the average New York Times and Fox News article and is 53% larger than outlet-level measures imply.

*Braghieri: Bocconi University, CEPR, CESifo, IGIER, and NHH. Email: luca.braghieri@unibocconi.it. Eichmeyer: Bocconi University, CEPR, CESifo, and IGIER. Email: sarah.eichmeyer@unibocconi.it. Levy: Tel Aviv University and CEPR. Email: roeelevy@tauex.tau.ac.il. Mobius: Microsoft Research. Email: mobius@microsoft.com. Steinhardt: UC Berkeley, BAIR, and CLIMB. Email: jsteinhardt@berkeley.edu. Zhong: UC Berkeley. Email: ruiqi-zhong@berkeley.edu. We thank Pablo Barberá, Rafael Jiménez-Durán, Stefano DellaVigna, Matthew Gentzkow, Carlo Schwarz, Jesse Shapiro, Mateusz Stalinski, and seminar participants at the Association for Comparative Economic Study and Society for Institutional and Organization Economics ASSA Session, Bocconi University, Carnegie Mellon University, the EJPE-CEPR Political Economy Conference (Naples), the Hebrew University of Jerusalem, the London School of Economics, the National University of Singapore, Northwestern University, NOVA University in Lisbon, the NYU Center for Social Media and Politics, the Polarize Conference, the TAD annual conference, the University of Cologne, the University of Copenhagen, the University of Essex, the University of Vienna, and the University of Warwick for helpful comments. We thank Maor Ankri, Oliver Kiss, Heitor Sousa Borges, Wei Pin, Uri Schwartz, and Daniel Wolkowicz for excellent research assistance. Funding for this project was provided by the European Research Council (ERC Grant ECHO - 101162903), the Foerder Institute for Economic Research, The Pinhas Sapir Center, and the Tel Aviv AI and Data Science Center.

1. Introduction

Over the last two decades, social media has become one of the primary ways in which people consume information and access news. As of 2025, around 54% of U.S. adults reported regularly consuming news through social media, with Facebook being the dominant social media platform for news access (Newman et al., 2025).

The increased reliance on social media for news consumption has generated widespread concern. In particular, many worry that the high personalization of content through the political pages a user follows, the structure of the social network, and the algorithm governing users’ feeds lead to the formation of “echo chambers” and “filter bubbles” that promote the consumption of pro-attitudinal news (Pariser, 2011; Sunstein, 2018).¹ The resulting limited exposure to counter-attitudinal news might hinder the formation of accurate political beliefs, and, ultimately, people’s ability to constructively participate in the democratic process (Downs, 1957; Becker, 1958).

Although almost ubiquitous in the popular press, worries about pro-attitudinal news consumption online, and especially on social media, have received relatively little empirical substantiation from the academic literature (Gentzkow and Shapiro, 2011; Bakshy, Messing and Adamic, 2015; Nelson and Webster, 2017; Guess and Nyhan, 2018; Eady et al., 2019; Yang et al., 2020; Guess, 2021). This lack of evidence may reflect a mismatch between the level of aggregation at which political slant is commonly measured in the academic literature—namely, news outlets—and the level at which content is curated on social media platforms—namely, news articles. When measuring polarization in news consumption on social media, much of the existing literature proxies the slant of an article with that of the outlet it came from. If, however, the social media curation process exposes users to ideologically congenial content even *within* outlets, measuring slant at the outlet level will underestimate the polarization of news consumption.

This paper addresses this limitation: we develop a content-based method to assign slant to individual news articles rather than to entire outlets; we then use it to estimate polarization in news consumption on Facebook.

Our analysis relies on three main datasets. The first dataset encompasses the near universe of news articles (~ 1 million) published online by the top ~ 100 U.S. outlets in terms of online visits

¹Pro-attitudinal news, also known as congenial or like-minded news, refers to news whose slant matches a person’s ideology.

in 2019. By “news articles,” we refer to coverage of politics and public affairs, including opinion pieces. The second dataset is the “Facebook Privacy-Protected Full URLs Dataset,” produced by a consortium of academics (Social Science One) in partnership with Facebook’s parent company Meta. The dataset comprises aggregated, anonymized Facebook activity data at the URL level for all articles publicly shared on Facebook at least 100 times. The third dataset is a panel of browsing behavior for a representative sample of U.S. adults.

The cornerstone of our analysis is a content-based measure of slant at the article level obtained by fine-tuning a large language model (LLM) pre-trained on a vast corpus of text data. To generate the training and validation data, we hired two expert raters (as well as a research assistant) and asked them to assign each article a measure of slant on a scale from -3 (very left-wing) to 3 (very right-wing). We use 3,033 expert-labeled articles to fine-tune the model and reserve 726 articles for model validation. Fine-tuning the LLM substantially increases its performance: the Spearman correlation between the model’s predictions and the expert labels in the holdout portion of the labels dataset rises from 0.502 to 0.711.

An array of validation exercises shows that our machine-learning-based measure of article-level slant produces predictions that closely match expert labels across multiple validation metrics. For instance, when we aggregate our measure at the outlet level and compare it to one of the most well-established and comprehensive measures of outlet-level slant ([Bakshy, Messing and Adamic, 2015](#)), we find a Pearson correlation of 0.87. Similarly, when we compare our article-level slant measure to the slant manually assigned by media watchdog Ad Fontes Media to a small subset of our articles, we find a Pearson correlation of 0.83.

Our first set of results focuses on the production of slant. We find that outlet averages miss most article-level variation: a variance decomposition shows that $\sim 70\%$ of the variance in news article slant occurs within outlets rather than across outlets. The overall distribution of articles is centered around moderate values, but the tails are substantively important and vary systematically by article type: opinion articles and national news are much more slanted than non-opinion, local, or international news.

Our main results concern the circulation and consumption of slanted articles on social media. We introduce and estimate a new measure of polarization in news consumption—the polarization index. We define it as a scaled difference, ranging from -1 to $+1$, between the average slant of the

articles that Republicans and Democrats consume. For Facebook news consumption, the index is 0.16. The underlying slant gap—the difference in average slant—equals 0.82 standard deviations of the article-level slant distribution and 89% of the average slant distance between the New York Times and Fox News. It slightly exceeds the distance between the average Washington Times and Washington Post article.² It is comparable to the gap between the news shared on X by Democratic Senator Chuck Schumer and then-Senator Marco Rubio, a Republican.

An independent browsing panel, in which partisanship is observed from voter-registration records rather than inferred from Facebook behavior, yields a very similar estimate. The polarization index is slightly—by 3%—larger than our URL-level Facebook activity dataset-based estimate.

Article-level measurement matters quantitatively. If each article is assigned the average slant of its outlet—thus shutting down the channel of pro-attitudinal news consumption within outlets—the polarization index falls from 0.16 to 0.10. Thus, our polarization index is 53% higher when calculated using article-level rather than outlet-level slant.

The browsing panel lets us ask whether this within-outlet sorting is specific to social media: using its referrer information, we assign each article visit to the entry channel of the browsing session in which it occurs. We can therefore compare news consumed in sessions initiated through social media with news consumed in sessions initiated through direct navigation to an outlet’s homepage—the closest online analog to traditional news consumption. We find that sessions initiated through social media are narrowly more polarized than those initiated via direct homepage navigation, but the underlying mechanisms differ: under direct navigation, partisans sort almost entirely across outlets (the within-outlet slant gap is close to zero), whereas on social media a meaningful within-outlet gap remains.

We document three mechanisms that drive polarized news consumption on social media. First, extreme articles are four times more likely to circulate widely on Facebook than moderate ones. As a result, the distribution of articles that are widely shared on Facebook has much thicker tails than the distribution of articles produced. Second, most of the consumption gap is already present in exposure: the articles appearing in users’ feeds account for 83% of the click-based polarization index, highlighting the important role of the feed-curation process. Third, a back-of-the-envelope

²Across the outlets in our dataset, the Washington Post, the New York Times, the Washington Times, and Fox News rank at percentiles 11, 31, 82, and 90 of the distribution of average outlet slant, respectively.

calculation suggests that politically homophilous friend networks combined with pro-attitudinal sharing—the phenomenon of “echo chambers”—can account for a large share of polarized exposure, though not all of it: if users’ feeds simply reflected the articles shared by their friends, the degree of polarization in news exposure on Facebook would be at least 60% of what it currently is.

This project contributes to the literature studying the differential consumption of news by partisanship or ideology (Bakshy, Messing and Adamic, 2015; Flaxman, Goel and Rao, 2016; Gentzkow and Shapiro, 2011; González-Bailón et al., 2023; Green et al., 2025). To help readers navigate this crowded literature, we provide a schematic overview of the research articles most closely related to our paper in [Appendix Table A.1](#). To the best of our knowledge, this paper is the first to employ a content-based measure of article-level slant at scale to estimate the degree of polarization in news consumption on social media.

We view each component of the sentence above as conceptually important. First, our measure of *polarization* in news consumption is an essential complement to the measures of segregation and favorability common in the literature (Gentzkow and Shapiro, 2011; González-Bailón et al., 2023): polarization captures whether Democrats and Republicans consume articles of different *slant*, whereas segregation and favorability capture whether they consume *different articles*—a related but distinct question.³

Second, for studying polarization of news consumption in the digital age, the relevant object is a *content-based* measure of slant at the *article level*, yet most existing measures populate the other cells of this design space: audience-based measures at the article level (as in González-Bailón et al., 2023; Green et al., 2025), content-based measures at the outlet level (as in Groseclose and Milyo, 2005; Gentzkow and Shapiro, 2010), or audience-based measures at the outlet level (as in Flaxman, Goel and Rao, 2016; Nyhan et al., 2023).⁴ Content matters because audience-based patterns may have little to do with slant; the article level matters because social-media curation operates on articles, not outlets. The small literature that does assign content-based slant to individual articles (Boxell and Conway, 2022; Budak, Goel and Rao, 2016; Garz et al., 2020; Green et al., 2025; Peterson, Goel and Iyengar, 2021) either does not deploy it at scale or addresses questions (e.g.,

³We define segregation and favorability, contrast them with our polarization measure, and show that the two indeed diverge—e.g., identical wire-service articles published by differently slanted outlets receive the same slant score but different favorability scores—in [Appendix B](#).

⁴Audience-based measures rely on differential consumption or sharing patterns by Democrats and Republicans.

the role of journalists in the production of slant) orthogonal to ours.

Beyond the slant measure itself, these analyses deliver three substantive contributions to the literature on partisan news consumption. We document that most slant variation lies within rather than across outlets; we show that polarization in news consumption on social media is substantial—contrary to much of the prior literature⁵—and is sizably understated when slant is proxied at the outlet level; and we identify the social-media-specific mechanisms, such as pro-attitudinal sharing and network homophily, that drive it.

The remainder of the paper provides background on news consumption on social media ([Section 2](#)), describes our data ([Section 3](#)), the slant measure and its validation ([Section 4](#)), the production-side results ([Section 5](#)), and the analysis of social media news consumption ([Section 6](#)).

2. Background

As of 2025, around 54% of U.S. adults reported consuming news through social media—more than through any other channel and, for the first time, more than through television ([Newman et al., 2025](#)). Facebook is the dominant platform for news access, owing to its sheer size (~ 190 million monthly active U.S. users; [Datareportal, 2023](#)): 32% of U.S. adults report regularly getting news there, compared to 23% on Twitter (now “X”) and 12% on TikTok ([Newman et al., 2025](#)).

The shift toward social media is not merely a change in where people get news; it also changes how the news bundle a reader encounters is assembled. Readers could always select articles within newspapers, but on social media the outlet-curated edition is dissolved: articles reach a user’s feed individually, as a function of the pages she follows, the friends she is connected to, and the platform’s ranking algorithm. The unit at which news is curated and selected is thus the article rather than the outlet, and the slant of what a user encounters from a given outlet need not resemble that outlet’s average article. A Democrat may disproportionately encounter New York Times columns by Paul Krugman (a left-wing columnist), while a Republican may encounter New York Times columns by Bret Stephens (a right-wing columnist). Studying the slant of the news that partisans are exposed to and consume on social media therefore requires assigning slant to individual articles rather than to entire outlets.

⁵This finding is consistent with the analysis in [González-Bailón et al. \(2023\)](#) showing that Democrats and Republicans tend to consume different outlets and, even more markedly, different articles.

3. Data

Our analysis relies on three main datasets: one that includes the URL, text, and meta-data of a large collection of news articles, one that details the number of Facebook views, clicks, and shares for a subset of those articles, and one that tracks the browsing behavior of a sample of U.S. adults.

3.1 Near Universe of News Articles

The first dataset we employ in our analysis comprises the near universe of articles covering politics and public affairs (~ 1 million) published online in 2019 by the top ~ 100 U.S. outlets in terms of online visits.⁶ Focusing on the top 100 outlets balances broad coverage of U.S. online news consumption against data collection costs: according to Comscore data, visits to these outlets account for 80% of all online consumption from U.S. outlets focused on political news (Comscore, 2019).

We automatically retrieved the English-language articles published online in 2019 by the top 100 U.S. news outlets directly from the outlets’ websites.⁷ We decided not to include articles containing fewer than 250 characters since these are typically not news articles (e.g., videos, alerts, or pages linking to other articles).⁸ This procedure generated a dataset of 3,268,760 articles.

In order to identify and isolate the subset of articles covering politics and public affairs—henceforth referred to as “news articles”—we employed a machine learning algorithm. Following the literature, we define news articles as political and public-affairs content—commonly referred to in the literature as “front-section” or “hard news”—including both reported news and opinion pieces (e.g., op-eds), consistent with the literature on political news consumption (Flaxman, Goel and Rao, 2016; Gentzkow and Shapiro, 2010; Prior, 2007). Thus, our definition of news articles excludes soft-news content such as sports, entertainment, weather, and similar categories. The reasons for restricting our focus to news articles are twofold: first, slant is not as well-defined for entertainment or “soft” news such as celebrity gossip and sports; second, news articles are likely more relevant in shaping people’s political behaviors like voting. To classify news articles, we re-

⁶We exclude popular outlets from regions other than the U.S. (e.g., the BBC), news agencies (e.g. the Associated Press), as well as outlets that do not focus on political news (e.g., Elle).

⁷To ensure that we were only collecting news articles, we restricted data collection to English articles where the HTML ‘pageType’ meta field is defined as ‘article’.

⁸We further excluded the following outlets where our automatic retrieval appeared not to capture the full article text: azfamily.com, dailykos.com, dailywire.com, inforum.com, komonews.com, post-gazette.com, today.com, triblive.com, usnews.com, and WSJ.com.

lied on a machine-learning algorithm trained on 2,470 labels annotated by the research team. The binary classifier achieved 90% sensitivity and 97% specificity.

Our final sample includes 1,054,124 news articles. We present the list of outlets included in our dataset, as well as the number of news articles published in 2019 by each outlet in [Appendix Table C.1](#).

To obtain an estimate of the degree to which our dataset is comprehensive, we make use of browsing data from a representative sample of 5,032 adults who visited at least one news article (described in detail in [Section 3.3](#)). To do so, we match the URLs that appear in our dataset of news articles with the news article URLs visited by the panelists. We find that only 4.6% of the news-article URLs from the top-100 outlets visited by these panelists do not appear in our dataset, i.e., our dataset covers 95.4% of them.

3.2 URL-level Facebook Activity Dataset

The second dataset we employ is the “Facebook Privacy-Protected Full URLs Dataset,” which results from a partnership between Social Science One (SS1), a consortium of academics hosted by Harvard’s Institute for Quantitative Social Science, and Facebook’s parent company Meta. The dataset, henceforth referred to as the URL-level Facebook Activity Dataset, contains information about various forms of engagement with all the URLs shared publicly on Facebook more than 100 times starting on January 1, 2017.⁹

The dataset provides the number of unique users who viewed, clicked, and shared a given URL on Facebook. A user counts as having “viewed” a URL if a post containing it was rendered on the user’s screen (e.g., appeared in their News Feed); views are thus a measure of passive exposure rather than active engagement. The engagement counts (i.e., views, clicks, shares) are broken down by calendar year-month, gender, age group, and users’ political affinity (discussed in detail below).¹⁰

⁹For privacy protection, Meta added Laplace-5 noise to the share counts used to determine which URLs are shared publicly on Facebook more than 100 times ([Messing et al., 2023](#)).

¹⁰We only consider engagement that occurred in the calendar year 2019. In principle, some articles published in 2019 could experience a degree of engagement in subsequent years. In practice, almost all engagement—more precisely 91%, 92% and 89% of views, clicks, and shares, respectively—occurs immediately, in the month in which an article first appears in the URL-level Facebook Activity Dataset. This means that using 2019 engagement data should provide relatively precise engagement measures since almost no engagement occurs in the calendar year following the year when the article first appears in the dataset.

For privacy protection, Meta added Gaussian noise to each engagement data point (i.e., counts of clicks, views, and shares) according to the dictates of epsilon-delta differential privacy ([Messing et al., 2023](#)). In all our analyses based on the Facebook engagement data, we aggregate engagement counts across thousands of news articles within our article slant bins, thereby averaging out the noise present in the engagement measures at the level of individual articles.

Of our news articles, 22% appear in the URL-level Facebook Activity Dataset (thus having been shared at least 100 times); of those, the average article received more than 200,000 views, 9,000 clicks, and 1,800 shares (see [Appendix Table A.2](#) for details).

Political Affinity Measure The URL-level Facebook Activity Dataset contains a political affinity measure that classifies Facebook users into five ideology buckets, along with a sixth bucket for users who are not assigned a political affinity score. For most of our analysis, we consolidate the measure into three buckets: liberals, moderates, and conservatives. The political affinity measure is derived from users’ connections (likes) to political pages on Facebook, employing a methodology similar to the one in [Barberá et al. \(2015\)](#).¹¹

The political affinity measure has several desirable features. First, [Bond and Messing \(2015\)](#) validate a measure based on the same model against existing measures of ideology for political actors and against the self-reported political views of users, and find Pearson correlations as high as 0.94. Second, the measure is defined on the URL-level Facebook Activity Dataset, which covers the universe of URLs shared on Facebook at least 100 times. Third, the measure can be applied to users who do not actively report their ideology or register with a party.

The political affinity measure also has three drawbacks, each of which we address directly.

First, it is not a canonical measure of partisanship: it does not track indicators such as party registration, and because we do not have access to the proprietary data underlying Meta’s measure, we cannot construct a confusion matrix linking affinity scores to registration status. Our solution

¹¹This process begins by constructing an adjacency matrix that represents monthly active U.S. users (aged 18 and over) as rows and a curated list of political pages as columns. The matrix tracks whether users “like” these pages on Facebook. Correspondence analysis is then applied to a subset of this matrix, focusing on users who “like” at least ten political pages, to estimate the political positions of pages and users on a common ideological scale. Subsequently, political page-affinity scores for all users who like at least one of the selected political pages are calculated by averaging the scores of the pages they “like.” These scores are then converted to percentiles to categorize users into quintiles based on their political affinity. See the codebook of the URL-level Facebook Activity Dataset for details ([Messing et al., 2023](#)).

exploits the fact that the measure is constructed almost identically to the Twitter-based (now “X”) measure of [Barberá et al. \(2015\)](#), whose underlying data are public and include party registration for a subset of users. We recover the affinity-to-registration confusion matrix from the Twitter data and, under the assumption that it carries over to the Facebook population, use it to estimate the degree of polarization in news consumption that would prevail if partisanship in the URL-level Facebook Activity Dataset were defined by party registration rather than by political affinity (see [Section 6.1](#)).

Second, the measure does not have universal coverage: it is assigned only to users who like at least one political page, roughly 25% of users ([Messing et al., 2023](#)). These users nonetheless account for 72% to 87% of engagement with news articles on Facebook (depending on the engagement measure; see [Appendix Table A.3](#)), so analyses that rely directly on the affinity score capture the vast majority of engagement. Moreover, our headline polarization estimate—being based on party registration—incorporates the news consumption of users with missing affinity scores, as described in [Section 6.1.2](#).

Third, the measure might suffer from a degree of circularity: among the more than 500 political pages used in its construction, 17 news outlets overlap with our list.¹² In [Appendix D.1](#), we show that dropping these outlets from the analysis leaves our results virtually unchanged, as does replacing the affinity measure with an alternative ideology proxy based on basic demographics.

Finally, and most importantly, we validate our results against a dataset that is immune to all three concerns: the panel of browsing behavior described below, in which partisanship is measured directly from party-registration records. Restricting attention to news visits with Facebook as the referrer yields a polarization estimate that is entirely independent of the URL-level Facebook Activity Dataset and, reassuringly, almost identical to it ([Appendix D.1](#)).¹³

¹²The other political pages are members of Congress, hosts of cable news shows, political parties, candidates in presidential primary elections, and other news outlets.

¹³The advantage of the browsing-panel estimate is that it relies directly on party-registration records; the disadvantage is that it is estimated on a relatively small sample that is not necessarily representative of the Facebook population.

3.3 Panel of Browsing Behavior

For additional robustness checks and analyses, we rely on browsing panel data from a representative sample of 35,444 U.S. adults, curated by market-research firm Luth Research.¹⁴ This dataset provides comprehensive URL-level browsing activity data on both desktop and mobile devices for all panel members during the last quarter of 2019. Among these panelists, 5,032 visited at least one news article published by the 100 outlets included in our study, generating a total of 88,167 clicks.

A unique feature of this dataset is that it includes party registration status for registered voters from almost all states where official party registration data is publicly available. Luth procures this information from TargetSmart’s voter registration database, which itself is sourced from information released by state election officials. Of the 5,032 panelists who visited at least one news article, 1,613 have a Democratic or Republican party-registration record; these panelists constitute the sample for our browsing-panel analyses.¹⁵

The two datasets play complementary roles. The Facebook URL-level data provide platform-wide engagement measures for the near-universe of news content, allowing us to study polarization and within-outlet selection at scale; the browsing panel is smaller but offers an independent benchmark with partisanship observed from voter-registration records, and its referrer information allows us to compare social media to other channels of online news access (e.g., direct homepage navigation).

4. Measurement of Slant at the Article Level

The cornerstone of our analysis is a content-based measure of slant for each article in our database. We construct this measure by fine-tuning a large language model on a training dataset with more than 3,000 articles whose slant was rated by experts. We discuss the procedure in detail below.

¹⁴The sample is recruited to be representative, on the first moments of gender, age, and household income, of the adult U.S. population.

¹⁵Party-registration records are available for panelists residing in the 30 states (plus the District of Columbia) that register voters by party—56.1% of our panelists live in these states, closely matching the 57.0% population share—as well as for some panelists in other states who were previously registered in a party-registration state. Coverage is partial in California owing to legal restrictions on the data vendor. Panelists without a Democratic or Republican registration record include unregistered individuals, voters registered without a party affiliation, and residents of states without party registration.

4.1 “Ground Truth” Expert Slant Labels

The credibility of our slant measure rests on the quality of the expert labels used to fine-tune the model. We therefore invested heavily in the labeling process: over a period of about six months, we employed two freelance workers with graduate degrees in political science and criminal justice, as well as a research assistant who worked as an editor for a Berkeley student-run newspaper and who was trained for the task. Henceforth, we refer to the two freelance workers as “expert raters.”¹⁶

Our two experts independently assigned slant labels to 3,759 news articles randomly selected from our database, with the randomization stratified by week of the year. For each article, raters received only the cleaned text and title; we did not provide any additional information, such as the URL of the article or the outlet from which it came. Raters were asked to assess the slant of each article on a 7-point scale, where 1 corresponded to a very left-wing article and 7 corresponded to a very right-wing article. For ease of interpretation, in the rest of the paper we recenter the scale so that negative numbers indicate left-leaning articles, positive numbers indicate right-leaning articles, and zero indicates a moderate article.¹⁷

Raters were instructed to consider three main criteria when assigning slant: language, political position, and issue coverage. Language refers to the usage of partisan terminology as in [Gentzkow and Shapiro \(2010\)](#). Political position refers to whether the article takes a stance that is closer to the Democratic or Republican position on the issue. Issue coverage refers to whether the article disproportionately emphasizes issues salient to either party’s base.¹⁸ Raters were asked to consider all three criteria and provide a holistic assessment for each article. The full instructions are reported in [Appendix E.1](#).

We took several steps to ensure high-quality ratings. First, we ordered articles by publication date, so that, when rating articles from the same period, experts could better account for contemporaneous context. Second, we elicited confidence levels and excluded 2% of articles for which

¹⁶The experts were recruited through Upwork, after careful review of their credentials, a Zoom interview, and the successful completion of a sample rating task.

¹⁷We do not interpret a “moderate” article as unbiased reporting or as a normative gold standard of journalistic objectivity. The existence and attainability of such a standard are contested in the journalism literature ([Hackett, 1984](#); [Schudson, 2001](#)). In our usage, an article is classified as moderate when expert raters place it near the center of the left–right scale.

¹⁸See, for instance, the example of selective coverage in [Chopra, Haaland and Roth \(2024\)](#) in which the Congressional Budget Office evaluates a prospective reform along two dimensions and news articles selectively report only one of the two dimensions. Other examples include the coverage of climate change ([Djourelouva et al., 2024](#)) and political scandals ([Puglisi and Snyder, 2011](#); [Galvis, Snyder Jr and Song, 2016](#)).

at least one rater reported no confidence. Third, in the rare cases where the two experts disagreed about the direction of an article’s slant (occurring 6% of the time), we referred the article to a research assistant (RA) hired specifically for the task of providing a third independent assessment. The RA reviewed the article, replaced the rating of the expert with whom he disagreed with his own independent rating, and documented the rationale for his assessment.

The final slant label used in the analysis is constructed as the average of the two expert ratings after RA correction. The distributions of the two experts’ raw ratings and of the resulting article-level labels are centered around moderate values with thick tails, as shown in [Appendix Figure A.1](#).

The Pearson (Spearman) correlation between the original two labels is 0.56 (0.49) and increases to 0.72 (0.68) after including the RA’s correction. The investment in expert raters paid off: previous work relying on crowd-sourced (mTurk) ratings found a Pearson correlation between raters’ labels of only 0.23 ([Peterson, Goel and Iyengar, 2021](#)).

Among pairs of holdout articles assigned distinct values by both raters, the two raters assign the same relative slant ordering in 78% of cases (and in 91% of cases after the RA’s correction). Quadratic-weighted Cohen’s Kappa, a commonly used measure of inter-rater reliability, is 0.68, which is typically interpreted as “substantial” agreement (the second best category out of 5).¹⁹ Disagreement is quantitatively small: the median and modal absolute difference between the two labels is 0, and the average absolute difference between them is 0.60 on our 7-point scale.

4.2 Machine Learning

Article cleaning We pre-processed the article text to ensure that our model assessed article slant based solely on the article text and title. We first used existing packages (Domdistiller, Readability, Trafilatura) designed to extract relevant text from HTML files. Before predicting slant, we further removed irrelevant text from the article (e.g., copyrights, links, captions, advertisements, and other boilerplate) using a large language model. Finally, we masked the article’s outlet name, so the model would not explicitly use the outlet to predict slant.²⁰ [Appendix E.2](#) provides more details on

¹⁹Cohen’s Kappa measures agreement between raters who each classify items into mutually exclusive categories while accounting for chance agreement, providing a more accurate assessment of inter-rater reliability than simple percent agreement.

²⁰For example, we replaced “Wyoming Sen. John Barrasso told CNN” with “Wyoming Sen. John Barrasso told [Outlet name].”

the text processing.

Slant prediction Following recent advances in natural language processing and machine learning, we use a mixture of general-purpose and task-specific machine learning to assign a measure of slant to each news article in our universe. The general purpose component consists of the generative pre-trained transformer open-weights model Gemma 3 by Google DeepMind. The task-specific component involves fine-tuning the model, using our expert labels as training data, for the specific task of assigning a measure of slant to a given news article based on its text. For the fine-tuning component, we provided the model with the same set of instructions as our expert raters, as well as with 3,033 articles from our training set containing the average of our two experts’ labels. A further 726 labeled articles were held out for validation. We only provided the model with each article’s headline and text; we did not provide any additional information, such as the URL of the article. [Appendix E.3](#) provides more details on fine-tuning the model and the prediction of slant.

Fine-tuning the model yielded sizeable improvements in the model’s ability to predict the average of our two experts’ labels. Among the 726 articles in the original holdout for which both models return parseable predictions, the Spearman correlation between the model’s predictions and the average of our two experts’ ratings is 0.502 without fine-tuning and 0.711 with fine-tuning.

Model output The fine-tuned model produces, for each article in our database, a slant prediction ranging from -3 (a very left-wing article) to 3 (a very right-wing article). Before constructing any analysis or validation statistic, we merge slant bucket -3 with bucket -2.5 and bucket 3 with bucket 2.5 , because the two most extreme buckets contain too few articles to draw meaningful conclusions, especially given the noise introduced into the URL-level Facebook Activity Dataset by differential privacy. Thus, every reported validation metric evaluates the same deployed measure used in the substantive analysis.

4.3 Validation Exercises

Because our article-level slant measure is novel, we validate it through an array of exercises showing that our methodology to estimate slant at the article level produced precise and high-quality estimates.

First, for the articles in the hold-out portion of our training dataset, we correlate the prediction from our machine learning algorithm with our expert slant label (given by the average rating of our two experts). We find a Pearson correlation of 0.76; we find a Spearman correlation of 0.71.²¹ Since our “ground truth” measure of slant is itself noisy (the correlation between the two experts’ labels is 0.72), the maximum achievable correlation between any model prediction and the expert labels is bounded away from one. Intuitively, even a perfect model of the latent “true” slant cannot fully match the expert ratings because those ratings themselves contain a degree of measurement error.

To quantify this ceiling, note that the two expert ratings can be viewed as independent noisy measures of the same latent slant, so their correlation pins down the share of signal in each rating—and hence the residual noise in their average. As shown in [Appendix F](#), the resulting upper bound is 0.91, of which our model attains 84%.

Second, we compare the absolute difference between the model prediction and the expert slant labels. In 83% of cases, the model prediction and the label fall within 0.5 points of each other. The mean absolute difference (MAE) is 0.41, which corresponds to only 8% of the length of the deployed scale. We also run this comparison separately for each expert. The absolute differences between our model and each expert label are smaller than the absolute difference between the expert labels; this is true for both the original and the RA-amended expert labels. Considering all pairs of rated articles in the holdout test set, the two measures assign the same relative slant ordering in 87% of cases. Lastly, among articles with a non-neutral expert label, the model avoids assigning slant in the opposite direction in 91% of cases; predictions of zero are included among the cases that avoid an opposite-direction assignment.

Third, we aggregate our measure at the outlet level and compare it to one of the most comprehensive existing measures of outlet-level slant, namely the measure by [Bakshy, Messing and Adamic \(2015\)](#). [Figure 1](#) plots the outlet-level slant rating against the outlet-level measures we obtain by averaging the slant of all the articles produced by the same outlet. We find a Pearson correlation of 0.87.

Fourth, we correlate our model predictions of slant with high-quality human slant ratings from

²¹The R^2 from a regression of a hold-out article’s model-predicted slant on its average expert label is 0.58 and the RMSE is 0.69.

an independent source: media watchdog Ad Fontes Media (henceforth “AFM”). AFM uses a panel of analysts from across the political spectrum to rate a representative sample of news articles every week. For the year 2019, we found a total of 512 news articles with publicly available expert slant ratings from AFM. Lending credence to our model predictions, we find a high correlation of 0.83 between the two slant measures. To illustrate the fit across the distribution, we plot the mean AFM slant rating across articles in each model-predicted slant bin in [Appendix Figure A.2](#). We find a smooth linear relationship between the two ratings. Given that the AFM raters are chosen to evenly represent political perspectives (each article gets rated by a panel of three individuals, one from the political left, one from the political center, and one from the political right), the strong alignment between our model predictions and the AFM labels should also alleviate concerns about our two expert raters having the same political orientation.²²

Taken together, these validation exercises give us confidence in the article-level slant measure used throughout the analysis. The high level of agreement among our raters, together with the measure’s close correspondence with the independent Ad Fontes ratings, provides an empirical basis for interpreting differences on the scale as differences in experts’ perceptions of political slant.

5. News Production

In order to provide useful context and to benchmark our polarization index for news consumption on social media, our first set of results analyzes the production of slant by the top 100 U.S. outlets in 2019. We begin this analysis by introducing a variance decomposition aimed at obtaining a quantitative measure of the importance of taking into account article-level slant instead of relying on slant measured at the outlet level.

5.1 Variance Decomposition: The Role of Outlets

One of the key arguments in this paper is that measuring slant at the outlet level might underestimate polarization in news consumption on social media. This argument is predicated on the intuition that the social media environment can induce partisans to sort ideologically within outlets and not just

²²Both our raters self-identify as weakly liberal.

across outlets. Of course, a necessary condition for meaningful partisan sorting within outlets is that there be variation in slant within outlets in the first place.

We first illustrate the role of within- vs. across-outlet variation in slant with a case study of two of the most influential U.S. news outlets: the New York Times and Fox News. [Figure 2](#) shows the distribution of article-level slant separately for each outlet. The two outlets visibly differ in their average slant, with the slant distribution of articles from the New York Times skewing visibly more left-wing than that of Fox News. The difference in average slant between the two outlets is 0.88 points on our scale. Despite the average difference in slant, however, there is large overlap in the two distributions: within each outlet, at least half of the articles are neutral or cross-cutting (i.e., exhibiting a slant of the opposite sign than the average slant of the outlet). Thus, the two outlets contain both liberal and conservative articles. Indeed, we find that there is sizeable within-outlet variation in slant for the vast majority of outlets.²³

To obtain a quantitative measure of the extent to which, in general, the slant distributions of articles from different outlets overlap, we turn to a variance decomposition. We note that, due to measurement error in our slant measure, a “naive” variance decomposition—obtained by first regressing our model-predicted article-level slant on outlet fixed effects, and then retrieving the R^2 as a measure of the fraction of the overall variation in slant that is accounted for by fixed differences across outlets—would produce biased estimates. Measurement error adds noise to article-level slant that is unrelated to outlet identity; this noise inflates variance without contributing to across-outlet differences, so a naive R^2 understates the importance of outlets.

To correct for this bias, we first note that our measure of article-level slant is affected by two separate sources of measurement error: first, the labels of our two experts exhibit a degree of noise in that they do not always agree. Second, our machine learning algorithm produces a noisy slant prediction. We deal with the latter source of noise by restricting the variance decomposition to our dataset of labeled articles and using the expert ratings directly, thereby avoiding introducing an additional layer of prediction error from the model.²⁴ We deal with rater noise by leveraging

²³See [Appendix Figure C.1](#) for the distribution of slant for each outlet in our sample.

²⁴For this exercise, we asked the experts to rate 800 additional randomly selected articles from outlets with fewer than 30 rated articles, with the goal of obtaining at least 30 rated articles per outlet. These additional ratings are included in the variance decomposition but not in model training or holdout validation, because they are selected according to a different criterion than our baseline set of article labels. Including these additional articles in the holdout validation set does not meaningfully alter the statistics reported in [Section 4.3](#).

the fact that we observe two independent ratings for each article. Intuitively, each rating can be viewed as the sum of a latent “true” article slant and an idiosyncratic measurement error. We use the disagreement between the two ratings to estimate the variance of this measurement error and subtract it from the total variance, thereby recovering the variance of the underlying slant measure. This allows us to compute the share of variance attributable to differences across outlets net of measurement error.²⁵ Conceptually, this approach is analogous to standard methods that estimate the variance of an underlying latent variable by considering the covariance of two independent but noisy draws of that variable.

Our variance decomposition shows that outlet differences account for only 30% of the variance in article-level slant. In other words, ignoring within-outlet variation in slant induces one to miss 70% of the overall variation in article-level slant—that is, all the variation that occurs *within outlets*.

More broadly, the decomposition above can shed new light on the existing literature and provide useful context for policy debates. First, it clarifies how our article-level evidence relates to classic “audience segregation” measures constructed at the outlet level ([Gentzkow and Shapiro, 2011](#)). When a meaningful share of slant variation occurs within outlets, low outlet-level segregation can coexist with sharp differences in what different groups actually read; conversely, high outlet-level segregation need not translate into comparably large differences in slant exposure when outlets regularly publish centrist or cross-cutting items. Second, the decomposition aligns naturally with an industrial organization view of news outlets as multi-product firms that choose a portfolio of items rather than a single editorial position. Recent models of media influence and capture allow equilibrium slant to vary widely within sources, with a nontrivial mass of centrist and cross-cutting content ([Alonso and Padró i Miquel, 2025](#)). Our estimate that outlet provenance accounts for only about one-third of the variance provides empirical evidence consistent with such models. Third, the decomposition sharpens the discussion around policies encouraging source diversity: if most variation is within outlets, increasing the number of sources consumers access need not yield a commensurate increase in the diversity of article-level content consumed.

²⁵The formal derivation is provided in [Appendix G](#).

5.2 Overall Distribution of Slant

[Figure 3](#) presents the overall slant distribution of the ~ 1 million news articles published online by the 100 most-visited U.S. news outlets in 2019. The distribution is unimodal and concentrated around the center: a slant of zero is by far the most common rating, accounting for 37% of all articles.²⁶ Among the remaining articles, 75% are left-leaning. While strongly slanted articles are a small share of production, the tails are thick: 5% of all articles are extremely left-leaning (slant of -2 or below) and 6% extremely right-leaning (slant of 2 or above)—which, given the ~ 1 million articles in our data, amounts to over 100,000 extreme articles published in a single year.

5.3 Distribution of Slant by News Category and Topic

Next, we analyze how slant varies by article type. We use GPT-4o-mini to classify articles along three dimensions: opinion versus non-opinion, local versus national versus international news, and 26 topic categories (e.g., immigration, environment, crime). The first two classifications agree with our expert labels in 95% and 92% of cases, respectively; [Appendix H](#) provides details on all three.

Two patterns stand out. First, opinion pieces—which make up 22% of all articles—are far more slanted than non-opinion articles: 43% of opinion articles have an absolute slant of 2 or more, compared to 2% of non-opinion articles. Panel (a) of [Figure 4](#) shows the corresponding distributions: non-opinion articles are concentrated around the center, while opinion pieces are strongly bimodal. The small share of moderate opinion pieces is consistent with arguments on the decline of moderates in US politics, both among political elites and engaged citizens ([Abramowitz, 2010](#)). Moreover, our finding that opinion pieces are particularly slanted, together with results from the literature showing that news consumers often trust opinion pieces as sources of fact ([Bursztyn et al., 2023](#)), suggests that opinion pieces might be an important contributor to political polarization among news readers.

Second, national news is much more slanted than local or international news (panel (b) of [Figure 4](#)): 17% of national-news articles have an absolute slant of 2 or more, compared to roughly 3% of local and international articles. Restricted to national news, the media landscape no longer

²⁶The independent Ad Fontes ratings described in [Section 4.3](#) corroborate the interpretation of a model-predicted slant of zero as centrist: among articles our model rates 0 that also have Ad Fontes ratings, the average Ad Fontes rating is -0.04 . The two measures also broadly agree on the overall fractions of left-leaning, moderate, and right-leaning articles.

appears primarily moderate: a majority of articles have an absolute slant of at least 1. Topic composition explains part of the gap: international news disproportionately covers foreign conflicts (37% of international articles), in which the U.S. was not directly involved in 2019, while local news leans toward bipartisan issues such as drug abuse—and both conflicts and drugs are among the topics with the lowest average absolute slant ([Appendix Figure A.3](#)).

The topic classification further reveals differences in what left- and right-leaning articles cover. The starkest pattern provides further validation for the slant measure: Democratic scandals are the most over-represented topic among right-leaning articles and Republican scandals the most under-represented ([Appendix Figure A.4](#)), in line with outlet-level evidence ([Puglisi and Snyder, 2011](#)).²⁷ Beyond scandals, left-leaning articles over-represent topics such as the environment, policing, and gender, while right-leaning articles over-represent the economy and social policy.²⁸ A complementary lasso analysis of article titles points to the same divide: the words most predictive of a right-leaning article include “aoc,” “illegal,” and “socialism,” while “climate,” “women,” and “lgbtq” are among the most predictive of a left-leaning one ([Appendix Figure A.7](#)). These patterns raise the question of whether polarization in consumption simply reflects partisans reading about different topics; we return to this in [Section 6.1.4](#), where we show that polarization arises overwhelmingly *within* topics.

6. News Consumption on Facebook

In the previous section, we provided a broad overview of the slant landscape of the news articles in our database. In this section, we first estimate the degree of polarization in news consumption on social media and then study the drivers of such polarization. The distribution of slant on the production side provides context for our polarization result and helps us describe the mechanisms driving it.

²⁷We classify an article as left-leaning if its slant is below -0.5 and right-leaning if above 0.5 . A topic is over- (under-) represented among left-leaning articles if the share of left-leaning articles within the topic exceeds (falls short of) the corresponding share in the overall sample.

²⁸Slant also varies across topics in level and dispersion: [Appendix Figure A.5](#) shows the full distribution of slant by topic, and [Appendix Figure A.6](#) shows the within-topic standard deviation as a measure of how ideologically divided coverage is. Cultural topics (gender, LGBTQ issues, race) as well as welfare, gun rights, immigration, and the environment are the most dispersed, while disasters, international conflicts, drugs, crime, housing, trade, and policing (our data pre-date the death of George Floyd in May 2020) are the least dispersed.

6.1 Polarization in News Consumption on Social Media

In order to study the degree to which news consumption on social media is polarized, we first introduce a novel measure of polarization in news consumption and then estimate it on the Facebook URL-level Activity dataset.

6.1.1 Polarization Measure: Definition

We summarize the polarization of news consumption with a single statistic we call the *polarization index*, $\mathcal{P}$, defined as the difference in click-weighted average slant between Republican and Democratic users, normalized to range between -1 and 1 :

$$\mathcal{P}(\overline{S}_R, \overline{S}_D) = \frac{\overline{S}_R - \overline{S}_D}{5}. \quad (1)$$

Here, $\overline{S}_j$ represents the click-weighted average slant for users registered with party $j \in \{D, R\}$, where D stands for Democratic and R for Republican. We divide the numerator by 5 because our effective slant scale ranges from -2.5 to $+2.5$.²⁹ Throughout the paper, we refer to the numerator of this expression—the difference in the average slant of the news consumed by Republicans and Democrats, expressed in units of the article-slant scale—as the *slant gap*. The polarization index is thus the slant gap divided by five, which normalizes it to the interval $[-1, 1]$.

Our measure of polarization has three desirable properties that make it an important complement to other measures developed in the literature, such as segregation and favorability (Gentzkow and Shapiro, 2011; González-Bailón et al., 2023).

First, the polarization index takes into account both the ideology of a user and the slant of the news articles that the user consumes. This allows us to study whether Democratic users tend to consume left-leaning articles and Republican users right-leaning articles. This distinguishes polarization from segregation, which instead captures whether the two groups consume *different articles*, regardless of those articles’ slant. To see the distinction, consider a world in which (i) people consume only local political news, (ii) individuals are geographically segregated by ideology, and (iii) every local-politics article is exactly moderate in slant. Democrats and Republicans then read en-

²⁹Recall that, in order to have a sufficient number of articles in each bucket to perform our analysis, we merge slant bucket -3 with slant bucket -2.5 and slant bucket 3 with slant bucket 2.5 . As a result, our effective scale ranges from -2.5 to 2.5 .

tirely different articles, so segregation is at its maximum; yet because every article is moderate, the average slant each group consumes is identical and our polarization measure is zero.

Second, the polarization index is straightforward to interpret: a value of 1 indicates extreme pro-attitudinal (like-minded) news consumption, a value of 0 means both groups consume news with, on average, an identical slant, and a value of -1 indicates extreme counter-attitudinal news consumption.

Third, the polarization index allows for an easily interpretable comparison between article-level and outlet-level polarization. To see this, note that when outlets contain both left- and right-leaning articles, the gap between article- and outlet-level polarization is determined by within-outlet selection. If Democrats disproportionately select relatively more left-leaning articles within each outlet and Republicans relatively more right-leaning ones, article-level polarization exceeds outlet-level polarization; the reverse holds under counter-attitudinal selection, and the two measures coincide if both groups assign equal weights to articles within outlets.³⁰

6.1.2 Polarization Index: Estimation

The goal of this section is to estimate our headline polarization index using a canonical indicator of partisanship, namely party registration. The Facebook URL-level Activity Dataset does not contain a party registration variable; it contains the less canonical political affinity variable described in [Section 3.2](#). To estimate polarization by party registration, we therefore construct a mapping from political affinity—liberal (L), moderate (M), and conservative (C)—to three registration categories: registered Democrat (RD), registered Republican (RR), and neither (NR).

Let

$$w_{g,p} = \Pr(p \mid g)$$

denote the probability that a user in political-affinity group $g \in \{L, M, C, \text{missing}\}$ belongs to registration category $p \in \{RD, RR, NR\}$.

For users with an observed political affinity score, we estimate $w_{g,p}$ using the Twitter (now “X”) data assembled by [Barberá et al. \(2015\)](#), which contain both political-affinity scores and

³⁰We formalize this relationship in [Appendix B](#), which shows that article-level polarization is strictly larger (smaller) than outlet-level polarization if and only if individuals engage in pro-attitudinal (counter-attitudinal) news consumption within outlets, and that the two coincide under random within-outlet selection.

party-registration information for a subset of users. We collapse the original affinity scores into the liberal, moderate, and conservative categories and calculate the registration shares within each category. We then assume that this Twitter-based affinity-to-registration mapping approximates the corresponding mapping among Facebook users. This cross-platform assumption is motivated by the fact that the underlying methodology employed to assign political affinity scores is analogous across the Twitter and Facebook datasets.³¹

For Facebook users without an affinity score, we cannot estimate a corresponding mapping from the Twitter data. We therefore assume that their registration distribution equals the population distribution reported in the 2020 American National Election Studies ([American National Election Studies, 2020](#)):

$$(w_{\text{missing, RD}}, w_{\text{missing, RR}}, w_{\text{missing, NR}}) = (0.23, 0.16, 0.61).$$

The resulting mapping from the Facebook political affinity variable to the three party registration categories is reported in [Appendix Table A.4](#).

We next use this mapping to allocate observed Facebook clicks across registration categories. Let $K_{a,g}$ denote the total number of clicks on article a attributed to affinity group g . We impute the number of those clicks attributable to registration category p as

$$\hat{K}_{a,p} = \sum_g w_{g,p} K_{a,g}.$$

Thus, within each affinity group, we allocate an article's clicks across registration categories according to the corresponding registration shares.

Let S_a denote the estimated slant of article a . The click-weighted average slant for registration category p is then

$$\bar{S}_p = \frac{\sum_a \hat{K}_{a,p} S_a}{\sum_a \hat{K}_{a,p}}.$$

6.1.3 Headline Polarization Results

Applying the procedure described in the previous section, we estimate that, in 2019, the polarization index, i.e. our measure of the degree of polarization in news consumption on Facebook, was 0.16 (see [Table 1](#)).

³¹Indeed, Facebook's methodology is based on [Barberá et al. \(2015\)](#).

As discussed, our headline polarization index uses a canonical notion of partisanship: party registration. If, instead, we defined partisanship using Facebook’s “political affinity” labels in the URL-level Facebook Activity Dataset, the polarization index would be higher: 0.27 (see [Appendix Table A.3](#)). The gap between the two estimates is consistent with Facebook’s “liberal” and “conservative” labels disproportionately selecting strong partisans.

The most direct validation of our headline estimate comes from the independent browsing panel described in [Section 3.3](#), in which partisanship is observed directly from voter-registration records rather than inferred from on-platform behavior, and news consumption is recorded at the individual level rather than obtained from platform-provided aggregates. When we harmonize the two datasets—restricting both to articles published in the final three months of 2019 (the period covered by the panel) and to articles appearing in the URL-level Facebook Activity Dataset—the browsing-panel polarization index differs from the Facebook-based index by only 3%—it is, in fact, slightly larger. Replicating our result this closely, in a dataset that shares neither its partisanship measure nor its data provider, suggests that our result is not an artifact of Facebook’s political-affinity measure, of our affinity-to-registration mapping, or of the platform-provided engagement data. [Appendix D.1](#) provides details.

6.1.4 Benchmarks

Is a polarization index of 0.16 large or small? To gauge its magnitude, we benchmark it against several reference points.

Article- vs. Outlet-level Polarization A natural benchmark is to recompute the polarization index after assigning to each article a measure of slant that equals the average slant of the outlet that the article came from. This exercise mimics the results one would obtain if, in line with much of the literature, one were to carry out the analysis of polarization in news consumption on Facebook at the outlet level rather than at the article level.

When we perform the exercise above, we obtain a polarization index of 0.10.³² Thus, the polar-

³²As in our preferred measure of polarization from the previous section, this measure employs, as a definition of partisanship, being a registered Democrat or a registered Republican. If we consider a polarization index based directly on the political affinity score, we obtain a similar result in terms of the fraction by which polarization in news consumption is underestimated when the channel of pro-attitudinal news consumption within outlets is shut down.

ization index estimate for Facebook news consumption increases by 53% (using unrounded values) when the channel of pro-attitudinal news consumption within outlets is taken into account. We also find a substantial gap between outlet-level polarization and article-level polarization when focusing on the articles individuals viewed or shared ([Appendix Figure A.8](#)).

Within-Topic Polarization In principle, polarization in news consumption could arise because Democrats and Republicans read about different topics as the distribution of slant varies across topics ([Appendix Figure A.5](#)). Three pieces of evidence indicate that topic selection is not the main driver. First, the partisan composition of readership is stable across topics: with the exception of scandals, the share of clicks accounted for by registered Democrats lies between 21% and 26% for every topic. Second, there is substantial polarization *within* topics: for nearly every topic Republicans consume considerably more conservative articles than Democrats ([Appendix Figure A.9](#)), and the polarization index exceeds 0.1 for almost every topic. Third, mirroring the outlet-level exercise, we recompute the polarization index after assigning each article the average slant of all articles on its topic; the resulting index of 0.01 is more than an order of magnitude smaller than the article-level index. Democrats and Republicans thus consume differently slanted articles about the *same* topics.

Additional Benchmarks An additional benchmark compares the mean difference in average slant between Republican and Democratic users (i.e., the numerator in [Equation 1](#)) to the standard deviation in the overall database of all news articles published by the top 100 outlets. The former is 0.78. The latter is 0.95. Therefore, the mean difference in slant for articles clicked on by Republican and Democratic users is 0.82 standard deviation units.

We can also interpret the difference from the previous section in standard deviation units of the article-slant distribution. The Republican–Democrat slant gap measured from article-level slant is 0.28 standard deviations larger than the outlet-level Republican–Democrat slant gap.

Another set of benchmarks compares the gap to the distance between ideologically contrasting outlets. The gap equals 89% of the distance in average slant between the New York Times and Fox News—the case study of [Figure 2](#)—and slightly exceeds the distance between the average Washington Times and Washington Post article.³³

³³Across the outlets in our dataset, the Washington Post, the New York Times, the Washington Times, and Fox

Two further comparisons help fix ideas. The gap in news consumption between Democrats and Republicans on Facebook is comparable to the slant difference between the average article by liberal New York Times columnist Paul Krugman and by his moderate/conservative colleague David Brooks, and slightly smaller than the average distance between the articles shared on Twitter (now “X”) by Democratic Senator Chuck Schumer and then-Senator Marco Rubio, a Republican.³⁴

Finally, [Appendix I](#) illustrates the magnitude with three articles covering the same event—Donald Trump discussing the border wall on Easter—with slants of -1 (liberal), 0 (moderate), and 1 (conservative). The distances between the liberal and the moderate article and between the conservative and the moderate article are comparable to our polarization measure, and the differences are palpable: the liberal article calls the administration’s efforts botched, the moderate article quotes Trump without positive or negative framing, and the conservative article discusses the administration’s accomplishments.

6.2 Mechanisms

What drives the high degree of polarization in news consumption on social media? We examine four mechanisms in turn: whether partisans select ideologically congenial articles within outlets; whether polarization stems from the articles users are shown on their feeds or from what they select among those articles; whether the articles circulating widely on Facebook are more slanted than the articles outlets produce; and whether echo chambers—the interaction of homophilous friend networks with pro-attitudinal sharing—can account for polarized exposure.

6.2.1 Pro-attitudinal news consumption within outlets

Partisan Facebook users consume pro-attitudinal news even within outlets. [Figure 5](#) shows this pattern for the 20 most visited outlets: within nearly every outlet, the average article consumed by users with liberal political affinity is virtually always more left-leaning than the average article consumed by users with conservative political affinity.³⁵ The within-outlet gaps can be large: at

News rank at percentiles 11, 31, 82, and 90 of the distribution of average outlet slant, respectively. Both pairs thus span most of the ideological range of mainstream U.S. outlets.

³⁴We match the news-article URLs shared by each senator’s Twitter account during 2019 to our corpus and compute the difference in the average predicted slant of the matched articles, which equals 0.85 slant points ([Litell, 2019](#)).

³⁵The same pattern holds across the full outlet set ([Appendix Table C.1](#)).

the New York Times, the average article clicked by conservative-affinity users is 0.22 slant points to the right of the average article clicked by liberal-affinity users (on our five-point scale), and at the New York Post the corresponding gap reaches 0.69 points—about 51% of the overall liberal–conservative consumption gap observed in the data.

6.2.2 Selection vs. Exposure

We next decompose polarization in news consumption into exposure and selection conditional on exposure. The URL-level Facebook Activity Dataset reports both the number of users who viewed a URL on Facebook and the number who clicked it. We use views as a proxy for exposure—what appears on users’ feeds—and clicks as our measure of consumption—which articles users click on and read among the articles that appear on their feeds.

The exposure-based polarization index is 0.13, which accounts for 83% of the click-based polarization index. Thus, most of the Republican–Democrat slant gap in consumption is already present in the articles that appear in users’ feeds. Put differently, if users clicked randomly among the articles they were shown, the polarization index would be only 17% smaller than it is in the data.^{36,37} But note that this thought experiment is partial equilibrium in nature: we take the set of articles shown to users as given and decompose polarization into the contribution of exposure versus selection conditional on that exposure. A full general equilibrium analysis would additionally account for feedback effects between exposure, consumption, and sharing behavior—for example, how users’ engagement shapes the content they are subsequently shown, or how sharing decisions influence others’ exposure.

6.2.3 Production vs. Distribution

The distribution of news circulating on Facebook differs sharply from the distribution of news produced by outlets. [Figure 6](#) shows a V-shaped relationship between article slant and the probability of appearing in the URL-level Facebook Activity Dataset, i.e., of being publicly shared at least 100 times. Centrist articles have an 11% probability of being publicly shared at least 100 times,

³⁶As in [Section 6.1](#), partisanship here is defined by party registration. Defining it instead by the political affinity score yields a similar exposure/selection split.

³⁷Consistent with our findings for news consumption, we also find strong pro-attitudinal news exposure within outlet, as depicted in [Appendix Figure A.10](#).

compared to more than 40% for articles with absolute slant of 2 or more. Extreme articles are therefore about 4 times more likely to circulate widely on Facebook. This V-shaped pattern reflects the combined effect of outlet posting decisions, user sharing, platform curation, and other factors shaping which articles circulate widely. As a result, the slant distribution of widely shared articles has thicker tails than the distribution of articles produced by outlets ([Appendix Figure A.11](#)) and the variance of the former is 64% higher than the variance of the latter.

Next, we document that article types that correlate with more extreme slant—opinion articles, national news, contentious topics—are much more likely to be shared on Facebook at least 100 times. Compared to the production side, opinion pieces and national news are 31% and 32% more common among articles widely shared on Facebook, respectively ([Appendix Figure A.12](#)). Contentious topics (i.e., topics with a relatively high standard deviation in slant on the production side) tend to also be over-represented among widely shared articles on Facebook ([Appendix Figure A.13](#)).

Finally, sharing itself is highly pro-attitudinal. [Figure 7](#) shows that among users with political affinity scores, liberals share left-leaning articles around three times more often than moderate ones and almost never share right-leaning articles; conservatives display the mirror image.³⁸ This pattern holds whether or not users clicked on the article before sharing it ([Appendix Figure A.16](#)), suggesting that headlines alone convey an article’s slant. Opinion-article sharing is especially polarized: the sharing-based polarization index for opinion articles is 2.9 times that for non-opinion articles ([Appendix Table A.5](#)). These sharing patterns are a key input into the echo-chamber bound we construct next.

6.2.4 The Role of Echo Chambers

Echo chambers can generate polarized exposure when politically homophilous friendship networks interact with pro-attitudinal sharing. We ask how much polarization in exposure would arise in a “friends-only” counterfactual in which users’ feeds reflected only the articles shared by their Facebook friends, shutting down personalized ranking and other forms of feed customization.³⁹ For

³⁸Residualizing on outlet fixed effects moves the two distributions closer together—outlets explain part of the differential sharing—but the liberal distribution remains markedly to the left of the conservative one ([Appendix Figure A.14](#), [Appendix Figure A.15](#)).

³⁹We note once again that this analysis is in “partial equilibrium” in the sense described in [Section 6.2.2](#).

this exercise, we use the political-affinity version of the polarization index, because the network-homophily estimates we use are reported by political affinity.⁴⁰ Let $x_{\{i \leftarrow j\}}$ denote the share of type- i users' friends who are of type j , where $i \in \{lib, con\}$ and $j \in \{lib, mod, con\}$. We take these network-composition terms from Table S2 of [Bakshy, Messing and Adamic \(2015\)](#). Let $\bar{y}_j$ denote the average slant of articles shared by users of type j , estimated from the sharing patterns documented in [Section 6.2.3](#). By the law of total expectation, the average slant a type- i user is exposed to under this counterfactual is $\sum_j \bar{y}_{j,i} \cdot x_{\{i \leftarrow j\}}$, where $\bar{y}_{j,i}$ is the average slant of the type- j -shared articles that reach type- i users. We do not observe $\bar{y}_{j,i}$, so we replace it with the unconditional average $\bar{y}_j$, which yields

$$\mathcal{P}_{exposure}^{echo} = \frac{1}{5} \left[\sum_j \bar{y}_j \cdot x_{\{con \leftarrow j\}} - \sum_j \bar{y}_j \cdot x_{\{lib \leftarrow j\}} \right], \quad j \in \{lib, mod, con\},$$

The first sum is a proxy for the average slant to which conservative users would be exposed if their feeds reflected only the articles shared by their friends; the second is the analogous object for liberal users. Their difference, normalized by five, is directly comparable to our baseline exposure-based polarization index measured with political affinity scores. This estimate provides a lower bound under the conservative assumption that, within each poster type, users with disproportionately liberal friend networks share weakly more left-leaning content than their type's average, and users with disproportionately conservative networks share weakly more right-leaning content—so that the content reaching liberal users is weakly more left-leaning, and the content reaching conservative users weakly more right-leaning, than $\bar{y}_j$.

The resulting lower bound equals 60% of the actual exposure-based polarization index measured using political affinity scores.⁴¹ Echo chambers alone can thus generate a majority of the observed polarization in news exposure on Facebook.⁴²

⁴⁰Using the polarization index based on party registration would have required imposing two sets of assumptions: one for translating polarization based on political affinity scores into polarization based on party registration as described in [Section 6.1](#), and another for estimating the role of echo chambers.

⁴¹The magnitude of this bound is not mechanical: it could in principle have been close to zero (had moderates' shares dominated friend-generated exposure) or exceeded one (had feed ranking moderated network-driven polarization).

⁴²This exercise is a back-of-the-envelope benchmark, not a causal decomposition of feed polarization.

6.3 Comparison of Social Media to Other Sources of Online News Consumption

The results in [Section 6.1](#) show that news consumption on Facebook is substantially polarized and that outlet-level slant measures miss a substantial part of the Republican–Democrat difference in average slant. We next ask whether this article-level component is distinctive to social media or also characterizes other ways of accessing online news.

We use the Luth browsing panel to compare news consumed in sessions initiated through social media with news consumed in sessions initiated through other channels, including direct navigation from an outlet’s homepage. Direct homepage navigation is our key benchmark because it is the closest online analog to traditional news consumption: readers enter through an outlet’s front page and then choose which of its articles to read.

To ensure that this comparison captures readers’ subsequent within-outlet article choices, rather than only the landing article, we assign access channels at the browsing-session level while retaining the news-article visit as the unit of observation. The panel records article-level browsing histories for registered Democrats and Republicans and contains referrer information that allows us to identify how each session begins. We treat an article reached from an external source or from the outlet’s homepage as the beginning of a session, classify the session according to that entry route, and assign the same channel to subsequent articles reached through within-outlet navigation. Thus, an article reached from social media and later articles reached by clicking within that outlet belong to the same social-media session; an article selected from the outlet homepage and subsequent within-outlet clicks belong to the same direct-homepage session. A session ends when the user leaves the outlet’s website and visits a different one. We group the session entry routes into social media, direct homepage navigation, search engines, news aggregators, and an “other” category, and we weight visits so that each individual receives equal total weight across observed news visits.

[Table 2](#) reports Republican–Democrat slant gaps by session entry channel—that is, the difference in the average slant of articles consumed between the average registered Republican and the average registered Democrat—on the article-slant scale that runs from -2.5 to $+2.5$. Column 1 shows that social media has the largest total slant gap, at 0.52 slant points. Direct homepage navigation is a close second, at 0.50, and the two estimates are not statistically distinguishable at

conventional levels. Search engines have a much smaller gap of 0.10, and news aggregators a gap of 0.17 (the coefficient for news aggregators is not statistically significant).⁴³

Column 2, which presents results after absorbing outlet fixed effects, paints a different picture. While social media and direct homepage navigation have similar total slant gaps, their sources differ. The direct-homepage gap falls from 0.50 to 0.01 and is statistically indistinguishable from zero, and the search-engine gap also falls close to zero. In contrast, the social-media gap remains positive at 0.12 slant points and statistically significant, albeit only at the 10% level. Thus, taking the point estimates at face value, the high total gap for direct homepage navigation is almost entirely accounted for by the outlet margin: in sessions initiated through outlet homepages, Republicans and Democrats reach outlets with different average slant but, conditional on outlet, read articles of similar average slant. Social media is distinctive because partisan slant differences remain within outlets. The difference between social media and direct homepage navigation is consistent with the mechanisms of [Section 6.2](#): features such as homophilous network structures and pro-attitudinal sharing behavior are relevant for news consumed through social media, but play no role in news consumed through direct homepage navigation.

These results sharpen the interpretation of the Facebook findings. Social media does not have a substantially larger total Republican–Democrat slant gap than every other way of accessing online news: direct homepage navigation is a close second. The source of the gap, however, differs: under direct navigation, partisans sort almost entirely across outlets; on social media, they also sort across articles within outlets. This within-outlet margin is precisely the component that outlet-level slant measures miss, and it is the reason why measuring slant at the article-level is important when studying polarization of news consumption on social media.

6.4 Robustness Checks

The estimates above rely on the URL-level Facebook Activity Dataset, which raises three natural concerns: that our conclusions might depend on Facebook’s political-affinity measure, on the focus on the top 100 outlets, or on the dataset’s inclusion threshold of 100 public shares. None of these

⁴³The magnitudes in [Table 2](#) are not directly comparable to the headline Facebook estimate, because the browsing-panel estimates include all social-media referrals, weight individuals equally, cover only the final three months of 2019, and include articles regardless of whether they appear in the URL-level Facebook Activity Dataset. [Appendix D.1](#) harmonizes these margins across the two datasets; the resulting polarization indices differ by only 3%.

drives our results; we summarize the evidence here and provide details in [Appendix D](#).

First, the results are robust to how partisanship is measured. Recomputing the polarization index in the browsing panel—where partisanship comes directly from party-registration records, restricting to Facebook-referred visits—yields an estimate just 3% larger than the baseline. Excluding the 17 outlets that overlap with the political pages used to construct Facebook’s political-affinity measure leaves both the polarization index and the article-vs.-outlet gap very similar. And when we proxy ideology with demographics instead—using the Republican lean of age-by-gender groups estimated from representative survey data—the average slant of the news a group is exposed to, consumes, and shares on Facebook increases steadily with the group’s Republican lean, mirroring our main pattern.

Second, the polarization index is not an artifact of the top-100-outlet sample: ranking outlets by online visits and extrapolating the index to progressively larger outlet sets implies that expanding to the top 200 outlets would change the estimate by only 6%.

Third, the inclusion threshold is unlikely to drive the results. Recomputing the browsing-panel index on all Facebook-referred news articles—including those shared fewer than 100 times—yields an estimate only 6% smaller than the baseline, and an extrapolation exercise within the URL-level data reaches the same conclusion. The threshold leads the main analysis to modestly overstate polarization, consistent with the pro-attitudinal sharing documented in [Section 6.2.3](#): articles clear the threshold partly through partisan amplification. However, the results do not change by much since the vast majority of clicks are concentrated among articles shared more than 100 times.

7. Conclusion

This paper develops a content-based measure of political slant at the article level—built by fine-tuning an open-weights large language model on expert ratings and validated against independent benchmarks—and applies it to the near universe of news articles published online by the top 100 U.S. outlets in 2019.

Three findings emerge. First, 70% of the variation in article-level slant occurs within rather than across outlets, so outlet-level measures miss most of the variation in the slant of the news people may consume. Second, news consumption on Facebook is substantially polarized: the

polarization index is 0.16, equivalent to 0.82 standard deviations of the slant distribution, to 89% of the New York Times–Fox News distance, and to slightly more than the distance between the average Washington Times and Washington Post article. Because partisans select ideologically congenial articles within outlets, this figure is 53% higher than an outlet-level analysis would deliver. Third, this within-outlet margin is distinctive to social media: under direct homepage navigation—the closest online analog to traditional news consumption—the partisan slant gap arises almost entirely across outlets, whereas on social media a significant gap remains within them, consistent with the pro-attitudinal sharing and homophilous networks we document.

Our analysis has limitations. It covers a single year, 2019, and one platform. Other platforms may weight friend networks, followed pages, and algorithmic recommendation differently, and any platform’s workings evolve over time; extending the polarization measure across time and platforms is a natural next step, and we release our fine-tuned model weights to facilitate such work.

For policy, our results show that various traditional remedies may be insufficient to address polarization on social media. Policy on both sides of the Atlantic has long pursued diversity at the level of the *source*—“diverse and antagonistic” outlets in the U.S. tradition ([Associated Press v. United States, 1945](#)), “external pluralism” in the European one ([Valcke, Sükösd and Picard, 2015](#)). On social media, however, a significant share of polarization arises within outlets: Democrats and Republicans reading the same outlet consume systematically different articles. An expansion of the set of sources does not necessarily address this margin. In contrast, interventions at the level of the feed—where the sorting occurs—can. As news consumption migrates to social media, the share of polarization beyond the reach of source-level policy will grow with it.

References

- Abramowitz, Alan.** 2010. *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy*. Yale University Press.
- Alonso, Ricardo, and Gerard Padró i Miquel.** 2025. “Competitive Capture of Public Opinion.” *Econometrica*, 93(4): 1265–1297.
- American National Election Studies.** 2020. “ANES 2020 Time Series Study Full Release.”
- Associated Press v. United States.** 1945. 326 U.S. 1.
- Bakshy, Eytan, Solomon Messing, and Lada Adamic.** 2015. “Exposure to Ideologically Diverse News and Opinion on Facebook.” *Science*, 348(6239): 1130–1132.
- Barberá, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau.** 2015. “Tweeting from left to right: Is online political communication more than an echo chamber?” *Psychological Science*, 26(10): 1531–1542.
- Becker, Gary S.** 1958. “Competition and democracy.” *The Journal of Law and Economics*, 1: 105–109.
- Bond, Robert, and Solomon Messing.** 2015. “Quantifying social media’s political space: Estimating ideology from publicly revealed preferences on Facebook.” *American Political Science Review*, 109(1): 62–78.
- Boxell, Levi, and Jacob Conway.** 2022. “Journalist Ideology and the Production of News: Evidence from Movers.” *Working Paper*.
- Budak, Ceren, Sharad Goel, and Justin M. Rao.** 2016. “Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis.” *Public Opinion Quarterly*, 80(S1): 250–271.
- Buntain, Cody, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker.** 2023. “Measuring the Ideology of Audiences for Web Links and Domains Using Differentially Private Engagement Data.” *Proceedings of the International AAAI Conference on Web and Social Media*, 17(1): 72–83.
- Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott.** 2023. “Opinions as Facts.” *Review of Economic Studies*, 90(4): 1832–1864.
- Chopra, Felix, Ingar Haaland, and Christopher Roth.** 2024. “The Demand for News: Accuracy Concerns Versus Belief Confirmation Motives.” *The Economic Journal*, 134(661): 1806–1834.
- Comscore.** 2019. “Comscore Web Behavior Database.” *Data set, accessed via Wharton Research Data Services (WRDS)*, Licensed for academic research use.
- Datareportal.** 2023. “Facebook Users, Stats, Data & Trends.” <https://datareportal.com/essential-facebook-stats>.
- Djourelouva, Milena, Ruben Durante, Elliot Motte, and Eleonora Patacchini.** 2024. “Experience, Narratives, and Climate Change Beliefs.” *Available at SSRN 4680430*.
- Downs, Anthony.** 1957. “An economic theory of political action in a democracy.” *Journal of Po-*

litical Economy, 65(2): 135–150.

- Eady, Gregory, Jonathan Nagler, Andy Guess, Jan Zilinsky, and Joshua A Tucker.** 2019. “How many people live in political bubbles on social media? Evidence from linked survey and Twitter data.” *Sage Open*, 9(1): 2158244019832705.
- Flaxman, Seth, Sharad Goel, and Justin M. Rao.** 2016. “Filter Bubbles, Echo Chambers, and Online News Consumption.” *Public Opinion Quarterly*, 80(S1): 298–320.
- Galvis, Ángela Fonseca, James M Snyder Jr, and BK Song.** 2016. “Newspaper market structure and behavior: Partisan coverage of political scandals in the United States from 1870 to 1910.” *The Journal of Politics*, 78(2): 368–381.
- Garz, Marcel, Gaurav Sood, Daniel F Stone, and Justin Wallace.** 2020. “The supply of media slant across outlets and demand for slant within outlets: Evidence from US presidential campaign news.” *European Journal of Political Economy*, 63: 101877.
- Gentzkow, Matthew, and Jesse M. Shapiro.** 2010. “What Drives Media Slant? Evidence From U.S. Daily Newspapers.” *Econometrica*, 78(1): 35–71.
- Gentzkow, Matthew, and Jesse Shapiro.** 2011. “Ideological Segregation Online and Offline.” *Quarterly Journal of Economics*, 126(4): 1799–1839.
- Gentzkow, Matthew, Jesse M. Shapiro, and Matt Taddy.** 2019. “Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech.” *Econometrica*, 87(4): 1307–1340.
- González-Bailón, Sandra, David Lazer, Pablo Barberá, Meiqing Zhang, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Deen Freelon, Matthew Gentzkow, Andrew M. Guess, Shanto Iyengar, Young Mie Kim, Neil Malhotra, Devra Moehler, Brendan Nyhan, Jennifer Pan, Carlos Velasco Rivera, Jaime Settle, Emily Thorson, Rebekah Tromble, Arjun Wilkins, Magdalena Wojcieszak, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker.** 2023. “Asymmetric ideological segregation in exposure to political news on Facebook.” *Science*, 381(6656): 392–398.
- Green, Jon, Stefan McCabe, Sarah Shugars, Hanyu Chwe, Luke Horgan, Shuyang Cao, and David Lazer.** 2025. “Curation bubbles.” *American Political Science Review*, 119(4): 1704–1722.
- Groseclose, Tim, and Jeffrey Milyo.** 2005. “A Measure of Media Bias.” *The Quarterly Journal of Economics*, 120(4): 1191–1237.
- Guess, Andrew, and Brendan Nyhan.** 2018. “Avoiding the Echo Chamber About Echo-Chambers: Why Selective Exposure to Like-Minded Political News Is Less Prevalent Than You Think.” *Knight Foundation*.
- Guess, Andrew M.** 2021. “(Almost) everything in moderation: new evidence on Americans’ online media diets.” *American Journal of Political Science*, 65(4): 1007–1022.
- Hackett, Robert A.** 1984. “Decline of a Paradigm? Bias and Objectivity in News Media Studies.” *Critical Studies in Mass Communication*, 1(3): 229–259.
- Levy, Ro’ee.** 2021. “Social Media, News Consumption, and Polarization: Evidence from a Field

- Experiment.” *American Economic Review*, 111(3): 831–70.
- Litell, Alex.** 2019. “Tweets of Congress.” *GitHub dataset*, Daily Twitter output of members of the U.S. Congress; accessed 2019.
- Messing, Solomon, Christina DeGregorio, Bennett Hillenbrand, Gary King, Saurav Mahanti, Zagreb Mukerjee, Chaya Nayak, Nate Persily, Bogdan State, and Arjun Wilkins.** 2023. “Facebook Privacy-Protected Full URLs Data Set.” <https://doi.org/10.7910/DVN/TDOAPG>, Harvard Dataverse, V10.
- Nelson, Jacob L, and James G Webster.** 2017. “The myth of partisan selective exposure: A portrait of the online political news audience.” *Social Media + Society*, 3(3).
- Newman, Nic, Amy Ross Arguedas, Craig T Robertson, Rasmus Kleis Nielsen, and Richard Fletcher.** 2025. “Reuters Institute Digital News Report 2025.”
- Nyhan, Brendan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y. Chen, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Andrew M. Guess, Edward Kennedy, Young Mie Kim, David Lazer, Neil Malhotra, Devra Moehler, Jennifer Pan, Daniel Robert Thomas, Rebekah Tromble, Carlos Velasco Rivera, Arjun Wilkins, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker.** 2023. “Like-minded sources on Facebook are prevalent but not polarizing.” *Nature*, 620: 137–144.
- Pariser, Eli.** 2011. *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. Penguin.
- 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.
- Prior, Markus.** 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. Cambridge:Cambridge University Press.
- Puglisi, Riccardo, and James M Snyder.** 2011. “Newspaper coverage of political scandals.” *The Journal of Politics*, 73(3): 931–950.
- Schudson, Michael.** 2001. “The Objectivity Norm in American Journalism.” *Journalism*, 2(2): 149–170.
- Sunstein, Cass R.** 2018. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Valcke, Peggy, Miklós Sükösd, and Robert G. Picard,** ed. 2015. *Media Pluralism and Diversity: Concepts, Risks and Global Trends*. London:Palgrave Macmillan.
- Wong, David W. S.** 2003. “Spatial Decomposition of Segregation Indices: A Framework Toward Measuring Segregation at Multiple Levels.” *Geographical Analysis*, 35(3): 179–194.
- Yang, Tian, Sílvia Majó-Vázquez, Rasmus K Nielsen, and Sandra González-Bailón.** 2020. “Exposure to news grows less fragmented with an increase in mobile access.” *Proceedings of the*

National Academy of Sciences, 117(46): 28678–28683.

Yao, Jing, David W. S. Wong, Nick Bailey, and Jonathan Minton. 2019. “Spatial segregation measures: a methodological review.” *Tijdschrift voor Economische en Sociale Geografie*, 110(3): 235–250.

Tables and Figures

Table 1: Polarization in News Exposure, Consumption, and Circulation on Facebook

	Average Slant Reg. Democrats	Average Slant Reg. Republicans	Average Slant Difference	Polarization Index
Views	-0.61	0.04	0.64	0.13
Clicks	-0.61	0.16	0.78	0.16
Shares	-0.69	0.46	1.15	0.23


Notes: The first two columns of the table report the views/clicks/shares-weighted average slant of the news diets of registered Democrat and registered Republican Facebook users imputed according to the methodology described in [Section 6.1.2](#). The third column shows the difference between the first two columns. The last column shows the polarization index: it normalizes the difference from the third column to be in the interval $[-1, 1]$, with -1 indicating extreme counter-attitudinal news consumption, 0 indicating no difference in average slant across the news consumed by Democrats and Republicans, and 1 indicating extreme pro-attitudinal news consumption. The news article sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook).

Table 2: Article-Level Republican–Democrat Slant Gaps by Online News Access Channel in the Browsing Panel

	(1) Rep-Dem Slant Gap	(2) Rep-Dem Slant Gap
Social media	0.515*** (0.095)	0.125* (0.068)
Direct homepage navigation	0.496*** (0.081)	0.011 (0.071)
Search engine	0.100** (0.050)	0.035 (0.041)
News aggregator	0.175 (0.113)	0.090 (0.100)
Observations	34,170	34,170
Individuals	1,613	1,613
Outlet FE	No	Yes


Notes: The dependent variable is the predicted slant of a news article consumed by a browsing-panel member, measured on a scale from -2.5 to $+2.5$. An observation is a news-article visit. The sample is restricted to registered Democrats and registered Republicans. The table reports regressions of article slant on session entry-channel indicators and a full set of interactions between those indicators and a registered-Republican indicator. Because the specifications are saturated in entry-channel-by-party cells, each reported coefficient is the Republican–Democrat gap in article slant for the corresponding access channel. Column 1 reports the overall Republican–Democrat slant gap for each channel. Column 2 adds outlet fixed effects, so the coefficients capture average within-outlet Republican–Democrat slant gaps by channel. Visits are weighted so that each individual receives equal total weight across observed news visits. Access channels are assigned at the browsing-session level: subsequent articles reached through within-outlet navigation inherit the channel through which the session began. The channels are news aggregators, search engines, direct homepage navigation, social media, and other/missing; the other/missing category is included but not reported. The browsing panel is described in [Section 3.3](#). Standard errors, two-way clustered by individual and by article (URL), are reported in parentheses.

Figure 1: Comparison of Our Outlet-Level Measure with the one from Bakshy, Messing and Adamic (2015)




Notes: This figure presents our outlet-level slant measure on the horizontal axis, calculated as the average article-level slant across all articles published by an outlet (with article slant determined by the fine-tuned Gemma 3 27B’s predictions, as described in [Section 4.2](#)). The vertical axis represents an alternative outlet-level slant measure for the same set of outlets. Specifically, it uses the measure developed by [Bakshy, Messing and Adamic \(2015\)](#), which defines an outlet’s slant in terms of ‘alignment’—the extent to which an outlet’s articles are predominantly shared by left- or right-leaning users on Facebook. The figure highlights a few well-known outlets.

Figure 2: Distribution of Slant: New York Times and Fox News



Notes: This figure shows the distribution of slant among all news articles ($N = 82,051$), published online by the New York Times and Fox News in 2019. Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).


Figure 3: Overall Distribution of Slant Across Articles




Notes: This figure shows the distribution of slant among all news articles ($N = 1,054,124$) published online by the top 100 U.S. news outlets in 2019. Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure 4: Distribution of Slant Across Article Types

(a) Opinion vs. Non-Opinion Articles




(b) National, Local, and International News




Notes: Panel (a) shows the distribution of slant among all opinion news articles ($N = 230,203$) published online in the top-100 most-read U.S. news outlets in 2019 (right panel), and their non-opinion counterparts (left panel). Panel (b) shows the distribution of slant among all news articles covering national news (left panel), local news (middle panel), and international news (right panel). Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The classification into opinion articles and into national/local/international news is described in [Appendix H](#).

Figure 5: Average Slant of Articles Clicked on Facebook by Outlet and User Political Affinity




Notes: This figure shows the clicks-weighted average slant of articles by Facebook users' political affinity and outlet. The sample includes all news articles published online by the top 20 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., articles shared at least 100 times on Facebook). Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure 6: Likelihood of Being Shared on Facebook > 100 Times by Article Slant



Notes: The gray bars in this figure show the distribution of slant among all news articles ($N = 1,054,124$) published online by the top 100 U.S. news outlets in 2019 (left y-axis). The blue line shows the fraction of articles in each bin that were shared at least 100 times on Facebook (right y-axis), as measured by whether they appear in the URL-level Facebook Activity Dataset. The sample includes all news articles ($N = 1,054,124$) published online by the top 100 U.S. news outlets in 2019. Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure 7: Empirical PDF of Facebook Shares by User Political Affinity and Article Slant



Notes: This figure shows the fraction of all shares by a given Facebook user political affinity group (liberal and conservative, respectively) that accrue to articles in a given slant bin. The dashed vertical line represents the mean of the distribution. The sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., articles shared at least 100 times on Facebook). Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Appendix For Online Publication

Table of Contents

A	Additional Tables and Figures	A.2
B	Polarization, Segregation, and Favorability	A.24
B.1	Empirics	A.24
B.2	Theory	A.27
C	Slant distribution for each outlet in the top 100	A.35
D	Robustness	A.41
D.1	Robustness: Political Affinity Measure	A.41
D.2	Extending Results to Additional Outlets	A.42
D.3	Robustness: Threshold in the URL-Level Facebook Activity Dataset	A.43
E	Slant Ratings	A.46
E.1	Expert Label Construction	A.46
E.2	Text Processing	A.47
E.3	Predicting Slant	A.51
F	Correlation Ceiling Between Model Predictions and Expert Labels	A.51
G	Variance Decomposition	A.53
H	Classifying Other News Characteristics	A.54
I	Slant Distance: Example Articles	A.57

A. Additional Tables and Figures

Table A.1: Papers Estimating the Extent of Pro-attitudinal News Engagement

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Paper	Primary Estimand	Slant Measure	Primary Level of Analysis	Browsing Data	Social Media Data	Type of Social Media Content	Number of Articles or Outlets
Bakshy, Messing and Adamic (2015)	Ratio of clicks on counter-attitudinal vs. pro-attitudinal content	Audience-based	Outlet		U.S. Facebook users who self-report their ideology	Clicks, views, shares	500 news outlets
Flaxman, Goel and Rao (2016)	Standard deviation of average slant consumed across individuals	Audience-based	Outlet	Bing Toolbar sample	Bing Toolbar sample (only clicks referred from Facebook and Twitter)	Clicks	100 outlets
Gentzkow and Shapiro (2011)	Segregation	Audience-based	Outlet	Comscore sample		Clicks	119 outlets
González-Bailón et al. (2023)	Segregation	Audience-based	Article (social media) and outlet (browsing and social media)	U.S. 2020 Facebook and Instagram Election Study sample	Near universe of US Facebook users	Clicks, views, shares	~640,000 URLs
Green et al. (2025)	Correlation between predicted ideology and slant	Content-based	Article		Subset of Twitter's 10% Decahose sample	Shares	1,000 URLs
Guess (2021)	Degree of overlap in the distributions of domains visited	Audience-based	Outlet	YouGov Pulse sample		Clicks	~500 news outlets
Levy (2021)	Segregation and standard deviation of average slant consumed across individuals	Audience-based	Outlet	Extension sample & Comscore sample	Extension sample	Clicks, views, shares	~500 news outlets
Nelson and Webster (2017)	Audience overlap (network analysis)	Audience-based	Outlet	Comscore sample		Clicks	72 outlets
Peterson, Goel and Iyengar (2021)	Segregation and standard deviation of average slant consumed across individuals	Audience-based	Outlet	Wakoopa Toolbar sample		Clicks	355 domains
Braghieri et al. (2026) [this paper]	Polarization (difference in mean slant consumed by Republicans and Democrats)	Content-based	Article	Luth sample	Universe of US Facebook users	Clicks, views, shares	~1 million articles of which ~230,000 appear in the social media data

Notes: Column 1 lists the paper. Column 2 describes the primary estimand used to determine the extent of pro-attitudinal news engagement (e.g., polarization, segregation). Column 3 describes how the slant of articles or domains is determined. Column 4 describes whether the analysis of pro-attitudinal engagement is measured at the outlet or the article level. Columns 5 and 6 describe the sources of data on browsing and social media behavior, respectively. Column 7 describes

whether social media engagement is analyzed at the clicks, views, or shares level. Finally, Column 8 presents the number of articles or domains used to compute the primary estimand. We define polarization as the difference in the mean slant consumed by registered Republicans and Democrats and define segregation as the difference between the share of Republican visits that visit outlets (or articles) predominantly visited by Republicans and the share of Republican visits that visit outlets (or articles) predominantly visited by Democrats. Note that different papers in the literature refer to different measures by the same name and to the same measure by different names. For instance, although [Flaxman, Goel and Rao \(2016\)](#) refer to their measure as segregation, for the purposes of this paper we refer to it as standard deviation in consumption. We reserve the term segregation for the measure described in [Appendix B](#). Similarly, [Gentzkow and Shapiro \(2011\)](#) refer to their measure as the isolation index, but [González-Bailón et al. \(2023\)](#) refer to it as the segregation index. In this paper, we refer to both as the segregation because they match the construction in [Appendix B](#). The table only refers to measures and data used for the primary estimates (e.g., segregation or polarization). For example, [Peterson, Goel and Iyengar \(2021\)](#) also have a crowd-sourced content-based measure of slant but do not use this measure to compute the primary estimand. [Green et al. \(2025\)](#) also study an audience-based measure of slant and analyze Facebook data; we focus on the content-based measure since it is the most similar to this paper.

Table A.2: Summary Statistics

	(1) All mean/sd	(2) Not in SS1 mean/sd	(3) In SS1 mean/sd
Opinion piece (binary)	0.22 (0.41)	0.20 (0.40)	0.29 (0.45)
Slant	-0.21 (0.95)	-0.20 (0.86)	-0.26 (1.22)
FB views (1000s)			219.40 (765.56)
FB clicks (1000s)			9.32 (33.11)
FB shares (1000s)			1.81 (7.81)
N	1054124	822303	231821

Notes: This table reports summary statistics for all news articles published by the top 100 U.S. news outlets in Column 1. See [Appendix Table C.1](#) for the full list of outlets. Column 2 shows summary statistics for the subset of articles not found in the URL-level Facebook Activity Dataset; Column 3 shows summary statistics for the articles that do appear in the URL-level Facebook Activity Dataset (the dataset only includes articles shared publicly at least 100 times). The table reports means, as well as standard deviations in parentheses.

Table A.3: Engagement and Polarization in News Exposure, Consumption, and Circulation on Facebook by Political Affinity

Panel A: Engagement Based on Political Affinity						
Political Affinity	Views		Clicks		Shares	
	Share(%)	Mean Slant	Share(%)	Mean Slant	Share(%)	Mean Slant
Liberal	36%	-0.83	38%	-0.87	47%	-0.94
Moderate	11%	-0.49	8%	-0.44	8%	-0.46
Conservative	25%	0.34	28%	0.46	32%	0.77
Unassigned	28%	-0.40	26%	-0.36	13%	-0.30

Panel B: Polarization Based on Political Affinity			
	Views	Clicks	Shares
Average Slant Difference	1.17	1.33	1.7
Polarization Index	0.23	0.27	0.34

Notes: This table shows the summary statistics of all views, clicks, and shares in the Social Science One data based on users’ political affinity. The table is based on the sample of $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., publicly shared on Facebook at least 100 times). Panel A presents two statistics for each user political affinity group and engagement type (views, clicks, shares): the share (%) of all news article engagements of a given engagement type that is accounted for by a given political affinity group (accordingly, columns sum to 100%), and the views-, clicks-, and shares-weighted average slant (“Mean Slant”) of news article views, clicks, and shares, respectively, by user political affinity. Liberal, moderate and conservative users are defined as those with a political affinity score below 0, equal to 0 and above 0, respectively; those with missing political affinity scores are grouped into the category “Unassigned.” Panel B reports, for each engagement type, the difference between the mean slant of news diets by liberal and conservative users on Facebook in the first row. The last row shows the polarization index: it normalizes the difference by dividing it by 5, so that -1 indicates extreme counter-attitudinal news consumption, 0 indicates no difference in average slant across the news consumed by liberals and conservatives, and 1 indicates extreme pro-attitudinal news consumption. Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Table A.4: Confusion Matrix: Political Affinity and Party Registration

Political Affinity	Party Registration		
	RR	RD	NR
Liberal	0.04	0.38	0.58
Moderate	0.11	0.24	0.65
Conservative	0.52	0.06	0.41
Missing	0.16	0.23	0.61

Notes: This table shows, for each political affinity group $g \in \{L, M, C, \text{missing}\}$, the share of users who have a given party registration $p \in \{RR, RD, NR\}$ (hence, each row sums up to 1). The values in the row with “Missing” political affinity are imputed using the population shares of party registration as per [American National Election Studies \(2020\)](#). The values in all other rows come from [Barberá et al. \(2015\)](#), based on Twitter (now “X”) users.

Table A.5: Polarization in News Exposure, Consumption, and Circulation on Facebook: Opinion vs. Non-Opinion Articles

	Average Slant Reg. Democrats	Average Slant Reg. Republicans	Average Slant Difference	Polarization Index
Panel A: Opinion pieces				
Views	-0.86	0.59	1.45	0.29
Clicks	-0.87	0.67	1.54	0.31
Shares	-0.99	1.22	2.21	0.44
Panel B: Non-opinion pieces				
Views	-0.55	-0.12	0.43	0.09
Clicks	-0.54	-0.01	0.53	0.11
Shares	-0.60	0.17	0.77	0.15

Notes: This table shows the same statistics as in [Table 1](#) for all opinion news articles ($N = 66,544$) published online in the top-100 most-read U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook) in Panel A, and their non-opinion counterparts in Panel B. The classification into opinion articles is described in [Appendix H](#).

Table A.6: Polarization in News Exposure, Consumption, and Circulation on Facebook: Excluding Outlets that Feed Into Facebook’s Political Affinity Measure

	Average Slant Reg. Democrats	Average Slant Reg. Republicans	Average Slant Difference	Polarization Index
Views	-0.50	0.15	0.65	0.13
Clicks	-0.52	0.23	0.75	0.15
Shares	-0.60	0.62	1.22	0.24


Notes: This table shows the same statistics as in [Table 1](#), only that when constructing the statistics, we omit all articles from the 17 outlets that feature in the computation of Facebook users’ political affinity scores in the URL-level Facebook Activity Dataset (see [Section 3.2](#) for details).

Table A.7: Polarization in News Exposure, Consumption, and Circulation on Facebook when Slant is Measured at Outlet-Level: Excluding Outlets that Feed Into SS1’s Political Affinity Measure

	Average Slant Reg. Democrats	Average Slant Reg. Republicans	Average Slant Difference	Polarization Index
Views	-0.35	0.01	0.36	0.07
Clicks	-0.39	0.08	0.47	0.09
Shares	-0.38	0.24	0.62	0.12


Notes: This table shows the same statistics as in [Table 1](#) with two differences. First, when constructing the statistics, we omit all articles from the 17 outlets that feature in the computation of Facebook users’ political affinity scores in the URL-level Facebook Activity Dataset (see [Section 3.2](#) for details). Second, we calculate slant at the outlet level (i.e., we assign each article the average slant across all articles published by the outlet).

Figure A.1: Distribution of Expert Slant Ratings




Notes: This figure shows the distribution of the expert slant ratings for the $N = 3,759$ articles they rated. The combined labels are created by taking the average raw score of the two raters after the RA-correction for cases where the raters did not agree on whether an article is liberal or conservative. More details are provided in [Section 4.1](#).

Figure A.2: Correlation between Model-Predicted Slant and Third-Party Expert Slant Ratings




Notes: This figure shows, for each model-predicted slant bin (horizontal axis), the average third-party slant ratings (by an expert panel curated by Ad Fontes Media) across all articles in that bin. For a total of $N = 512$ articles from 2019 for which Ad Fontes Media ratings could be retrieved. To render the two scales comparable, each type of slant measure is standardized to mean 0 and standard deviation 1.

Figure A.3: Average Absolute Slant by Topic




Notes: This figure shows the average absolute slant of the articles falling in each topic category listed on the vertical axis. Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The absolute value of the slant scale goes from 0 (a centrist article) to 2.5 (a very extreme article). The classification into topics is described in detail in [Appendix H](#).

Figure A.4: Over- vs. Under-Representation of Topics among Right-Leaning Articles




Notes: This figure shows the degree of over- vs. under-representation of topics among right-leaning articles. The figure is constructed as follows. First, all “moderate” articles with a slant between -0.5 and 0.5 are dropped. Second, the remaining articles are defined as “right-leaning” if they have a slant larger than 0.5 and “left-leaning” otherwise. Third, we compute the fraction of right-leaning articles overall. Fourth, we compute the fraction of right-leaning articles for each topic. Fifth, we divide the fraction of right-leaning articles within a topic by the overall fraction of right-leaning articles and subtract one. The number we obtain—and that is plotted in the figure—is an estimate of the degree to which a particular topic is over- vs. under-represented among right-leaning articles.

Figure A.5: Slant Distribution within Topics




Notes: This figure shows the distribution of slant within news topics. The vertical blue line indicates the mean slant for the given topic. Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The classification into topics is described in detail in [Appendix H](#).

Figure A.6: Standard Deviation of Slant by Topic




Notes: This figure shows the standard deviation of the slant of the articles falling in each topic category listed on the vertical axis. Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The classification into topics is described in detail in [Appendix H](#).

Figure A.7: Terms Most Predictive of Article Slant




Notes: This figure presents the top ten n-grams, including both unigrams and bigrams, identified through LASSO regression as predictive of slant in news articles, based on our slant measure. These terms were extracted from article titles, with a minimum occurrence threshold of 1,000. The horizontal bar plot displays the Lasso coefficients of these n-grams, with positive values indicating a right-leaning slant and negative values indicating a left-leaning slant.

Figure A.8: Polarization Index for News Exposure, Consumption, and Circulation on Facebook Measured at Article and Outlet Level




Notes: This figure shows the polarization index—i.e. views-/clicks-/shares-weighted average difference in the slant of news diets by registered Republican and Democratic users on Facebook divided by 5, so that -1 indicates extreme counter-attitudinal news consumption, 0 indicates no difference in average slant across the news consumed by Democrats and Republicans, and 1 indicates extreme pro-attitudinal news consumption. The blue bars on the left show the polarization index when slant is calculated using our main article-level slant measure (as in [Table 1](#)). The white bars on the right show the polarization index when we assign each article its outlet-level slant measure (constructed as the average slant across all articles published by the outlet). The news article sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset.

Figure A.9: Average Slant of Articles Clicked on Facebook by Topic and User Political Affinity




Notes: This figure shows the clicks-weighted average slant of articles by Facebook users' political affinity and topic. The sample includes all news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook). Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure A.10: Average Slant of Articles Viewed on Facebook by Outlet and User Political Affinity




Notes: This figure shows the views-weighted average slant of articles by Facebook users' political affinity and outlet. The sample includes all news articles published online by the top 20 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook). Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure A.11: Empirical PDF of Article Slant in Full Sample vs. SS1 Sample




Notes: This figure shows the distribution of article-level slant for two samples: the full sample of all 1,054,124 news articles published online by the top-100 U.S. news outlets in 2019 (left panel) and the sub-sample of all 231,821 articles shared more than 100 times on Facebook (right panel)—i.e., the articles included in the Facebook engagement dataset (SS1). Article Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure A.12: Shares of Opinion Pieces and National News




Notes: This figure presents the fraction of articles, both in our dataset (“Production”) and in the URL-level Facebook Activity dataset (“Shared on FB > 100x”), classified as national news and as opinion pieces. The classification into national, international, and local news, as well as the one into opinion vs. non-opinion pieces, is described in [Appendix H](#).

Figure A.13: Under- vs. Over-representation of Topics among Articles Shared on FB more than 100x




Notes: This figure shows that topics exhibiting a high standard deviation of slant tend to be over-represented in the URL-level Facebook Activity dataset. The x-axis captures the standard deviation of article-level slant for each topic. The y-axis captures the degree of under- vs over-representation. Specifically, we first compute the overall fraction of articles in each topic. Second, we compute the same fractions for articles in the URL-level Facebook Activity dataset. Third, we divide the latter number by the former and subtract one. This way, the quantity we obtain can be thought of as the degree of under- vs over-representation of articles in the URL-level Facebook Activity Dataset compared to our overall database of articles. The classification into topics is described in detail in [Appendix H](#).

Figure A.14: Empirical PDF of Facebook Shares by Facebook User Political Affinity and Article Slant Net of Outlet Fixed Effects



Notes: This figure shows the fraction of all shares by a given Facebook user political affinity group (liberal and conservative, respectively) that accrue to articles in a given residualized slant bin. The dashed vertical line represents the mean of the distribution. The sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook). Residualized slant is constructed as follows: We take our article-level slant measure (given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#)), regress it on outlet fixed effects, and obtain the residuals, to which we add the unconditional mean slant across all articles from the full sample. The (residualized) slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).


Figure A.15: Average Slant of Articles Shared on Facebook by Outlet and User Political Affinity




Notes: This figure shows the shares-weighted average slant of articles by Facebook users' political affinity and outlet. The sample includes all news articles published online by the top 20 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook). Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure A.16: Empirical PDF of the Slant of Articles Shared on Facebook With and Without Clicks by User Political Affinity

(a) Shares with no Clicks



(b) Shares with Clicks



Notes: This figure shows the fraction of all shares by a given Facebook user political affinity group (liberal and conservative, respectively) that accrue to articles in a given slant bin. The figure shows the results separately for news article URLs that users clicked on before sharing on Facebook and for news article URLs that users shared without clicking. The sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

B. Polarization, Segregation, and Favorability

B.1 Empirics

In the empirical part of this appendix, we substantiate the claims made in the introduction about the difference between our content-based article-level slant measure and audience-based article-level favorability scores, as well as the difference between our polarization index and the segregation (or isolation) index.⁴⁴

To begin with, we provide definitions of the notions of favorability score and segregation (or isolation) index. [González-Bailón et al. \(2023\)](#) define the favorability score of URL α as follows

$$fav_{\alpha} = \frac{C_{\alpha} - L_{\alpha}}{C_{\alpha} + L_{\alpha}}$$

where α indicates a particular URL, C_{α} indicates the number of conservative clicks on URL α , and L_{α} indicates the number of liberal clicks on URL α . Intuitively, this is an audience-based measure of the URL’s leaning, where URLs with more conservative clicks getting higher score.

[Gentzkow and Shapiro \(2011\)](#) define the segregation (or isolation) index as:

$$S = \sum_{\alpha \in A} \left(\frac{C_{\alpha}}{\bar{C}} \times \frac{C_{\alpha}}{v_{\alpha}} \right) - \sum_{\alpha \in A} \left(\frac{L_{\alpha}}{\bar{L}} \times \frac{C_{\alpha}}{v_{\alpha}} \right)$$

where A is the set of URLs under consideration, $\bar{C} = \sum_{\alpha \in A} C_{\alpha}$ indicates the total number of conservative clicks for any URL, $\bar{L} = \sum_{\alpha \in A} L_{\alpha}$ indicates the total number of liberal clicks for any URL, and $v_{\alpha} = C_{\alpha} + L_{\alpha}$ indicates the total number of liberal and conservative clicks for URL α . Intuitively, segregation measures the difference in the share of conservatives among the audiences of outlets (or articles) visited by conservatives and the share of conservatives among the audiences of outlets (or articles) visited by liberals.

Favorability Scores We first analyze the differences between our measure of slant and the favorability scores, and then segue into the differences between our polarization index and the segregation (or isolation) index.

In order to study the relationship between our content-based article-level slant measure and the audience-based article-level favorability scores, we proceed in two steps. First, following the procedure outlined in [Buntain et al. \(2023\)](#) to account for differential privacy in the URL-level Facebook Activity dataset, we restrict our attention to the set of articles clicked on Facebook more than 8,962 times. This restriction ensures that the signal-to-noise ratio is sufficiently high to estimate meaningful favorability scores in spite of differential privacy. Second, we calculate the favorability scores for those articles, as described in [González-Bailón et al. \(2023\)](#).⁴⁵

⁴⁴The names segregation index and isolation index are sometimes used interchangeably in the literature. See, for instance, [Gentzkow and Shapiro \(2011\)](#) and [González-Bailón et al. \(2023\)](#).

⁴⁵The version of the favorability score calculated in [Buntain et al. \(2023\)](#) and in [González-Bailón et al. \(2023\)](#) are slightly different but extremely highly correlated (Pearson correlation: 0.993). The threshold we calculate to address the issues that arise due to differential privacy is based on the measure in [Buntain et al. \(2023\)](#). In spite of that, we calculate and discuss the favorability score from [González-Bailón et al. \(2023\)](#) throughout this appendix, as it is more easily interpretable and more standard. Given the extremely high correlation between the two versions of the favorability scores, this decision should make little practical difference.

Our main empirical result when comparing the audience-based favorability scores to our content-based slant measure is that these two measures capture distinct concepts and one cannot easily be considered a proxy for the other.⁴⁶

Appendix Figure B.1 overlays the distribution of our slant measure and the favorability scores for the set of articles clicked on Facebook more than 8,962 times. The figure shows that the favorability score exhibits a strongly bimodal distribution, whereas our slant measure exhibits a distribution that is much closer to unimodality.

The relationship between the two measures is further clarified in Appendix Figure B.2, which presents a binned scatter plot of our slant measure on the horizontal axis versus the favorability score on the vertical axis. The figure exhibits a marked S-shaped pattern, indicating that favorability scores exhibit compression among both left-leaning and right-leaning articles.

Appendix Figure B.3 plots the distribution of favorability scores, after restricting attention to the set of articles to which our slant measure assigns a value of zero. Even among articles that are considered to be neutral based on our content-based slant measure, the distribution of favorability scores is strongly bimodal. This finding echoes the concern raised in the introduction that favorability scores might assign partisan value to centrist articles published in partisan outlets simply because partisan outlets tend to have a partisan readership. In other words, an article from an outlet that is visited primarily by liberals (conservatives) is likely to be visited and shared primarily by liberals (conservatives) even when the article is moderate. Indeed, among the set of articles to which our measure of slant assigns a value of zero, we find that the correlation between the favorability score and the audience-based measure of slant at the outlet level from Bakshy, Messing and Adamic (2015) is 0.85.⁴⁷ Thus, articles that, according to our measure, are moderate appear to inherit their favorability score from the partisanship of the outlet they come from.

Finally, Appendix Figure B.4 examines articles with favorability scores near zero (between -0.5 and 0.5) and plots the distribution of their slant values. Unlike Appendix Figure B.1, this distribution is unimodal with a moderate mode. Thus, articles that our measure classifies as moderate are generally not classified as moderate by the favorability score, whereas articles that are generally classified as moderate by the favorability score are also largely classified as moderate by our slant measure.⁴⁸

As an extreme example of articles inheriting their favorability scores from the partisanship of the outlet they come from, let us consider a specific article written by the Associated Press (AP), a

⁴⁶Throughout this section, we use a click-based favorability score to ensure that the favorability score is consistent with the segregation index and our main polarization index, which are based on clicks. The results are very similar when using a shares-based favorability score.

⁴⁷Bakshy, Messing and Adamic (2015) determine the slant of a domain based on the ideology of individuals who shared news from the domain, similarly to a favorability score.

⁴⁸A research assistant manually reviewed the 37 articles that our measure of slant classified as extreme (i.e., with an absolute slant measure greater than or equal to 2) and the favorability scores classified as moderate (with a favorability score between -0.5 and 0.5). The research assistant, who was informed neither about how the articles were selected nor about the purpose of the task, was simply asked to assign a measure of slant to each of the 37 articles according to the criteria specified in Appendix E.1. For 70% of the articles, the RA's ratings confirmed both the direction of slant and the fact that the article was extreme. For example, the following opinion piece by Bret Stephens was considered moderate by the favorability score (probably because it appeared in the New York Times), but we considered it clearly conservative: <https://www.nytimes.com/2019/01/25/opinion/venezuela-maduro-socialism-government.html>. In all but two of the remaining articles, the RA confirmed the direction of slant but not the fact that the article was extreme. This exercise shows that, for articles that are considered extreme by our measure of slant and moderate by the favorability scores, the vast majority of discrepancies are adjudicated in favor of our measure of slant.

portion of which is reported at the end of this section. This article appeared, verbatim, on the websites of ABC News, Business Insider, NBC News, and the New York Post. The title of the article varied slightly depending on the outlet.⁴⁹ Our slant measure assigns to the article, independently of the title, a slant rating of zero. Conversely, the favorability score of the article varies greatly depending on the outlet that published it, as shown in the table below.

	Favorability Score	Content-based Outlet Slant
Business Insider	-0.53	-0.65
NBC News	-0.50	-0.62
ABC News	-0.09	-0.50
The New York Post	0.55	0.40

Note that an article with the exact same text receives a liberal score when appearing in more liberal outlets (NBC News and Business Insider) and a conservative score when appearing in a more conservative outlet (The New York Post). More generally, once again, the correlation between the favorability score and the audience-based measure of slant at the outlet level from [Bakshy, Messing and Adamic \(2015\)](#) is high (0.84). The high degree of correlation between the favorability scores of the same AP article published by different outlets and a measure of outlet-level slant suggests that, indeed, articles can inherit their favorability scores from the partisanship of the audience of the outlet they are published in.

Segregation Index The example above allows us to segue into the difference between our polarization index in news consumption and the segregation (or isolation) index. Specifically, in line with the discussion in the introduction, we can calculate both our polarization index and the shares-based segregation index for the articles from the example above. We find that the polarization index equals 0, whereas the segregation index equals 0.10, which corresponds to 21% of the segregation index reported in [González-Bailón et al. \(2023\)](#).⁵⁰ This example shows how the segregation index can in principle be inflated in the presence of articles with a similar slant (in this case, identical articles) that appear in outlets with a different partisan composition of readership. By being based on a purely content-driven measure of slant at the article level, our index of polarization does not run into this problem.

To more broadly illustrate the point that the segregation index can be inflated when articles with a similar slant appear in outlets that differ in the partisanship of their readers, we carry out an additional exercise. First, we calculate the clicks-based segregation index on the overall set of articles that meet the differential privacy threshold described at the beginning of this appendix; second, we calculate the clicks-based segregation index on the subset of those articles that our

⁴⁹ABC News titled the article “APNewsBreak: US approved thousands of child bride requests.” Business Insider titled it “The US government approved thousands of requests for men to bring child brides to the US.” NBC News titled it “U.S. approved thousands of child bride requests, lives ruined.” Lastly, the New York Post titled it “US approved thousands of child bride requests over last 10 years.”

⁵⁰[González-Bailón et al. \(2023\)](#) do not provide a click-based segregation index so we use their exposure-based index instead for this comparison. This is a conservative comparison as views are typically less segregated than clicks. When using a shares-based index, the segregation among the articles from the example above is 0.19 and corresponds to 35% of the engagement-based segregation index reported in [González-Bailón et al. \(2023\)](#).

content-based measure of slant classifies as moderate (i.e., articles that are assigned a slant rating of zero). We find that clicks-based segregation index on the overall set of articles is 0.63 and that, when we limit our attention to articles that are classified as moderate by our slant measure, the index takes a value of 0.51. Thus, the clicks-based segregation index on moderate articles is 81% of the clicks-based segregation index on all articles. This finding, once again, suggests that the segregation index can be driven by moderate articles published on outlets with a partisan readership.

Article “WASHINGTON — Thousands of requests by men to bring in child and adolescent brides to live in the United States were approved over the past decade, according to government data obtained by The Associated Press. In one case, a 49-year-old man applied for admission for a 15-year-old girl.

The approvals are legal: The Immigration and Nationality Act does not set minimum age requirements. And in weighing petitions for spouses or fiancées, US Citizenship and Immigration Services goes by whether the marriage is legal in the home country and then whether the marriage would be legal in the state where the petitioner lives.

But the data raises questions about whether the immigration system may be enabling forced marriage and about how US laws may be compounding the problem despite efforts to limit child and forced marriage. Marriage between adults and minors is not uncommon in the United States and most states allow children to marry with some restrictions.

[...]”

B.2 Theory

In the theoretical part of this appendix, we introduce a very simple example to illustrate two key points. The first point is that the relationship between our polarization index at the outlet and at the article levels depends, in an intuitive fashion, on whether partisans consume pro- or counter-attitudinal news within outlets. Specifically, if partisans consume pro-attitudinal news within outlets, polarization in news consumption at the article level will be higher than polarization in news consumption at the outlet level; conversely, if partisans consume counter-attitudinal news within outlets, polarization at the article level will be smaller than at the outlet level. Lastly, if partisans read randomly within outlets, polarization will be the same at the article and at the outlet levels.

The second key point is that measures of segregation suffer from a small-sample bias. As a result, it has been argued that, in small samples, measures of segregation at finer levels of aggregation tend to be higher than measures of segregation at coarser levels of aggregation ([Gentzkow, Shapiro and Taddy, 2019](#); [Yao et al., 2019](#); [Wong, 2003](#)).

Formally, let M denote the set of news articles, with $M = \{1, 2, 3, 4\}$. Let N denote the set of news outlets, with $N = \{\alpha, \beta\}$. Let articles 1 and 2 belong to outlet α and articles 3 and 4 belong to outlet β . Let there be two Democrats and two Republicans and suppose each reads a single article picked at random according to categorical distribution $Categorical(p_1^z, p_2^z, p_3^z, p_4^z)$, where the distribution is allowed to differ between Democrats and Republican (i.e., $z \in \{D, R\}$). Let $s_i \in \mathbb{R}$ denote the slant of article $i \in M$ and assume $s_1 < s_2$ and $s_3 < s_4$. Thus, each outlet has a relatively more liberal and a relatively more conservative article. Let $\mathbf{S}_a = (s_1, s_2, s_3, s_4)$ be a vector describing the slant of each article and $\mathbf{S}_o = (\frac{s_1+s_2}{2}, \frac{s_3+s_4}{2})$ be a vector describing the slant of each outlet.

Let $\mathbf{X}^z \in \mathbb{R}^4$ for $z \in \{D, R\}$ be a random vector describing the number of Democrats or Republicans who are exposed to each article $i \in M$. $\mathbf{X}^z$ has the following distribution

$$\mathbf{X}^z = \begin{cases} [1, 1, 0, 0] & \text{w.p. } 2 \cdot p_1^z \cdot p_2^z \\ [1, 0, 1, 0] & \text{w.p. } 2 \cdot p_1^z \cdot p_3^z \\ [1, 0, 0, 1] & \text{w.p. } 2 \cdot p_1^z \cdot p_4^z \\ [0, 1, 1, 0] & \text{w.p. } 2 \cdot p_2^z \cdot p_3^z \\ [0, 1, 0, 1] & \text{w.p. } 2 \cdot p_2^z \cdot p_4^z \\ [0, 0, 1, 1] & \text{w.p. } 2 \cdot p_3^z \cdot p_4^z \\ [2, 0, 0, 0] & \text{w.p. } (p_1^z)^2 \\ [0, 2, 0, 0] & \text{w.p. } (p_2^z)^2 \\ [0, 0, 2, 0] & \text{w.p. } (p_3^z)^2 \\ [0, 0, 0, 2] & \text{w.p. } (p_4^z)^2 \end{cases}$$

Thus $E[\mathbf{X}^z] = (2p_1^z, 2p_2^z, 2p_3^z, 2p_4^z)$. The average expected slant that Democrats are exposed to is, therefore, $\frac{1}{2}\{E[\mathbf{X}^D] \cdot \mathbf{S}'_a\}$; similarly, the average expected slant that Republicans are exposed to is $\frac{1}{2}\{E[\mathbf{X}^R] \cdot \mathbf{S}'_a\}$. But then, in this simple example, expected polarization at the article level can be written as

$$E[\mathcal{P}_{article}] = \frac{\frac{1}{2}\{E[\mathbf{X}^R] \cdot \mathbf{S}'_a\} - \frac{1}{2}\{E[\mathbf{X}^D] \cdot \mathbf{S}'_a\}}{5} = \frac{1}{10}\{E[\mathbf{X}^R] - E[\mathbf{X}^D]\} \cdot \mathbf{S}'_a$$

Let $\mathbf{Y}^z \in \mathbb{R}^2$ for $z \in \{D, R\}$ be a random vector describing the number of Democrats or Republicans who are exposed to each outlet $j \in N$. $\mathbf{Y}^z$ has the following distribution

$$\mathbf{Y}^z = \begin{cases} [1, 1] & \text{w.p. } 2(p_1^z \cdot p_3^z + p_1^z \cdot p_4^z + p_2^z \cdot p_3^z + p_2^z \cdot p_4^z) \\ [2, 0] & \text{w.p. } 2(p_1^z \cdot p_2^z) + (p_1^z)^2 + (p_2^z)^2 \\ [0, 2] & \text{w.p. } 2(p_3^z \cdot p_4^z) + (p_3^z)^2 + (p_4^z)^2 \end{cases}$$

Thus:

$$E[\mathbf{Y}^z] = [2(p_1^z \cdot p_3^z + p_1^z \cdot p_4^z + p_2^z \cdot p_3^z + p_2^z \cdot p_4^z) + 4(p_1^z \cdot p_2^z) + 2(p_1^z)^2 + 2(p_2^z)^2, \\ 2(p_1^z \cdot p_3^z + p_1^z \cdot p_4^z + p_2^z \cdot p_3^z + p_2^z \cdot p_4^z) + 4(p_3^z \cdot p_4^z) + 2(p_3^z)^2 + 2(p_4^z)^2] \quad (2)$$

The average expected slant that Democrats are exposed to is, therefore, $\frac{1}{2}\{E[\mathbf{Y}^D] \cdot \mathbf{S}'_o\}$; similarly, the average expected slant that Republicans are exposed to is $\frac{1}{2}\{E[\mathbf{Y}^R] \cdot \mathbf{S}'_o\}$. But then, in this simple example, expected polarization at the outlet level can be written as

$$E[\mathcal{P}_{outlet}] = \frac{\frac{1}{2}\{E[\mathbf{Y}^R] \cdot \mathbf{S}'_o\} - \frac{1}{2}\{E[\mathbf{Y}^D] \cdot \mathbf{S}'_o\}}{5} = \frac{1}{10}\{E[\mathbf{Y}^R] - E[\mathbf{Y}^D]\} \cdot \mathbf{S}'_o$$

As a result,

$$E[\mathcal{P}_{article} - \mathcal{P}_{outlet}] = \frac{1}{10}\{E[\mathbf{X}^R] - E[\mathbf{X}^D]\} \cdot \mathbf{S}'_a - \frac{1}{10}\{E[\mathbf{Y}^R] - E[\mathbf{Y}^D]\} \cdot \mathbf{S}'_o$$

The expression above can be simplified to:

$$E[\mathcal{P}_{article} - \mathcal{P}_{outlet}] = \frac{1}{10} \left[\underbrace{(s_2 - s_1)}_{>0} [(p_1^D - p_2^D) - (p_1^R - p_2^R)] + \underbrace{(s_4 - s_3)}_{>0} [(p_3^D - p_4^D) - (p_3^R - p_4^R)] \right]$$

To illustrate how the relationship between our polarization index at the outlet and at article levels depends on whether partisans consume pro- or counter-attitudinal news within outlet, we consider three cases:

1. Pro-attitudinal news consumption within outlets, which corresponds to the following assumption: $p_1^D > p_2^D, p_3^D > p_4^D, p_1^R < p_2^R, p_3^R < p_4^R$
2. Neutral news consumption within outlets, which corresponds to the following assumption: $p_1^D = p_2^D, p_3^D = p_4^D, p_1^R = p_2^R, p_3^R = p_4^R$
3. Counter-attitudinal news consumption within outlets, which corresponds to the following assumption: $p_1^D < p_2^D, p_3^D < p_4^D, p_1^R > p_2^R, p_3^R > p_4^R$

It is easy to show that, for pro-attitudinal news consumption within outlets $\mathcal{P}_{article} - \mathcal{P}_{outlet} > 0$, for neutral news consumption within outlets $\mathcal{P}_{article} - \mathcal{P}_{outlet} = 0$, and for counter-attitudinal news consumption within outlets $\mathcal{P}_{article} - \mathcal{P}_{outlet} < 0$. Thus, the relationship between our polarization index at the outlet and at the article levels depends, in an intuitive fashion, on whether partisans consume pro- or counter-attitudinal news within outlets. This illustrates the first point.

The second point that this appendix aims to illustrate is that, in our very simple example, segregation at the article level is weakly higher than segregation at the outlet level and, generically, strictly higher. This result reflects a general concern in the spatial segregation literature, namely that, in small samples, measures of segregation at finer levels of aggregation tend to be higher than measures of segregation at coarser levels of aggregation (Yao et al., 2019). Thus, in small samples, it is hard to obtain an apple-to-apple comparison of segregation at the outlet and at the article level.

Consider, once again, the environment set up at the beginning of this appendix. The segregation index at the article level can be written as

$$\mathcal{S}_{article} = \frac{1}{2} \sum_{m \in M} \left(\frac{R_m^2}{v_m} \right) - \frac{1}{2} \sum_{m \in M} \left(\frac{D_m \cdot R_m}{v_m} \right)$$

where R_m is the m^{th} component of $\mathbf{X}^R$, D_m is the m^{th} component of $\mathbf{X}^D$, and $v_m = R_m + D_m$. Thus, the expected segregation index at the article level can be written as

$$E[\mathcal{S}_{article}] = \sum_{m \in M} \left[p_m^R - 2 \cdot p_m^R \cdot p_m^D + \frac{2}{3} \cdot p_m^R \cdot (p_m^D)^2 + \frac{2}{3} \cdot (p_m^R)^2 \cdot p_m^D - \frac{1}{3} \cdot (p_m^R \cdot p_m^D)^2 \right]$$

The segregation index at the outlet level can be written as

$$\mathcal{S}_{outlet} = \frac{1}{2} \sum_{n \in N} \left(\frac{R_n^2}{v_n} \right) - \frac{1}{2} \sum_{n \in N} \left(\frac{D_n \cdot R_n}{v_n} \right)$$

where R_n is the n^{th} component of $\mathbf{Y}^R$, D_n is the n^{th} component of $\mathbf{Y}^D$, and $v_n = R_n + D_n$. Thus, the expected segregation index at the outlet level can be written as

$$\begin{aligned} E[\mathcal{S}_{outlet}] &= \frac{2}{3}(p_1^R + p_2^R) + \frac{2}{3}(p_1^D + p_2^D) + \frac{1}{3}(p_1^R + p_2^R)^2 + \frac{1}{3}(p_1^D + p_2^D)^2 \\ &\quad - \frac{8}{3}(p_1^R + p_2^R) \cdot (p_1^D + p_2^D) + \frac{2}{3}(p_1^R + p_2^R) \cdot (p_1^D + p_2^D)^2 \\ &\quad + \frac{2}{3}(p_1^R + p_2^R)^2 \cdot (p_1^D + p_2^D) - \frac{2}{3}(p_1^R + p_2^R)^2 \cdot (p_1^D + p_2^D)^2 \end{aligned}$$

The following proposition holds.

Proposition: $E[\mathcal{S}_{article}] - E[\mathcal{S}_{outlet}] \geq 0$. Furthermore, $E[\mathcal{S}_{article}] - E[\mathcal{S}_{outlet}] = 0$ if and only if $p_1^R p_2^D = p_2^R p_1^D = p_4^D p_3^R = p_4^R p_3^D = 0$.

Proof

It is possible to simplify the expression for $E[\mathcal{S}_{article}] - E[\mathcal{S}_{outlet}]$ as follows:

$$E[\mathcal{S}_{article}] - E[\mathcal{S}_{outlet}] = \frac{1}{3}(G_1 + G_2)$$

, where

$$\begin{aligned} G_1 &= ((p_1^D)^2 + (p_2^D)^2)p_1^R p_2^R \\ &\quad + p_1^R p_2^D (p_3^D + p_4^D + 1 - p_1^D)p_3^R \\ &\quad + p_1^R p_2^D (p_3^D + p_4^D + 1 - p_1^D)p_4^R \\ &\quad + p_1^R p_2^D (p_3^D + p_4^D + 1 - p_1^D) \\ &\quad + p_2^R p_1^D (p_3^D + p_4^D + 1 - p_2^D)p_3^R \\ &\quad + p_2^R p_1^D (p_3^D + p_4^D + 1 - p_2^D)p_4^R \\ &\quad + p_2^R p_1^D (p_3^D + p_4^D + 1 - p_2^D) \\ &\quad + 2p_1^R p_2^D (1 - p_2^R) + 2p_2^R p_1^D (1 - p_1^R) \end{aligned}$$


and

$$\begin{aligned} G_2 &= ((p_3^D)^2 + (p_4^D)^2)p_3^R p_4^R \\ &\quad + p_3^R p_4^D (p_1^D + p_2^D + 1 - p_3^D)p_1^R \\ &\quad + p_3^R p_4^D (p_1^D + p_2^D + 1 - p_3^D)p_2^R \\ &\quad + p_3^R p_4^D (p_1^D + p_2^D + 1 - p_3^D) \\ &\quad + p_4^R p_3^D (p_1^D + p_2^D + 1 - p_4^D)p_1^R \\ &\quad + p_4^R p_3^D (p_1^D + p_2^D + 1 - p_4^D)p_2^R \\ &\quad + p_4^R p_3^D (p_1^D + p_2^D + 1 - p_4^D) \\ &\quad + 2p_3^R p_4^D (1 - p_4^R) + 2p_4^R p_3^D (1 - p_3^R) \end{aligned}$$

Since all of the terms in G_1 and G_2 are non-negative, $G_1 \geq 0$, $G_2 \geq 0$, and, thus, $E[\mathcal{S}_{article}] - E[\mathcal{S}_{outlet}] \geq 0$

Furthermore, it is easy to see that $G_1 = 0$ if and only if $p_1^R p_2^D = p_2^R p_1^D = 0$. Similarly, it is easy to see that $G_2 = 0$ if and only if $p_4^D p_3^R = p_4^R p_3^D = 0$. Therefore, $E[\mathcal{S}_{article}] - E[\mathcal{S}_{outlet}] = 0$ if and

Figure B.1: Distribution of Favorability Scores and Slant



Notes: This figure shows the distribution of favorability scores (left panel) and slant (right panel) among news articles clicked on Facebook more than 8,962 times ($N = 38,064$) and published online by the top 100 U.S. news outlets in 2019. Our measure of Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The construction of the favorability scores is described in detail in [Appendix B](#).


only if $p_1^R p_2^D = p_2^R p_1^D = p_4^D p_3^R = p_4^R p_3^D = 0$. $\square$

This shows that, in the simple example above, the segregation index at the article level is always at least as large as the segregation index at the outlet level. In fact, generically, the segregation index at the article level is always strictly larger than the segregation index at the outlet level.⁵¹

Furthermore, the proposition shows that a necessary condition for segregation at the outlet level to equal segregation at the article level is that partisans read some articles with zero probability—a knife-edge case. Conversely, the difference between our polarization index at the article and at the outlet level equals zero precisely when partisans read randomly within outlets and, thus, consume the average slant of the outlet.


⁵¹The result holds generically, in the sense that the set of probability distributions $(p_1^z, p_2^z, p_3^z, p_4^z)$ for $z \in \{D, R\}$ where the result does not hold has measure zero when seen as a subset of $\mathbb{R}^8$.

Figure B.2: Correlation between Favorability Scores and Slant




Notes: This figure shows, for each slant bin according to our slant measure (horizontal axis), the average favorability score across all articles in that bin. Our measure of Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The construction of the favorability scores is described in detail in [Appendix B](#). The sample includes all news articles shared on Facebook more than 8,962 times ($N = 38,064$) and published online by the top 100 U.S. news outlets in 2019.

Figure B.3: Distribution of Favorability Scores: Moderate Slant



Notes: This figure shows the distribution of favorability scores among news articles that are classified as moderate (i.e., assigned a value of zero) according to our measure of slant. Our measure of Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The construction of the favorability scores is described in detail in [Appendix B](#). The sample includes all news articles shared on Facebook more than 8,962 times ($N = 38,064$) and published online by the top 100 U.S. news outlets in 2019.

Figure B.4: Distribution of Slant: Moderate Favorability Score



Notes: This figure shows the distribution of our measure of slant among news articles that are classified as moderate (i.e., assigned a value between -0.5 and 0.5) according to their favorability score. Our measure of Slant is given by the fine-tuned Gemma 3 27B’s prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The construction of the favorability scores is described in detail in [Appendix B](#). The sample includes all news articles shared on Facebook more than 8,962 times ($N = 38,064$) and published online by the top 100 U.S. news outlets in 2019.

C. Slant distribution for each outlet in the top 100

In this appendix, we present summary statistics and histograms of the within-outlet slant distribution for each of the top 100 U.S. outlets in terms of online visits in 2019.

Table C.1: Summary Statistics within Outlets

Outlet	Obs.		Mean SS1 Slant			Sd. Slant	% Lib Clicks
	All	SS1	All	Lib	Cons		
6abc	3,617	233	-0.44	-0.42	-0.43	0.44	52.19
9news	3,166	63	-0.45	-0.57	-0.27	0.60	54.17
ABC News	20,994	2,150	-0.45	-0.52	-0.35	0.52	63.60
Abc13	3,548	302	-0.43	-0.42	-0.40	0.46	43.72
Abc7	4,741	431	-0.44	-0.45	-0.40	0.45	61.73
Ajc	11,596	868	-0.60	-0.68	-0.46	0.66	58.76
Al	3,968	1,056	-0.58	-0.76	-0.43	0.88	50.50
American Thinker	7,040	2,577	2.27	2.21	2.28	0.64	9.69
Azcentral	4,071	651	-0.68	-0.83	-0.49	0.90	63.85
Baynews9	12,634	85	-0.38	-0.38	-0.32	0.52	45.30
Bloomberg	6,334	1,407	-0.53	-0.62	-0.31	0.66	72.56
Boston	5,692	134	-0.69	-0.83	-0.60	0.74	78.35
Boston Globe	17,317	802	-1.03	-1.14	-0.67	0.90	83.31
Breitbart	40,956	15,496	0.89	1.10	0.86	1.26	1.65
Buffalonews	2,772	183	-0.50	-0.65	-0.37	0.78	51.24
Business Insider	20,938	4,357	-0.64	-0.73	-0.47	0.62	72.13
Buzzfeed	1,265	353	-1.51	-1.50	-1.49	0.70	82.80
CBS News	10,324	5,051	-0.52	-0.56	-0.42	0.59	66.15
CNBC	15,034	3,593	-0.26	-0.34	-0.13	0.62	63.11
CNN	23,786	13,880	-0.63	-0.70	-0.36	0.69	82.32
Cbslocal	3,781	2,446	-0.41	-0.50	-0.31	0.42	51.16
Chicago Tribune	20,631	1,565	-0.67	-0.89	-0.23	0.99	70.72
Cincinnati	2,030	260	-0.25	-0.35	-0.02	0.81	58.94
Cleveland	7,973	461	-0.58	-0.74	-0.46	0.68	70.08
Click2Huston	9,608	291	-0.24	-0.28	-0.16	0.50	38.29
Clickondetroit	11,584	403	-0.31	-0.34	-0.28	0.48	56.44
Dallasnews	4,779	767	-0.47	-0.71	-0.07	0.86	62.06
Detroitnews	11,017	827	-0.10	-0.27	0.18	0.91	52.41
Forbes	2,644	1,036	-0.68	-0.91	-0.19	1.08	66.55
Fox News	34,809	15,458	0.50	0.48	0.47	1.16	5.98
Foxbusiness	7,825	1,296	0.44	0.53	0.42	0.94	9.20
Freep	2,718	834	-0.58	-0.60	-0.57	0.65	68.59

Table C.1: Summary Statistics within Outlets (Continued)

Outlet	Obs.		Mean SS1 Slant			Sd. Slant	% Lib Clicks
	All	SS1	All	Lib	Cons		
Heavy	2,387	184	-0.75	-1.04	-0.26	1.02	68.01
Hot Air	20,169	1,303	1.54	1.48	1.53	1.07	3.97
Houston Chronicle	29,641	248	-0.50	-0.67	-0.32	0.61	60.04
Huffington Post	13,622	9,517	-1.22	-1.24	-1.17	0.62	93.53
Indystar	1,570	273	-0.56	-0.62	-0.27	0.69	63.32
Inquirer	5,818	595	-0.89	-1.00	-0.69	0.54	74.18
Jsonline	2,047	438	-0.44	-0.47	-0.32	0.56	69.14
Khou	3,076	123	-0.26	-0.33	-0.23	0.54	36.97
Ksl	11,416	71	-0.34	-0.26	-0.30	0.66	45.14
Ktla	6,024	898	-0.49	-0.49	-0.48	0.54	54.23
LA Times	14,695	3,272	-1.01	-1.15	-0.65	0.77	78.23
MSNBC	8,376	3,986	-1.40	-1.44	-1.14	0.61	96.42
Masslive	5,242	248	-0.44	-0.49	-0.46	0.54	56.56
Mediaite	8,459	708	-0.46	-0.83	0.07	1.14	57.64
Mlive	4,655	1,012	-0.44	-0.47	-0.35	0.59	57.98
NBC News	12,156	7,657	-0.62	-0.65	-0.51	0.64	80.13
National Review	5,923	1,584	1.30	0.79	1.38	1.34	18.05
New York Post	15,701	4,859	0.44	0.06	0.75	1.10	34.82
New York Times	47,242	14,300	-0.90	-0.93	-0.70	0.74	88.17
Newsday	9,387	157	-0.48	-0.71	-0.35	0.76	44.12
Newser	3,615	10	-0.58	-1.41	-0.35	0.92	23.87
Newsmax	10,774	228	0.87	0.11	0.93	1.12	7.34
Newsweek	7,300	4,862	-0.87	-0.91	-0.66	0.56	84.53
Nj	6,860	885	-0.44	-0.52	-0.37	0.79	57.48
Nola	3,166	372	-0.65	-0.81	-0.54	0.72	57.00
Nydailynews	12,426	2,694	-0.67	-0.80	-0.29	0.87	79.61
Omaha	3,780	71	-0.39	-0.78	0.04	0.69	49.57
Oregonlive	5,101	843	-0.47	-0.61	-0.26	0.70	67.49
Palmerreport	4,925	2,406	-1.98	-1.97	-2.31	0.44	98.60
Patch Media	23,650	1,540	-0.37	-0.41	-0.34	0.62	55.12
Pennlive	4,990	424	-0.18	-0.42	0.16	0.80	36.44
Pjmedia	5,491	2,492	1.46	1.53	1.43	1.18	7.09
Politico	12,155	5,427	-0.60	-0.62	-0.42	0.69	85.36
Raw Story	19,900	7,808	-1.14	-1.14	-0.99	0.75	97.03
Realclearpolitics	1,465	248	1.98	0.43	2.18	1.25	11.38
San Francisco Gate	35,049	279	-0.42	-0.48	-0.21	0.63	75.57
Seattle Times	16,069	695	-0.71	-0.84	-0.45	0.73	74.83
Slate	4,634	1,802	-1.50	-1.52	-1.55	0.78	93.61

Table C.1: Summary Statistics within Outlets (Continued)

Outlet	Obs.		Mean SS1 Slant			Sd.	% Lib
	All	SS1	All	Lib	Cons	Slant	Clicks
Star Tribune	26,283	480	-0.43	-0.68	0.04	0.74	64.30
Staradvertiser	2,336	128	-0.48	-0.46	-0.43	0.57	68.06
Stltoday	8,858	411	-0.43	-0.56	-0.28	0.74	60.73
Syracuse	4,550	220	-0.36	-0.49	-0.13	0.70	53.19
The Atlantic	3,847	2,163	-1.53	-1.54	-1.32	0.82	90.65
The Blaze	7,471	4,417	0.97	1.03	0.96	1.13	3.35
The Epoch Times	14,353	3,832	0.82	0.75	0.91	1.01	16.81
The Hill	43,870	14,904	-0.32	-0.43	0.17	0.84	80.05
Thedailybeast	4,290	1,936	-1.08	-1.08	-1.02	0.88	91.79
Thefederalist	3,378	1,404	2.03	1.93	2.04	0.91	9.38
Time	6,368	2,081	-0.97	-1.04	-0.67	0.67	82.11
Townhall	12,385	4,351	1.70	1.80	1.67	1.00	4.40
Twitchy	9,315	684	1.88	1.76	1.89	0.99	3.16
USA Today	13,442	5,395	-0.51	-0.71	-0.13	0.86	64.35
Voanews	13,196	156	-0.65	-0.75	-0.49	0.53	81.87
Vox	5,475	2,977	-1.46	-1.47	-1.38	0.61	92.62
WND	14,221	1,933	1.25	1.34	1.27	1.02	3.38
WRAL	13,543	290	-0.43	-0.53	-0.35	0.64	47.51
WTOP	19,573	165	-0.41	-0.47	-0.24	0.53	68.09
Washington Examiner	17,783	4,010	0.60	0.17	0.65	1.17	10.22
Washington Post	38,331	14,577	-1.01	-1.06	-0.68	0.86	88.05
Washington Times	20,168	1,362	0.43	0.22	0.47	1.15	11.16
Wfaa	2,942	111	-0.39	-0.50	-0.25	0.52	55.28

Notes: This table shows the number of articles, the clicks-weighted average slant by all users, by users with a liberal political affinity score, and by users with a conservative political affinity score, the standard deviation of slant, and the percentage of liberal visits defined as the percentage of clicks by liberal users divided by the sum of the percentage of clicks by liberal and conservative users for the outlets included in our analysis. “SS1” denotes the URL-level Facebook Activity Dataset (which is curated by Social Science One). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure C.1: Slant Distribution within Outlets




Figure C.1: Slant Distribution within Outlets (Continued)





Figure C.1: Slant Distribution within Outlets (Continued)



Notes: This figure shows the distribution of slant within the outlets included in our analysis. The vertical blue line indicates the mean slant for all articles within a given outlet. Slant is given by the fine-tuned Gemma 3 27B's prediction as described in [Section 4.2](#). The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

D. Robustness

This appendix provides additional robustness checks for the Facebook analysis in [Section 6](#). We address three concerns: whether our conclusions depend on the set of outlets in our sample, on the construction of Facebook’s political-affinity measure, or on the fact that the URL-level Facebook Activity Dataset only includes URLs that were publicly shared on Facebook at least 100 times.

D.1 Robustness: Political Affinity Measure

In this section, we present a battery of robustness checks to address the limitations of the political affinity measure in the URL-level Facebook Activity Dataset. As discussed in [Section 3](#), the political affinity measure in the URL-level Facebook Activity Dataset has three limitations: first, it does not directly track canonical indicators of partisanship such as party registration; second, it covers only 25% of adult Facebook users from the U.S.; third, it might suffer from a circularity problem, because a small set of news outlets is included in the set of political pages used in the construction of the measure.


Our first and main robustness check addresses all three problems at once. Specifically, it calculates polarization in news consumption on Facebook using our independent dataset of individual-level browsing behavior. This dataset includes a highly accurate measure of partisanship in the form of administrative records on panel members’ party registration status. We restrict attention to the 1,613 individuals who are registered Democrats or registered Republicans.

To facilitate comparability between the two polarization indices, we re-compute our polarization index based on the URL-level Facebook Activity Dataset when including only those articles that were published in the last three months of 2019 (that is, covering only those months that are also covered in the browsing behavior dataset). Correspondingly, in the browsing panel dataset, we restrict attention to all clicks on news articles that i) appear in the URL-Level Facebook Activity Dataset (i.e. were shared on Facebook at least 100 times) and ii) for which the session-level referrer is Facebook. The browsing panel data gives rise to a polarization index that is almost identical to its equivalent based on the URL-level Facebook Activity Dataset (the browsing panel measure is slightly larger, by 3%). Thus, our central finding is robust to using a completely different dataset in which partisanship is measured directly from administrative voting records.

Our second robustness check aims to address the potential circularity in the political affinity variable from the URL-level Facebook Activity Dataset. Specifically, we recalculate our preferred polarization index estimate after excluding all articles from the 17 outlets in our database that are also used in the construction of the political affinity variable. We report the results in [Appendix Table A.6](#) and find a pattern of results virtually identical to the one described in [Section 6.1](#). Furthermore, when omitting those 17 outlets, our polarization index becomes 60% higher when the analysis is carried out at the article level rather than at the outlet level (see [Appendix Table A.7](#)).

Our third robustness check involves proxying a person’s political ideology using her demographic characteristics. Specifically, the URL-level Facebook Activity Dataset contains engagement numbers by age bracket and gender. Based on representative survey data from the ANES 2020-2022 Social Media Study, we can determine, for each age-gender bucket, the difference in the share of the demographic group’s members who identify as or lean Republican and the share who identify as or lean Democrat. Using that difference as a proxy for the ideological leaning of individuals in each age-gender bucket, we can then relate the average ideology in an age-gender

Figure D.1: Average Slant in News Diet By Demographic Ideology Proxy



Notes: This figure shows the mean slant of a given Facebook user group’s news diet (separately for views, clicks, and shares) on the vertical axis, for 12 different user groups sorted according to the group’s mean ideology proxy shown on the horizontal axis. User groups are given by gender-by-age cells, with two genders and six age groups. The ideology proxy is given by the difference in the share of the demographic group’s members who identify as or lean Republican and the share who identify as or lean Democrat (such that larger numbers mean more Republican-leaning). The shares are estimated based on representative survey data from the ANES 2020-2022 Social Media Study (2020 wave), among all 3,483 respondents who reported having a Facebook account that they used in the past month.


bucket to the average slant that individuals in that bucket are exposed to and consume on Facebook. The results are shown in [Appendix Figure D.1](#). In line with the results in previous sections, we find that, as we move towards age-gender buckets with relatively higher Republican vote shares, the average slant of the articles individuals in those buckets are exposed to, consume, and share on Facebook becomes more right-leaning according to our slant measure.

D.2 Extending Results to Additional Outlets

Our main analysis focuses on the 100 most visited U.S. outlets that cover news articles. As discussed in [Section 3](#), these top 100 outlets account for 80% of all online news consumption in the U.S. according to Comscore data ([Comscore, 2019](#)). Thus, they provide very extensive coverage. Nonetheless, it is natural to wonder as to how our polarization index estimate on Facebook would change if we included more outlets.

To explore this, [Appendix Figure D.2](#) extrapolates our results by extending the analysis to the top 200 U.S. outlets by online visits. Our extrapolation works as follows. First, we rank outlets by size (using Comscore data) and calculate the Republican–Democrat slant gap in clicks for

Figure D.2: Difference in the Average Slant of News Consumption by Party Registration Status, Extending Beyond Top 100 Outlets



Notes: This figure presents the clicks-weighted average difference in the slant of news diets by registered Republican and Democratic users on Facebook by extending our analysis beyond top 100 U.S. news outlets as discussed in [Appendix D.2](#). The news article sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook).

ever-increasing subsets of these outlets.⁵² Based on the trend observed, we assume a logarithmic relationship between the number of outlets included in the subset and the slant gap. The logarithmic extrapolation suggests that expanding to 200 outlets would increase the slant gap by 6%; because the polarization index is the slant gap divided by five, the index increases by the same proportion.

D.3 Robustness: Threshold in the URL-Level Facebook Activity Dataset

As discussed in [Section 3.2](#), a limitation of the URL-level Facebook Activity Dataset is that it only includes articles shared publicly on Facebook at least 100 times (after adding Laplace-5 noise to the share counts). The polarization index derived in [Section 6.1](#) is therefore based on the subset of news articles that clear this threshold. Naturally, polarization in news exposure and consumption on social media might differ if one were to consider all articles consumed through Facebook, including those shared less than 100 times. In this subsection, we address this concern in two complementary ways.

First, we recompute the browsing-panel-based polarization index using all Facebook-referred

⁵²We first only consider the most visited U.S. outlets, and then progressively expand the subset to less visited outlets.

news articles, rather than only those that also appear in the URL-level Facebook Activity Dataset. The resulting estimate is only 6% smaller than the baseline one. Relaxing the threshold thus decreases the polarization estimate only marginally. Two observations help explain this. First, articles shared widely on Facebook account for the large majority of Facebook news consumption: in the browsing data, articles that appear in the URL-level Facebook Activity Dataset account for 84% of all news article clicks through Facebook. Second, the polarization estimate is slightly *smaller* when including less widely shared articles. This is consistent with the strong pro-attitudinal sharing patterns we document in [Section 6.2.3](#): conservatives disproportionately share articles with an extreme conservative slant, and liberals disproportionately share articles with an extreme liberal slant. Articles that clear the 100-share threshold are therefore selected in part through partisan amplification, which makes the widely shared subset of articles more conducive to polarized consumption than the universe of all Facebook-referred news. Therefore, the inclusion threshold leads our main analysis to modestly overstate polarization relative to all Facebook news consumption—but the magnitude of this difference is small.

Second, we perform an extrapolation exercise using only the URL-level Facebook Activity Dataset, described in the following paragraphs.

The goal of this exercise is to assess how our polarization index would change if the dataset included not only articles shared at least 100 times, but all articles shared on Facebook at least once. We proceed as follows. We examine how the Republican–Democrat slant gap in clicks—the numerator of our polarization index—changes as we progressively add the least-shared articles in the URL-level Facebook Activity Dataset. Starting from the sample that excludes articles in the bottom 50% of total shares in the URL-level Facebook Activity Dataset, we lower the exclusion threshold one percentile at a time until no articles are excluded, recomputing the slant gap at each step. [Appendix Figure D.3](#) presents the resulting pattern: the gap computed on more widely shared articles is generally somewhat higher, but the differences are small. For example, excluding articles in the bottom 20% of total shares increases the gap by only 0.002 slant points, or 0.3% of the full-sample gap. Because the polarization index equals the slant gap divided by five, the index is correspondingly insensitive to the inclusion threshold.

In the next step, we impose assumptions that allow us to extrapolate this pattern and estimate how the slant gap would change if the dataset included additional articles. We assume that (i) every article in our database is shared on Facebook at least once, and (ii) starting from the 50th percentile in [Appendix Figure D.3](#), the relationship between the share of articles included and the slant gap is linear. The first assumption yields a conservative estimate of the share of articles shared fewer than 100 times; the second permits the extrapolation.

When we merge our article database to the URL-level Facebook Activity Dataset, we find a match rate of 22%: that share of the articles in our database is shared on Facebook at least 100 times (subject to Laplace-5 noise), and we assume the remaining articles are shared at least once but fewer than 100 times. Given this match rate, 100% in [Appendix Figure D.3](#) corresponds to only 22% of all articles shared at least once, so we extend the horizontal axis to 455% to cover them all.


The linear extrapolation implies a slant gap of 0.718 if the dataset included all articles shared at least once rather than only those shared at least 100 times—equivalently, a polarization index of 0.14, compared to our baseline of 0.16 ([Section 6.1](#)). This corresponds to a decrease of 8% relative to the baseline.

The intuition is straightforward: articles shared fewer than 100 times account for a small share of Facebook news consumption. Consistent with this, using a similar extrapolation for views, we

find that including all articles shared at least once would increase total views by only 26.9%, even though these articles represent roughly 78% of the articles in our database.

Taken together, the browsing-panel exercise and this extrapolation suggest that the 100-share threshold has only a limited effect on our polarization estimates.

Figure D.3: Difference in the Average Slant of News Consumption by Party Registration Status, Extending to Articles Shared less than 100x



Notes: This figure shows the click-weighted average difference in slant of news diets by registered Republican and Democratic users on Facebook by total shares, as discussed in [Appendix D.3](#). The news article sample includes all $N = 231,821$ news articles published online by the top 100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., shared at least 100 times on Facebook).

E. Slant Ratings

E.1 Expert Label Construction

We provided our expert human raters with the following instructions. Note that raters were asked to set slant at a scale of 1 to 7, which we recentered to a scale of -3 to +3.

Our goal is to determine how biased a newspaper article is towards a left or right-wing point of view, with the Democratic Party broadly representing the left and the Republican Party broadly representing the right. We will use a 7-point scale to measure the degree of bias, ranging from extreme left to extreme right.

There are three main factors that we consider when assessing the bias of an article: language, political position, and coverage.

(1) When we look at the language of an article, we consider whether it uses words and phrases that are typically used by members of the Republican or Democratic parties. If an article uses more Republican-like language, we consider it biased towards the right. If it uses more Democratic-like language, we consider it biased towards the left.

For example, terms like “death tax” would be considered slanted towards the right, whereas “estate tax” or “inheritance tax” would be considered slanted towards the left.

(2) Next, we look at the political position presented in the article. If the article aligns with the political views of the Republican party, we consider it biased towards the right. If it aligns with the political views of the Democratic party, we consider it biased towards the left.

As an example, an article that presents an anti-abortion stance would be viewed as having a right-leaning bias, while an article that presents a pro-abortion stance would be considered to have a left-leaning bias.

(3) Finally, we consider the issues covered in the article. If an article covers issues that are important to Republican voters and ignores issues that are important to Democratic voters, we consider it biased toward the right. If it covers issues that are important to Democratic voters and ignores issues that are important to Republican voters, we consider it biased toward the left.

For example, suppose one of the political parties proposes a minimum wage bill, and the Congressional Budget Office (CBO), Congress’s official nonpartisan provider of cost and benefit estimates for legislation, publishes a report about the bill’s potential impact. Suppose the CBO’s report states that the bill could lift 1 million people out of poverty but could also lead to a reduction of 1.5 million jobs. If a news article only reports on the job loss without mentioning the poverty reduction, it could be considered to have a right-wing bias. Conversely, if a news article only reports on the reduction in poverty without mentioning the potential job loss, it could be considered to have a left-wing bias.

By considering these factors, we can assign a score to each article that reflects its degree of bias towards the left or right. To help you understand the slant rating scale, we’ve provided rankings for a few well-known politicians (we did not come up with

those rankings; we took them from political scientists Keith Poole and Howard Rosenthal). You can find the rankings in the table below.

1	Very left-wing	Alexandra Ocasio-Cortez
2	Left-wing	Cory Booker
3	Somewhat left-wing	Chuck Schumer
4	Neutral	
5	Somewhat right-wing	Mitt Romney
6	Right-wing	Marco Rubio
7	Very right-wing	Ted Cruz

E.2 Text Processing

We first clean the articles using existing text extractor libraries after manually checking articles. The articles of approximately 70 outlets are processed using ‘readability’, Mozilla’s Readability algorithm, approximately 20 outlets are processed using ‘domdistiller’, Google’s DOM Distiller (useful for sites with complex layouts), and approximately 5 outlets are processed with ‘trafilatura’, a specialized web scraping parser.

Before feeding the text into the model, we further process the text. First, we review the articles and remove additional irrelevant text using a large language model (while human experts would probably not be affected by an advertisement in the middle of the article, it could affect our model). Second, we mask the outlet name when it appears in the article text. We provide more details on each of these steps below.

Text Cleaning We use GPT-5-mini (2025-08-07) to further clean the article text. We first tokenize the text into sentences using spaCy and then ask the LLM to classify every sentence into one of three categories: keep (preserve the sentence exactly as-is), remove (delete the sentence entirely), clean (the LLM provides both the cleaned text and the specific bad text to strip). We use the following system prompt:

“You are an expert at cleaning news articles by identifying and removing boilerplate content while preserving actual news content. You analyze articles sentence by sentence and make precise decisions about what to keep, remove, or clean.

TASK: Clean news articles by removing ONLY boilerplate content and cleaning sentences that contain embedded non-news elements.

RULES:

1. Mark for REMOVAL: ONLY sentences that are clearly not actual article content, such as:

- Copyright notices (e.g. “© Provided by Fox News Network LLC”)
- Social media links (e.g. “Follow us on...”)
- Subscription prompts (e.g. “Subscribe to our newsletter”)
- Author bylines (e.g. “By [Name]” or contact info)
- Timestamps: (e.g. “Published 3 hours ago”)

- Advertisements (e.g. "\$5 can Feed a Family for a Day. Can you give?")
 - Image caption (e.g. "Image: A view of the city skyline")
 - Website artifacts from web scraping (e.g. "Buy Photo", "1 of 8", "Autoplay", single words from a navigational menu like "Home", "News", or misparsed graph data, etc.)
 - Links to related articles or external sites (e.g. "Read more at...", or headlines of other articles that don't fit with the article text)
 - DO NOT remove actual article content even if it mentions social media or websites - be CONSERVATIVE
 - DO NOT remove sentences that are part of a quotation that may appear to be boilerplate (e.g. a tag on Twitter "@user" within a quoted tweet)
 - DO NOT remove attributions of quotes/tweets/social media posts (e.g. "–John Doe, spokesperson for XYZ on May 8, 2019")
2. Mark for CLEAN: Sentences that contain legitimate news content mixed with boilerplate elements that should be removed
- News content with embedded ads or promotional text as in the list above
 - Article content with web artifacts like "Click to expand" mixed in
 - Provide the cleaned version by specifying exactly what unwanted text should be stripped out
 - Do NOT edit actual article content or details - be CONSERVATIVE
3. Mark for KEEP: All legitimate news content
- All sentences that contain actual article information, are part of quotes, or provide commentary
 - All sentences that attribute quotes/tweets/social media posts or provide necessary context (e.g., "– John Doe, spokesperson for XYZ on May 8, 2019")

IMPORTANT: Be CONSERVATIVE - when in doubt, KEEP the sentence rather than removing actual news content.

We use the following user prompt:

"TASK: Clean the following news article by analyzing each sentence to remove boilerplate content while preserving actual news content.

CONTEXT BEFORE (for reference only): [Previous sentences as plain text, if chunked]

TARGET SENTENCES (provide decisions for these): 1: [sentence text] 2: [sentence text] ...

CONTEXT AFTER (for reference only): [Following sentences as plain text, if chunked]

CRITICAL: CONTEXT sections are for reference only - provide decisions ONLY for the numbered TARGET sentences.

IMPORTANT: You must provide a decision for EVERY sentence (1 to N INCLUSIVE). Each sentence must be categorized as "keep", "remove", or "clean".

Return a JSON object with a "decisions" array where each element corresponds to a sentence. Each element must include:

- action: "keep", "remove", or "clean"
- cleaned_text: Required for "clean" actions - the sentence with unwanted text stripped out
- bad_text: Required for "clean" actions - the specific, exact unwanted text that should be removed
- reason: Clear explanation for your decision

Example format:

```
{
  "decisions": [
    {
      "action": "keep",
      "reason": "Reasoning for keeping the sentence because it contains news information"
    },
    {
      "action": "clean",
      "cleaned_text": "The mayor announced the new policy.",
      "bad_text": "Click here to subscribe",
      "reason": "Reasoning for removal of portion of sentence that is boilerplate"
    },
    {
      "action": "remove",
      "reason": "Reasoning for removal of boilerplate sentence"
    }
  ]
}
```

““

Masking outlet names After cleaning, we replace publisher names with “[Outlet Name]”⁵³ to anonymize the text. We use the same model (gpt-5-mini, 2025-08-07). The cleaned sentences from the previous step are sent to the LLM in chunks of 20 sentences. The LLM identifies every mention of the article’s publisher (exact name, domain, abbreviations, colloquialisms, self-references), with location-based publisher names (e.g., “boston.com”) distinguished from geographic references to avoid false positives. We use the following system prompt

““ You are an expert at identifying publisher name mentions in news articles for anonymization. Your Goal: Identify every instance where the article’s publisher is mentioned so it can be replaced with “[Outlet Name]”.

MATCHING RULES:

1. Exact Matches: Look for the full publisher name (e.g., “The New York Times”).

⁵³Article titles (headlines) are processed separately from body text using a regex-based approach. Titles are short, single-line strings where publisher self-references follow predictable patterns (e.g., appended domain names, brand abbreviations). Regex patterns are sufficient for this narrower task.

2. Domains: Look for the domain name (e.g., "nytimes.com").
3. Variations: Look for common abbreviations, acronyms, or colloquialisms (e.g., "NYT", "The Times", "WaPo", "Fox").
4. Self-References: Be alert for phrases like "reporting by [Publisher]" or "according to [Publisher]".
5. Location-Based Publisher Names: If the publisher name is location-based (e.g., "boston.com", "Miami Herald"), distinguish between references to the PUBLISHER and general references to the LOCATION, only replacing when clearly referring to the publisher and not the geographic location itself.

OUTPUT REQUIREMENTS:

- Analyze EVERY sentence in the provided list.
- For each sentence, determine if it contains a mention ("replace") or not ("keep").
- If "replace", list the EXACT text strings of the mentions found.

““

We use the following user prompt:

““TARGET PUBLISHER: "[publisher_domain]", or its readable name "[Publisher Readable Name]"

SENTENCES TO ANALYZE: 1: [sentence text] 2: [sentence text] ...

INSTRUCTIONS: Provide a JSON response with a decision for every sentence (1 to N).

Example JSON output:

```
{
  "decisions": [
    {
      "action": "keep", \
      "publisher\_mentions": [],
      "reason": "No publisher name found"
    },
    {
      "action": "replace",
      "publisher\_mentions": ["Publisher Name", "Variation"],
      "reason": "Found publisher mention"
    }
  ]
}
```

““

E.3 Predicting Slant

We predict slant using the *Gemma 3 27B* open-weights model released by Google DeepMind (gemma-3-27b-it). To fine-tune Gemma efficiently, we use *LoRA* (Low-Rank Adaptation). LoRA is a lightweight fine-tuning method in which we *freeze* the original model weights and instead train a small number of additional low-rank matrices (“adapters”) that are inserted into selected layers (e.g., attention and MLP layers). Intuitively, LoRA learns a low-dimensional update to the model’s behavior without modifying the base model itself.

We use the following model parameters:

- Quantization: 4-bit NF4, double quantization, bfloat16 compute via bitsandbytes
- LoRA rank: $r = 32$
- LoRA alpha: 64
- LoRA dropout: 0.05
- LoRA target modules: q_proj, k_proj, v_proj, o_proj, gate_proj, up_proj, down_proj
- Task type: CAUSAL_LM
- Bias: none
- Epochs: 2
- Learning rate: $2e-4$, cosine schedule
- Warmup ratio: 0.03
- Batch size: 4, gradient accumulation 4 (effective batch = 16)
- Optimizer: adamw_torch_fused
- bf16: True
- Gradient checkpointing: True
- Max grad norm: 0.3

F. Correlation Ceiling Between Model Predictions and Expert Labels

This appendix shows how to derive an upper bound on the maximum achievable correlation between the predictions of our model and our expert labels. This appendix underlies the claim in [Section 4.3](#) that our model attains at least 84% of the correlation ceiling between model predictions and expert labels.

Consider the model from the [Appendix G](#). For simplicity, let $u = 0$. This assumption is without loss of generality because one can in principle always subtract the average slant across all articles from an article-level slant measure. Furthermore, let $Z_{i,j} = O_i + A_{i,j}$, so that we don’t have to separately carry around both O_i and $A_{i,j}$. Notice $E(Z_{i,j}) = E(O_i + A_{i,j}) = E(A_{i,j}) = E(E(A_{i,j}|O_i)) = 0$. Therefore,

$$R_{i,j,k} = Z_{i,j} + \varepsilon_{i,j,k}.$$

We impose the additional assumption of equal variance in errors across the two raters:

$$\text{Var}(\varepsilon_{i,j,1}|Z_{i,j}) = \text{Var}(\varepsilon_{i,j,2}|Z_{i,j}) = \text{Var}(\varepsilon_{i,j}|Z_{i,j}).$$

Let ρ_0 denote the Pearson correlation between the ratings of our two expert raters. In light of our assumptions, we can write ρ_0 as

$$\rho_0 = \frac{\text{Var}(Z_{i,j})}{\text{Var}(Z_{i,j}) + \text{Var}(\varepsilon_{i,j})},$$

which, rearranging, allows us to write

$$\text{Var}(\varepsilon_{i,j}) = \frac{1 - \rho_0}{\rho_0} \text{Var}(Z_{i,j}).$$

Let ρ_1 denote the Pearson correlation between our model predictions F and the average rating of the two expert raters.

We can write ρ_1 as

$$\rho_1 = \frac{\text{Cov}\left(F, \frac{R_{i,j,1} + R_{i,j,2}}{2}\right)}{\sqrt{\text{Var}(F) * \text{Var}\left(\frac{R_{i,j,1} + R_{i,j,2}}{2}\right)}},$$

where

$$\text{Var}\left(\frac{R_{i,j,1} + R_{i,j,2}}{2}\right) = \text{Var}(Z_{i,j}) + \frac{1}{2} \text{Var}(\varepsilon_{i,j}).$$

But then, using the expression for $\text{Var}(\varepsilon_{i,j})$ obtained above, we have

$$\text{Var}\left(\frac{R_{i,j,1} + R_{i,j,2}}{2}\right) = \text{Var}(Z_{i,j}) + \frac{1}{2} \frac{1 - \rho_0}{\rho_0} \text{Var}(Z_{i,j}) = \frac{1 + \rho_0}{2\rho_0} \text{Var}(Z_{i,j}).$$

Furthermore, it can be shown that

$$\text{Cov}\left(F, \frac{R_{i,j,1} + R_{i,j,2}}{2}\right) = \text{Cov}(F, Z_{i,j}).$$

In light of the above, we can write ρ_1 as

$$\rho_1 = \frac{\text{Cov}(F, Z_{i,j})}{\sqrt{\text{Var}(F) * \frac{1 + \rho_0}{2\rho_0} \text{Var}(Z_{i,j})}}.$$

By the Cauchy-Schwarz inequality, we have:

$$\text{Cov}(F, Z_{i,j}) \leq \sqrt{\text{Var}(F) * \text{Var}(Z_{i,j})}.$$

Therefore, we can obtain the following upper bound for ρ_1 :

$$\rho_1 = \frac{\text{Cov}(F, Z_{i,j})}{\sqrt{\text{Var}(F) * \frac{1+\rho_0}{2\rho_0} \text{Var}(Z_{i,j})}} \leq \frac{\sqrt{\text{Var}(F) * \text{Var}(Z_{i,j})}}{\sqrt{\text{Var}(F) * \frac{1+\rho_0}{2\rho_0} \text{Var}(Z_{i,j})}} = \sqrt{\frac{2\rho_0}{1+\rho_0}}.$$

The actual Pearson correlation between the ratings of the two expert raters is $\rho_0 = 0.72$. Therefore, an upper bound for ρ_1 is:

$$\rho_1 \leq \sqrt{\frac{2 * 0.72}{1 + 0.72}} \approx 0.91.$$

The actual correlation between the model prediction and our slant rating (which is simply the average of the two raters' ratings) is 0.76.

Therefore, our model achieves at least $\frac{0.76}{0.91} = 0.84$, that is 84%, of the correlation ceiling between our model prediction and our slant labels.

G. Variance Decomposition

This appendix formalizes the variance decomposition discussed in [Section 5.1](#). Its empirical implementation supplements the original week-stratified labeling sample with 800 additional randomly selected articles from outlets that had fewer than 30 rated articles, with the goal of obtaining at least 30 rated articles per outlet. These additional ratings are included in this decomposition, which is restricted to outlets with at least 30 rated articles; they are not used to train the slant model or to calculate its holdout-validation statistics.

Let $R_{i,j,k}$ denote the rating that rater $k \in \{1, 2\}$ gives to article $j \in J$ coming from outlet $i \in I$. We model $R_{i,j,k}$ as

$$R_{i,j,k} = u + O_i + A_{i,j} + \varepsilon_{i,j,k}$$

where u is a constant that captures the average slant across all articles, O_i is a random variable that captures the deviation between the slant of the average article in outlet i and the average slant across all articles, $A_{i,j}$ is a random variable that captures the deviation between the slant of article j within outlet i and the slant of the average article in outlet i , and $\varepsilon_{i,j,k}$ captures the noise that comes from a rater's imperfect measurement. We assume $E(\varepsilon_{i,j,k}|O_i, A_{i,j}) = 0$, $E(A_{i,j}|O_i) = 0$, and $E(O_i) = 0$. Note that, as a consequence, $\text{Var}(\varepsilon_{i,j,k}|O_i, A_{i,j}) = E(\varepsilon_{i,j,k}^2|O_i, A_{i,j})$, $\text{Var}(A_{i,j}|O_i) = E(A_{i,j}^2|O_i)$, and $\text{Var}(O_i) = E(O_i^2)$.

In light of the way in which we modeled $(R_{i,j,k})$, we can write $\text{Var}(R_{i,j,k})$ as follows:

$$\text{Var}(R_{i,j,k}) = \text{Var}(O_i) + E(\text{Var}(A_{i,j}|O_i)) + E(\text{Var}(\varepsilon_{i,j,k}|O_i))$$

Similarly, we can write $E(\text{Var}(R_{i,j,k}|O_i))$ as:

$$E(\text{Var}(R_{i,j,k}|O_i)) = E(\text{Var}(A_{i,j}|O_i)) + E(\text{Var}(\varepsilon_{i,j,k}|O_i))$$

Note that, if we were to perform a naive variance decomposition that does not take into account

measurement error, we would obtain:

$$\begin{aligned} & \frac{\text{Var}(R_{i,j,k}) - E(\text{Var}(R_{i,j,k}|O_i))}{\text{Var}(R_{i,j,k})} = \\ & = \frac{\text{Var}(O_i)}{\text{Var}(O_i) + E(\text{Var}(A_{i,j}|O_i)) + E(\text{Var}(\varepsilon_{i,j,k}|O_i))} \end{aligned}$$

Thus, measurement error mechanically reduces the fraction of the variance in ratings that can be explained by different outlets differing in terms of their average slant.

To eliminate measurement error in the variance decomposition, we consider

$$\text{Var}(R_{i,j,1}) + \text{Var}(R_{i,j,2}) - E\left((R_{i,j,1} - R_{i,j,2})^2\right) = 2[\text{Var}(O_i) + E(\text{Var}(A_{i,j}|O_i))]$$

Therefore, an expression that gives us the fraction of the variance in ratings (absent measurement error) that is explained by different outlets differing in terms of their average slant is:

$$\begin{aligned} & \frac{\sum_{k=1}^2 [\text{Var}(R_{i,j,k}) - E(\text{Var}(R_{i,j,k}|O_i))]}{\text{Var}(R_{i,j,1}) + \text{Var}(R_{i,j,2}) - E\left((R_{i,j,1} - R_{i,j,2})^2\right)} = \\ & = \frac{\text{Var}(O_i)}{\text{Var}(O_i) + E(\text{Var}(A_{i,j}|O_i))} \end{aligned}$$

Strictly speaking, the additive-noise representation is best interpreted as applying to an underlying continuous latent slant variable, with the observed 7-point ratings viewed as a discretization of that latent variable.

H. Classifying Other News Characteristics

This appendix describes how we classified news articles into opinion- vs. non-opinion pieces, into local, national, and international news, and into news topics. To do so, we used GPT-4o-mini. We supplied the whole article text as a user message and submitted the following prompt:

“You are a helpful assistant. I have provided you with a news article above. Its title is: [title]. It was written on [date]. It was published at the following url: [URL].

We want to decide if this article is an opinion piece or not. We define opinion pieces the following way: An opinion piece is an article that mainly reflects the author’s opinion about a subject.

We also want to determine the topic of the article. Each article belongs to one of the following categories:

- Budget and Shutdown: Issues related to the government’s budget, including disagreements over spending priorities that can lead to temporary shutdowns of government operations and services.

- **Racial Relations:** Discussions around race, ethnicity, and social dynamics related to race, including perspectives on topics like systemic racism, affirmative action, identity politics, equal opportunity, law and order, discrimination, colorblind approaches, and the role of government and society in addressing racial issues.
- **Crime, Law, and Justice:** The establishment and/or statement of the rules of behavior in society, the enforcement of these rules, breaches of the rules, the punishment of offenders and the organizations and bodies involved in these activities.
- **Conflict, War, and Peace:** acts of politically motivated protest or violence, military activities, geopolitical conflicts, as well as resolution efforts.
- **Democratic Scandal:** Allegations or controversies involving Democratic politicians or party members, covering potential misconduct, corruption, or political missteps.
- **Disaster, Accident and Emergency Incident:** Coverage of natural or human-made events that result in loss of life, injury, or property damage, such as natural disasters, industrial accidents, and emergency responses.
- **Drugs:** All matters related to legal and illegal drug use, drug trafficking, addiction, and drug policy, including debates over legalization, regulation, and public health impacts.
- **Economy, Business, and Finance:** Issues related to economic trends, business practices, market dynamics, financial markets, and related policies, including analyses of economic indicators, corporate news, investment strategies, stock market updates, fiscal policies, and global economic developments
- **Education:** All aspects of furthering knowledge, formally or informally through schools, universities, etc.
- **Election:** Matters concerning elections, including campaigns, voter turnout, polling, electoral procedures, and debates among candidates for public office.
- **Environment:** All aspects of protection, damage, and condition of the ecosystem of the planet Earth and its surroundings, including climate change and debates over policies aimed at protecting the environment versus economic growth.
- **LGBTQ:** Issues related to the legal status, rights, and social acceptance of LGBTQ individuals. Includes discussions on marriage equality, anti-discrimination protections, LGBTQ and religion, transgender issues, education on gender and sexuality, and debates over the role of LGBTQ issues in public policy and cultural norms.
- **Gun Rights/Regulation:** Matters related to the ownership, use, and regulation of firearms, as well as debates over gun control measures, firearm regulations, gun rights and the Second Amendment, background checks, concealed carry laws, and gun-related public safety concerns.
- **Health and Healthcare System:** All aspects of physical and mental well-being, as well as the healthcare system, pharmaceutical companies, etc.

- **Housing and Urban Development:** Issues related to housing affordability, availability, urban planning, and community development. Encompasses debates over rent control, zoning laws, the role of government in addressing housing needs, market-driven solutions, homelessness, and the effects of housing policies on economic growth and social dynamics.
- **Immigration:** Matters concerning immigration policy, border control, treatment of immigrants and refugees, legal status, and debates over immigration and national identity, the economic impact of immigrants, and cultural integration.
- **Labor and Employment:** Topics related to labor rights, workplace conditions, wages, employment law, minimum wage policy, union activities, labor policy and business interests/growth, strikes, and workforce trends, including job creation and unemployment.
- **National Security:** Issues related to the protection of the nation's borders, counterterrorism, intelligence activities, and measures to ensure the safety of citizens.
- **Policing Practices and Use of Force:** Coverage of incidents involving the use of force by law enforcement and discussions around policing practices. Includes debates on police accountability, law enforcement reform, public safety, community policing, "law and order" perspectives, and support for police officers as well as discussions on systemic challenges in policing.
- **Republican Scandal:** Allegations or controversies involving Republican politicians or party members, covering potential misconduct, corruption, or political missteps.
- **Science and Innovation:** Coverage of scientific research, discoveries, and technological advancements that are not necessarily commercial, including space exploration, health science, climate research, and general progress in knowledge and innovation.
- **Social Policy and Welfare:** Issues related to government social welfare programs and policies, such as food assistance, public healthcare, unemployment benefits, and debates on the social and economic impacts of welfare policies.
- **Technology:** All aspects pertaining to new products, commercial and non-commercial, that result from the application of recently gained scientific knowledge.
- **Gender:** Topics related to gender issues, including debates on women's rights, men's rights, workplace equality, reproductive rights, family policies, discrimination, and gender roles.
- **Trade:** Topics concerning international trade policies, tariffs, trade agreements, and economic relationships between countries, as well as their impacts on businesses and consumers.
- **Other:** Articles or topics that do not fit neatly into any of the predefined categories but are still of public interest or relevance to the period being analyzed.

We also want to know if the article is local, national or international politics.

- Local politics: Issues and developments related to U.S. state, county, city, or municipal government policies and political dynamics.
- National politics: Issues and developments related to U.S. federal government policies, political parties, and national leadership. Encompasses discussions on federal legislation, executive actions, Supreme Court decisions, national party debates, national public policy issues, etc.
- International politics: Matters concerning diplomatic relations and policy decisions between the U.S. and other countries, as well as political dynamics and developments in countries/regions other than the U.S.

Read the text of the article carefully and determine if it is an opinion piece (yes or no), its topic and whether it is local, national or international news.

Always add your chain of thought in the reasoning field.”

I. Slant Distance: Example Articles

To illustrate the magnitude of our polarization measure, this appendix reproduces excerpts from three articles published in April 2019 about the same event—an exchange regarding the border wall between President Donald Trump and a child attending the White House Easter Egg Roll. The slant distance between the liberal and moderate articles (1) and between the moderate and conservative articles (1) is comparable to our polarization measure. We do not reproduce the full text for copyright reasons.

Liberal (slant = −1): President Trump Bragged About the Border Wall to a Group of Children at the Easter Egg Roll — *Time*, April 22, 2019

President Donald Trump used the opportunity of the White House Easter Egg Roll to address some of his youngest supporters about one of his signature policy initiatives: building a wall on the southern U.S. border.

“Oh it’s happening, it’s being built now,” Trump told a group of children as he colored pictures with them on Monday in celebration of Easter on the White House’s South Lawn.

“Here’s a young guy just said ‘keep building that wall.’ Do you believe it? He’s going to be a conservative someday,” the president said.

Trump continues to push controversial policies to halt migration from Central America. The administration’s botched efforts have ranged from separating migrant children from their families to ultimatums against Mexico that have proved too costly to enforce.

[...]

Moderate (slant = 0): Trump talks up border wall with small child at Easter Egg roll: 'He's going to be a conservative some day' — Raw Story, April 22, 2019

President Donald Trump on Monday promised a child at the White House Easter event that a wall on the U.S.-Mexico borders is being “built now.”

According to a White House pool report, the boy asked Trump about his plans for the wall.

“I will. Oh It’s happening. It’s being built now,” Trump told the child. “Here’s a young guy who said ‘Keep building that wall can you believe that?’ He’s going to be a conservative some day!”

[...]

Conservative (slant = 1): Trump plugs border wall in exchange with young Easter egg roll attendee — The Hill, April 22, 2019

[...]

The president joined children to color cards for members of the military when he said a young man told him, “Keep building that wall.”

“It’s happening. It’s being built now,” Trump said before looking up to other attendees at the annual White House Easter celebration.

“Can you believe that?” Trump said. “He’s going to be a conservative some day.”

After Congress did not appropriate the \$5.7 billion Trump had requested for the wall in its latest round of funding, the president in February declared a national emergency at the border to secure money for his project.

[...]

“Our country is doing fantastically well,” Trump said. “Probably the best it’s ever done economically. We’re setting records on stock markets. We’re setting records with jobs.”

“Regulations, low taxes, our country has never done better,” he continued. “And do we love our military? Our military is literally being completely rebuilt. We are completely rebuilding our military. It was very depleted, as you know.”