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

February 2025

Abstract

We fine-tune a large language model to assign a content-based measure of slant to the near universe of hard-news articles published online by the top 100 U.S. outlets in 2019. We find that 67% of the article-level variation in slant is within rather than across outlets, implying substantial scope for ideological selection within outlets. Turning to social media, we find that polarization in news consumption, defined as the difference in the average slant of articles consumed by Republicans and Democrats, is 50% higher when slant is measured at the article level rather than proxied using outlet-level averages.

*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: ruiqizhong@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 (Newman et al., 2023). As of 2023, around 50% of U.S. adults reported regularly consuming news through social media, with Facebook being the dominant social media platform for news access (Pew Research Center, 2023*a*; Newman et al., 2023).

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). Such seeming lack of evidence might partly 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.[2] This approach does not account for the possibility that the social media curation process exposes users to ideologically congenial content even *within* outlets. If the social media ecosystem promotes the consumption of pro-attitudinal news within outlets, measuring slant at the outlet level will underestimate, perhaps severely, the degree of polarization of news consumption on social media.

This paper addresses the limitation above by means of a two-pronged approach. First, we pro-

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

[2] González-Bailón et al. (2023) and Green et al. (2023) are notable exceptions. We discuss our contribution relative to them towards the end of the introduction.

pose and implement a novel content-based method to assign slant to individual news articles rather than to entire news outlets. Second, we apply our article-level slant measure to estimate the degree of polarization in news consumption on Facebook.

Our analysis relies on three main datasets. The first dataset encompasses the near universe of hard news articles ($\sim$ 1 million) published online by the top 100 U.S. outlets in terms of online visits in 2019. 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 data on which to fine-tune the model, we hired two expert raters (as well as a research assistant), gave them a sample of more than 4,400 articles randomly drawn from our database of hard news articles, and asked them to assign to each article a measures of slant on a scale from -3 (very left-wing) to 3 (very right-wing). Fine-tuning the LLM substantially increases its performance, from a prediction accuracy of 0.72 to a prediction accuracy of 0.83.

An array of validation exercises shows that our machine-learning-based measure of article-level slant is well-calibrated. 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 0.89. 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 0.82.

Our first set of results focuses on the production of slant by the top 100 U.S. outlets in terms of online visits in 2019. Analyzing the landscape of slant on the production side helps contextualize our main result about polarization in news consumption and provides a useful benchmark against which to measure the distribution of slant on social media. To obtain a quantitative measure of the importance of taking into account article-level slant, we perform a variance decomposition estimating the degree to which variation in slant at the article level can be accounted for by outlet

provenance.[3] We find that the vast majority ($\sim$ 67%) of the variance in slant at the article level is within outlets rather than across outlets. This finding highlights the importance of measuring slant at the article level and the potential for social media to expose users to ideologically congenial news even within outlets.

Next, we analyze the distribution of slant on the production side. Based on the near-universe of articles published by the top 100 outlets in 2019, we find that the overall slant distribution is unimodal and centered around moderate. However, there is substantial variation in slant by article type and topic. For example we find that opinion article have a bimodal distribution and that articles about welfare, gender, LGBTQ issues, gun rights, and racial relations are more slanted than articles about international conflicts, natural disasters, drugs, and trade.

Our second set of results focuses on the circulation and consumption of slanted articles on social media. For our main result, we introduce and estimate a new measure of polarization in news consumption. We define polarization in news consumption as a scaled difference between the average slant of the articles that Republicans and Democrats consume on Facebook. We find that the degree of polarization in news consumption on social media is substantial: the distance on our scale between the average slant that Republicans and Democrats consume on Facebook corresponds to the distance in slant between the average article from the Washington Times and the average article from the Washington Post.[4] We also find that the news consumed through social media is substantially more polarized than the online news accessed through other means; relative to news accessed through search engines, polarization is 2.4 times as high, a gap equivalent to a 0.33 standard-deviation difference in relative slant.

To provide an additional benchmark for our polarization result, we calculate a version of our measure where, following the approach employed by much of the literature, we proxy the slant of each article by the slant of the outlet it came from. This way, we mechanically shut down the channel of pro-attitudinal news consumption within outlets. We find that the degree of polarization in news consumption is 47% higher when we take into account the possibility of pro-attitudinal news consumption within outlets—as per our original measure of polarization—than when we

[3] Our variance decomposition is designed to be robust to the fact that we measure slant at the article level with a degree of noise.

[4] Where the Washington Post and the Washington Times rank at the 11^{th} and 82^{nd} percentile of the overall distribution in average slant across the outlets in our dataset.

shut down that channel by proxying the slant of an article by that of the outlet it comes from. Thus, employing measures of slant at the outlet rather than the article level leads to a substantial underestimation of the degree of polarization in news consumption on social media.

We highlight three additional important drivers of polarized news consumption on social media. First, we show that extreme articles are much more likely to be shared widely on Facebook than moderate articles. As a result, the distribution of articles that are widely shared on Facebook has much thicker tails than the distribution of articles produced. Second, we examine whether polarization in news consumption is primarily driven by the articles that Democrats and Republicans are exposed to on their feeds or by the articles that they select among the articles they are exposed to. We find that 82% of our measure of polarization can be explained by the articles that Democrats and Republicans are exposed to on their feeds, thus highlighting the role of the feed-curation process in driving polarization in news consumption on social media. Third, in an attempt to further unpack the role of the feed-curation process, we explore the degree to which echo chambers—which arise from the interaction between the politically homophilous structure of the Facebook network and pro-attitudinal sharing patterns—can explain polarization in news consumption on Facebook. A back-of-the-envelope calculation finds results consistent with echo chambers being a substantial, although not exclusive, driver of polarization in news exposure on social media. Specifically, we find that if users' feeds simply reflected the articles shared by their friends, the degree of polarization in news exposure on Facebook would be around 59% of what it currently is.

This project contributes to the literature trying to measure 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., 2023). To help readers navigate this crowded literature, Appendix Table A.1 provides a schematic overview of the research articles most closely related to our paper. 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, we believe that our measure of *polarization* in news consumption is an essential complement to other measures employed in the literature, such as measures of segregation and favorability (Gentzkow and Shapiro,

2011; González-Bailón et al., 2023).[5] Specifically, an important hypothesis that researchers and policymakers are interested in assessing is whether Democrats and Republicans consume articles with different slant—as captured by our measure of polarization—as opposed to whether they simply consume different articles—which would be captured by measures of segregation or favorability.[6] The conceptual distinction between the polarization measure we introduce and measures of segregation and favorability can be readily seen in the following thought experiment. Consider a world in which: i) people only follow local politics, ii) individuals are geographically segregated by political ideology, and iii) the articles discussing local politics are all exactly moderate. In that world, our measure of polarization would equal zero, whereas measures of segregation would be at their maximum.[7]

Second, we believe that, for studying polarization of news consumption in the digital age, it is particularly important to develop a *content-based* measure of slant at the *article level*. The existing alternatives are mainly audience-based measures of slant at the article level (as in Bakshy, Messing and Adamic, 2015; González-Bailón et al., 2023; Green et al., 2023), content-based measures of slant at the outlet level (as in Groseclose and Milyo, 2005; Gentzkow and Shapiro, 2010), or audience-based measures of slant at the outlet level (as in Flaxman, Goel and Rao, 2016; Nyhan et al., 2023).[8] The importance of developing a content-based measure of slant rather than an audience-based measure of slant follows from the observation in the previous paragraph that the differential consumption and sharing patterns across Democrats and Republicans captured by audience-based measures of slant might, in principle, have little to do with slant. The importance of developing a content-based measure of slant at the article level rather than at the outlet level follows from the observation that the curation process on social media occurs at the article level and that

[5] See Appendix G for definitions of measures of segregation and favorability, and for a detailed discussion of the difference between those measures and the one we develop and employ in this paper.

[6] Of course, whether Democrats and Republicans consume the same or different articles is itself an important, albeit separate, policy-relevant question.

[7] As another example, consider outlets that source their content from news agencies like the Associated Press (AP). Suppose the same article from a news agency is published by both a liberal and a conservative outlet. Under our content-based measure of slant, the two identical articles would, by definition, receive the same slant score regardless of the outlet of provenance. As a result, our measure of polarization would remain unchanged even if Democrats predominantly read the article in the liberal outlet and Republicans predominantly read the article in the conservative outlet. In contrast, measures of segregation would increase if Democrats accessed the article through the liberal outlet and Republicans accessed the same article through the conservative outlet. Section G.1 shows that this phenomenon does indeed occur in the data.

[8] Audience-based measures of slant are those that rely on differential consumption or sharing patterns by Democrats and Republicans.

individuals might consume pro-attitudinal articles even within outlets. A relatively small literature does attempt to assign a content-based measure of slant to individual news articles (Boxell and Conway, 2022; Budak, Goel and Rao, 2016; Garz et al., 2020; Green et al., 2023; Peterson, Goel and Iyengar, 2021). Such papers, however, do not deploy their content-based article-level measure of slant at scale and/or focus on topics (e.g., the role of journalists in the production of slant) that are largely orthogonal to the ones considered in this paper.

Overall, the content-based measure of slant at the article level that we introduce in this paper allows us to make the following qualitative contributions to the literature. First, we show for the first time that most of the variance in article-level slant occurs within rather than across outlets. Second, we introduce and estimate a novel measure of polarization in news consumption on social media, which captures the degree to which Democrats and Republicans consume articles with different slant. Contrary to much of the literature, we find that polarization in news consumption on social media is substantial.[9] Third, we provide unique evidence on the mechanisms that produce polarized news consumption on social media. Specifically, we compare the slant landscape on the production side to that of articles that circulate widely on social media, highlighting special features of the social media environment—e.g., pro-attitudinal sharing patterns and network homophily—that drive polarization in news consumption.

The rest of this paper is organized as follows: Section 2 provides some background information about the prevalence of news consumption on social media and about the curation process on social media platforms; Section 3 introduces the datasets we use for our analysis; Section 4 describes our methodology to assign slant at the article level and a battery of validation exercises; Section 5 describes the slant landscape of the news produced by the top 100 U.S. outlets in 2019; Section 6 presents our main result about polarization of news consumption on social media and describes the mechanisms driving it; Section 6.4 presents extensions and robustness analyses; Section 7 concludes.

[9]This finding dovetails 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.

2. Background

As of 2023, around 50% of U.S. adults reported consuming news on social media "often" or "sometimes" as opposed to "rarely" or "never" (Pew Research Center, 2023*a*). Furthermore, individuals aged 19 to 29 reported preferring to consume news from social media more than from any other medium and were more likely to consume news on social media compared to visiting news sites or apps (Pew Research Center, 2023*b*).

Due to a combination of its sheer size ($\sim$ 190 million monthly active users in the U.S.) and the fact that many of its users get news on the platform, the dominant platform for news consumption on social media in the U.S. is Facebook (Datareportal, 2023). According to 2023 Pew Research Center data, 30% of U.S. adults report regularly getting news from Facebook, more than from any other social media platform. As a point of comparison, only 14% and 12% of U.S. adults report regularly getting news from TikTok and Twitter (now "X"), respectively (Pew Research Center, 2023*a*).

A striking difference between news consumption on social media and more traditional forms of news consumption such as newspapers is the degree of personalization of the news bundles that readers are exposed to. In particular, traditional newspapers distribute the same content to virtually all consumers; therefore, a person reading a newspaper top to bottom (e.g., the New York Times) is exposed to a bundle of news articles whose slant is curated by the news outlet. Conversely, the bundles of articles that social media users are exposed to are tailored to their individual tastes: news articles appear on users' feeds as a function of the pages that they follow on the platform, of the network of friends they are connected to, and of the platform's content-ranking algorithm.

The highly personalized process of news content curation on social media makes it so that the slant of an article that appears on a user's feed need not reflect the average slant of the news outlet that the article came from. For instance, a Democrat Facebook user might be more likely to encounter New York Times opinion columns by Paul Krugman (a relatively left-wing columnist according to our slant measure), whereas a Republican Facebook user might be more likely to encounter New York Times opinion columns by Bret Stephens (a relatively right-wing columnist according to our slant measure). For this reason, in order to study the slant of the bundle of news that Democrats and Republicans are exposed to and consume on social media, it is necessary to

assign a measure of slant to individual news articles rather than to entire outlets.

3. Data

Our analysis relies on three main datasets: one that includes the URL, text, and meta-data of a large collection of hard 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 Hard News Articles

The first dataset we employ in our analysis comprises the near universe of hard news articles ($\sim$ 1 million) published online in 2019 by the top 100 U.S. outlets in terms of online visits.[10] Our choice of focusing on the top 100 U.S. outlets is the result of a trade-off between covering as large a fraction of online news consumption in the U.S. as possible and data collection costs. According to Comscore data, visits to the top 100 U.S. outlets cover 93% of all online consumption from U.S. outlets focused on hard news, thus yielding a database of articles with very extensive coverage.

We automatically retrieved the English-language news articles published online in 2019 by the top 100 U.S. news outlets directly from the outlets' websites.[11] 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).[12] This procedure generated a dataset of 3,268,760 articles.

In order to identify and isolate the subset of hard news from the dataset of articles described in the previous paragraph, we employed a machine learning algorithm. Following the literature, we define hard news as news typically displayed in the first section of a newspaper. Thus, our definition of hard news includes both reported news articles and opinion pieces on hard-news topics; conversely, it excludes soft-news content such as sports, entertainment, weather, and similar

[10] We 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 hard news (e.g., Elle).

[11] To 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'.

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

categories. The reasons for restricting our focus to hard news are twofold: first, slant is not as well-defined for soft news such as celebrity gossip and sports; second, hard news is likely more relevant in shaping people's political behaviors like voting. To classify hard news, we relied 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 1,057,124 hard-news articles. We present the list of outlets included in our dataset, as well as the number of hard news articles published in 2019 by each outlet in Appendix Table J.1.

Using browsing data from a representative sample of 5,115 adults described in Section 3.3, we can obtain an estimate of the degree to which our dataset is comprehensive. Specifically, we can match the URLs that appear in our dataset of hard news with the hard news URLs visited by the panelists. We find that our dataset has very high coverage. Specifically, we find that only 4.6% of all hard news articles that appear in the browsing histories of the panelists and that originate from the top-100 news outlets do not find a match in our dataset. In other words, our dataset of hard news articles comprises 95.4% of all the news articles visited at least once by at least one of the 5,115 panelists in our dataset of browsing behavior.

Article cleaning The objective of our cleaning procedure was to ensure that our machine learning algorithm assessed article slant based solely on the article text and title. To clean the articles, we first tested existing packages (Domdistiller, Readability, Trafilatura) designed to extract relevant text from HTML files. We compared the differences across these packages and selected the most effective one for each outlet. Additionally, we removed extra white space and excluded paragraphs that contained only a short line of text (up to 100 characters) without a period, as such lines were typically advertisements or captions rather than part of the article body.

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.[13] The dataset provides the number of unique users who viewed, clicked, and shared a given URL on Facebook. 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).[14]

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—a mathematical framework designed to quantify and manage the privacy risk associated with the release of statistical summaries of datasets containing individual information (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. We provide summary statistics for the articles in the URL-level Facebook Activity Dataset in Appendix Table A.2.

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).[15]

[13]For 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).

[14]We 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 0.92, 0.93 and 0.91 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.

[15]This 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. We refer the reader to the codebook of the URL-level Facebook Activity Dataset for additional

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, however, has also some drawbacks. First, it is not a canonical measure of partisanship, in the sense that it does not directly track canonical indicators such as party registration. Second, it does not have universal coverage, because it is assigned only to individuals who "like" at least one political page on Facebook. 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. Below, we elaborate on each drawback in turn, as well as on our proposed solutions.

The first drawback of the political affinity measure in the URL-level Facebook Activity Dataset is that it is not a canonical measure of partisanship. Specifically, if we define partisanship status as being a registered member of the Democratic or Republican party, it is not immediately clear how the political affinity measure in the URL-level Facebook Activity dataset maps to party registration. Furthermore, since we do not have access to the underlying proprietary data that Meta uses to construct the political affinity measure, we cannot create a confusion matrix to understand the relationship between political affinity scores and being a registered member of the Democratic or Republican party. Luckily, however, the construction of the political affinity score in the URL-level Facebook Activity Dataset is almost identical to the construction of the political affinity measure developed in Barberá et al. (2015) using Twitter (now "X") data. Since the dataset used in Barberá et al. (2015) is publicly available, we can recover the confusion matrix that describes how voter registration maps to political affinity in the Twitter dataset. Assuming the confusion matrix for the Twitter and the Facebook datasets are the same, we leverage that confusion matrix to estimate the degree of polarization in news consumption that would prevail if, in the URL-level Facebook Activity Dataset, partisanship was defined according to party registration, as opposed to being defined according to the political affinity measure.

details.

Second, since the political affinity score is only assigned to individuals who like at least one news page on Facebook, the political affinity measure does not have universal coverage: it only covers around 25% of users (Messing et al., 2023). Despite being a minority of Facebook users, individuals who are assigned a political affinity score are responsible for the lion's share of engagement—72% to 87%, depending on the engagement measure—with hard news articles on Facebook, as shown in Table A.3. Thus, whenever we report results that rely directly on the political affinity score, we capture the vast majority of engagement. As described in the previous paragraph, our headline estimate of polarization is based on party registration as opposed to the political affinity score directly. Besides having the upside of tracking a more canonical definition of partisanship, basing our preferred polarization estimate on party registration also has the advantage of allowing us to take into account the news consumption behavior of users with a missing political affinity score as described in Section 6.1.2.

Lastly, the political affinity score might suffer from a small degree of circularity. Specifically, among the more than 500 political pages employed in its construction, there are 18 news outlets that overlap with the ones in our list.[16] In Section F.2, we assuage concerns about such potential circularity by providing a robustness check that drops from our analysis the set of outlets that are both in our universe and among the political pages used in the construction of the political affinity variable. Furthermore, we show that our results go through when we use an alternative ideology proxy that relies on basic demographics rather than on Facebook's political affinity measure.

Besides the proposed solutions above, the key robustness check we run to assuage potential concerns about the political affinity measure is to leverage a panel dataset of browsing behavior, described in detail below, in which individuals are assigned a measure of partisanship based directly on whether they are registered Democrats or Republicans. Specifically, by considering news visits that have Facebook as a referrer, we can calculate a measure of polarization in news consumption that is completely independent from the one calculated using the URL-level Facebook Activity Dataset. Reassuringly, as shown in Section F.2, we find that the measures of polarization calculated on our panel of browsing behavior and on the URL-level Facebook Activity Dataset are almost identical.[17]

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

[17]The advantage of measuring polarization using our panel of browsing behavior is that it relies directly on party

3.3 Panel of Browsing Behavior

For additional robustness checks and analyses, we rely on browsing panel data from a representative sample of 35,000 U.S. adults, curated by market-research firm Luth Research.[18] 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, 5,115 visited at least one hard news article published by the 100 outlets included in our study, generating a total of 121,703 clicks.

A unique feature of this dataset is that it includes party registration status for registered voters from the approximately 30 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 5,115 panelists who visit at least one hard news article, 1,637 have a Democratic or Republican party registration record, and form the sample for our browsing-panel analyses.

The Facebook URL-level dataset and the browsing panel play complementary roles in our analysis. The Facebook data provide platform-wide URL-level measures of views, clicks, and shares for the near universe of news content, which allows us to study polarization, within-outlet selection, and exposure patterns at scale. The browsing panel, in contrast, is smaller, but it provides an independent benchmark in which partisanship is observed from party-registration records rather than inferred from Facebook's political-affinity measure. Furthermore, it allows us to compare an estimate of polarization in news consumption on social media to corresponding estimates for other referral sources (e.g., search engines).

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 of more than

registration records; the disadvantage is that it is estimated on a sample that is not necessarily representative of the Facebook population.

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

4,400 articles whose slant was rated by experts. We discuss the procedure in detail below.

4.1 "Ground Truth" Expert Slant Labels

To successfully fine-tune a large language model to assign a measure of political slant to news articles, it is paramount to obtain a high-quality training dataset with accurate labels. To construct this dataset, we employed, over a period of about six months, two freelance workers with graduate degrees in political science and criminal justice, as well as a research assistant. Henceforth, we refer to the two freelance workers as "expert raters".[19]

Our two expert raters were independently asked to assign slant labels to 4,434 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.

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.[20] Raters were asked to consider all three criteria and provide a holistic assessment for each article. The full instructions are reported in Appendix B.

We chose to rely on expert raters rather than crowd-sourced ratings (e.g., via mTurk), because we believed our experts would produce higher quality ratings than mTurk workers. In fact, previous

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

[20]See, 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 (Djourelova et al., 2023) and political scandals (Puglisi and Snyder, 2011; Galvis, Snyder Jr and Song, 2016).

papers that employed mTurk workers found a correlation between the labels of different raters of only 0.26 (Peterson, Goel and Iyengar, 2021). As discussed below, the correlation between our expert raters' labels is much higher.

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 the less than 2% of articles for which 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, provided a written explanation for why he thought the two raters might disagree about the slant of the article, replaced the rating of the expert he disagreed with 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 substantial mass in the tails, as shown in 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. Considering all pairs of rated articles, the two raters assign the same relative slant ordering in 79% of cases (and in 91% of cases after the RA's correction). 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).[21] Disagreement is quantitatively small: the median and modal absolute difference between the two labels is 0, and the average absolute difference between them is 0.59 on our 7-point scale.

4.2 Machine Learning

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

[21] Cohen'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.

article in our universe. The general purpose component consists of the generative pre-trained transformer model GPT-4o by OpenAI. The task-specific component involves fine-tuning GPT-4o, 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 GPT-4o with the same set of instructions as our expert raters, as well as with 2,958 articles from our training set containing the average of our two experts' labels. The remaining 1,476 labeled articles were held out for validation. We only provided GPT-4o with each article's headline and text; we did not provide any additional information, such as the URL of the article or the outlet it came from.

Fine-tuning GPT-4o yielded substantial improvements in the model's ability to predict the average of our two experts' labels. In the hold-out portion of our training dataset, the Pearson correlation between the model's predictions and the average of our two experts' ratings is 0.72 without fine-tuning and 0.83 with fine-tuning.

4.3 Model Predictions

Our machine learning model outputted, for each article in our database, a slant prediction that ranged from -3 (indicating a very left-wing article) to 3 (indicating a very right-wing article). For all analyses other than the validation analysis below, we merge slant bucket -3 with slant bucket -2.5 and slant bucket 3 with slant bucket 2.5. This is because slant buckets -3 and 3 contained too few articles to draw meaningful conclusions, especially given the noise introduced to the URL-level Facebook Activity Dataset due to differential privacy.

For reference, Table A.2 provides some summary statistics about the mean and standard deviation of the slant of articles in our datasets, and Table A.4 lists the titles of the ten most clicked articles on Facebook in 2019 from the most extreme left-wing and right-wing slant buckets.

4.4 Validation Exercises

Our paper introduces the first large-scale content-based measure of slant at the article level built on expert labels. Since the measure is novel, it is important to carefully validate it before relying on it in our analysis. In this section, we describe the results of an array of validation 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). As mentioned above, we find a Pearson correlation of 0.83; we find a Spearman correlation of 0.77.[22] Since our "ground truth" measure of slant is itself noisy (the correlation between the two experts' labels is 0.72) we can derive an upper bound on the maximum achievable correlation between our model prediction and our slant labels. The analysis, described in detail in Appendix E, shows that this upper bound is 0.91. Therefore, our model predictions attain 91% of this correlation ceiling.

Second, we compare the absolute difference between the model prediction and the expert slant labels. The mean absolute difference (MAE) is 0.36, which corresponds to only 5% of the length of our scale.[23] 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. Furthermore, across all pairs of articles in our hold-out dataset, we find that our model prediction and the ground truth assign the same relative slant ordering in 90% of cases. Lastly, we check how often the predicted slant and the expert slant label have the same direction of slant (among all non-neutral articles). We find an agreement rate of 95%.

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 0.89.

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

[22] The R^2 from a regression of a hold-out article's model-predicted slant on its average expert label is 0.69 and the RMSE is 0.55.

[23] The model prediction and the label agree perfectly in 53% of cases; in 89% of cases, the model prediction and the label fall within 0.5 points of each other (on the scale of length six).

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 0.82 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 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.[24]

Taken together, these validation exercises solidify our confidence in the measure of article-level slant assigned to the near universe of hard news articles analyzed in this study. Furthermore, they show that the measure has a stable interpretation: it is anchored in a 7-point scale used by expert raters, the raters show substantial agreement in applying that scale, and the resulting ratings line up closely and approximately linearly with independent article-level ratings from Ad Fontes Media.

5. News Production

In order to provide useful context and to benchmark our measure of polarization in 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

[24] Both 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.

To assess the degree of variation in slant within outlets we perform a variance decomposition. The following two illustrative cases might help build intuition. At one extreme, one can imagine a world in which all articles from the same outlet have the same slant. In that world, the entirety of the variance in slant across articles would be explained by outlet provenance. At the other extreme, one can imagine a world in which all outlets have the same average slant, and articles differ in slant only within outlets. In that world, none of the variance in slant across articles would be explained by outlet provenance. Proxying the slant of an article with the average slant of the outlet it comes from is perfectly accurate in the first illustrative case and it becomes progressively less accurate the closer reality is to the second illustrative case. In other words, the smaller the fraction of variance in slant that can be accounted for by outlet provenance, the more important it is to measure slant at the article level.

We begin our analysis by illustrating 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 clearly 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 0.98 points on our scale. Despite the average difference in slant, however, there is substantial overlap in the two distributions: within each outlet, approximately 50% of 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. We present the distribution of slant for all outlets in Appendix Figure J.1. The figure demonstrates that there is substantial within-outlet variation in slant for virtually all outlets.

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 begin by noting 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. In particular, in the presence of measurement error in our slant measure, the naive variance decomposition above would necessarily underestimate the true amount of variation in article-level slant that can be accounted for by outlet provenance, as shown in Appendix D.

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 limiting our variance decomposition to our dataset of labeled articles, and we deal with the former source of noise by leveraging the fact that we have two independent labels per article and by employing a measurement-error-correction technique. The details of the variance decomposition are provided in Appendix D. Conceptually, our measurement-error-correction technique is similar to the standard procedure that estimates 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 33% of the variance in article-level slant. In other words, ignoring within-outlet variation in slant induces one to miss 67% of the overall variation in article-level slant —that is, all the variation that occurs *within outlets*.

The variance decomposition above has two direct implications for the argument developed in this paper: first, it emphasizes the importance of considering slant at the article level. Second, it shows that there are ample opportunities for ideological sorting 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.

5.2 Overall Distribution of Slant

Figure 3 presents the overall slant distribution of the ∼1 million hard news articles published online by the 100 most-visited U.S. news outlets in 2019. The distribution of slant on the production side is unimodal and centered around moderate values. Notably, 40% of all articles have a slant of zero, classifying them as centrist. Among the remaining articles, 68% are left-leaning, meaning left-leaning articles are around twice as common as right-leaning ones.[25] When comparing slant intensity, however, right-leaning articles are on average more extreme. In fact, the share of extreme right-leaning articles, defined as articles with a slant of 2 and above, among all article in our data is 7%, while the share of extreme left-learning, defined as articles with a slant of -2 and below, is only 3%.

Figure J.1 presents the distribution of slant for each outlet and shows that there is a clear asymmetry in how news outlets produce left- and right-leaning content. Only 1 outlet produces articles among which the majority is extremely left-leaning, whereas a considerably larger number of 5 outlets predominantly generate extremely right-leaning articles. These patterns dovetail with findings from the literature that point to the presence of a small segment of the news ecosystem that caters almost exclusively to conservatives (González-Bailón et al., 2023; Guess, 2021).

[25]The independent ratings by AdFontes Media described in Section 4.4 support these findings. Specifically, our model's 0-centrist designation closely matches AdFontes: for articles rated 0 by our model that also have AdFontes ratings, the average AdFontes rating is -0.11 (after we re-scale their scale to -2.5 to 2.5 with 0 unchanged), indicating these articles are indeed likely to be centrist. Moreover, among articles with AdFontes ratings, both slant measures broadly agree on the fraction that are left-leaning, moderate, and right-leaning.

5.3 Distribution of Slant by News Category and Topic

Next, we analyze how slant varies by news category and topic. This analysis compares opinion versus non-opinion pieces, local versus national versus international news, and slant differences across various topics. We use GPT-4o-mini to extract these characteristics. To evaluate the classification of the article, we compare the LLM's prediction to our expert classification, and find an agreement rate of 92% and 95% for location (local, national, or international) and whether the article is an opinion piece, respectively. autorefclassifyNewsCharacteristics provides more details on the prompt used.

Opinion vs. Non-Opinion As discussed in Section 3, our definition of hard news includes opinion pieces. The large language model classifies 22% of articles as opinion pieces.

We find that opinion pieces are significantly more slanted than non-opinion articles. Specifically, 42% of opinion articles have an absolute slant of 2 or more on our scale, compared to only 1% of non-opinion pieces. As illustrated in Figure A.3, the distribution of slant in non-opinion articles is unimodal and centered around moderate values, with relatively thin tails on both sides. In contrast, opinion pieces exhibit a strongly bimodal distribution, with peaks at both ends of the slant spectrum, indicating a tendency towards more extreme positions.

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), suggest that opinion pieces might be an important contributor to political polarization among the news-reading electorate.

National vs. Local vs. International News Figure A.4 presents clear differences across the three news categories of articles and shows that national news is substantially more slanted than international and local news. When we limit our analysis to national news, the media landscape no longer appears to be primarily moderate, as a majority of articles have an absolute slant of at least 1 on our scale. Furthermore, 18% of national news articles have an absolute slant of 2 or

more, whereas only approximately 3% of local and international news do. Interestingly, when focusing only on national news, we still observe more left-leaning than right-leaning articles and more extreme right-leaning than extreme left-leaning articles.

The discrepancy in slant across national, local, and international news can partly be explained by the fact that national news tends to cover relatively more contentious topics. For instance, 40% of international news articles cover international conflicts. Since in 2019 the U.S. was not directly involved in many of those conflicts, articles covering them tend to feature more factual and neutral coverage, as shown below. Similarly, bipartisan topics such as drug abuse are primarily covered in local news, and such coverage once again tends to be relatively neutral.

Slant across News Topics Our last piece of analysis in this section investigates the distribution of slant across news topics. Appendix C provides a list of the news topics and descriptions thereof that we used when prompting GPT-4o to classify our articles into topics.

We begin our analysis by showing the degree to which various topics are over vs. under-represented among left-leaning articles relative to right-leaning articles.[26] We consider a topic as being over- (under-) represented among left-leaning articles if, among all left- and right-leaning articles about that topic, the proportion of left-leaning articles is higher (lower) than the proportion of left-leaning articles across our entire sample of left- and right-leaning articles. Figure 4 presents the results. Perhaps unsurprisingly, the most over-represented topic among left-leaning articles is the topic of Republican scandals, and the most under-represented topic is that of Democrat scandals. This result aligns with prior evidence at the outlet level (Puglisi and Snyder, 2011) and provides further validation for our slant measure. Besides scandals involving Republicans, the most over-represented topics among left-leaning articles include natural disasters, the environment, policing practices, housing, drugs, gender, race, and LGBTQ issues. Besides scandals involving Democrats, the most over-represented topics among right-leaning articles include technology, the economy, international conflicts, welfare, trade, national security, and innovation.

A lasso regression analysis of article slant on the unigrams and bigrams in the article titles provides evidence in line with the rest of our topic analysis. The results, depicted in Figure A.5,

[26] For this analysis, we consider an article as being left-leaning if it has a slant below -0.5 and right-leaning if it has a slant above 0.5. We drop articles that are broadly moderate, namely articles with a slant between -0.5 and 0.5.

show that the most predictive unigrams for right-leaning articles are "dems", "aoc", "hillary", "democrats", "illegal", "socialism", "left", "china", and "beto", while those for left-leaning articles are "trump", "climate", "women", "lgbtq", "black", "globe", "pride", and "gay". These language-based signals map closely onto the topic-level findings, with right-leaning articles highlighting immigration and trade-related terms, and left-leaning articles emphasizing climate change, gender, race, and LGBTQ issues.[27]

Next, in order to obtain an indication of which topics are more contentious in terms of news coverage, we study the dispersion in slant across topics. Figure A.7, which displays the standard deviation of slant by topic, shows that cultural topics such as gender, LGBTQ issues, and race, together with the topics of welfare, gun rights, the environment, and immigration have the most dispersion in coverage in terms of slant. Conversely, topics such as international conflicts, drugs, natural disasters, trade, and crime have the least dispersion in terms of slant.[28] Figure A.8 shows the full distribution of slant by topic.

6. News Consumption on Facebook

In the previous section, we provided a broad overview of the slant landscape of the hard 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.

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 measure of polarization in news consumption and then estimate it on the Facebook URL-level Activity dataset.

[27] To give a sense of the most slanted topics, overall, we show the mean absolute slant by topic in Figure A.6.

[28] The topic of policing practices also has a relatively low degree of dispersion in slant. We note that our dataset pre-dates the death of George Floyd in May 2020, which sparked a series of protests about police brutality throughout the U.S.

6.1.1 Polarization Measure: Definition

We define polarization in news consumption $\mathscr{P}(\overline{S_R}, \overline{S_D})$ as the difference in click-weighted average slant between Republican and Democrat users, normalized to range between -1 and 1:

$$\mathscr{P}(\overline{S_R}, \overline{S_D}) = \frac{\overline{S_R} - \overline{S_D}}{5}. \tag{1}$$

Here, $\overline{S_j}$ represents the click-weighted average slant for users registered with party $j \in \{\text{D}, \text{R}\}$, 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]

Our measure of polarization has three desirable features that make it an important complement to other measures developed in the literature, such as segregation and favorability. First, our measure 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 Democrat users tend to consume left-leaning articles and Republican users right-leaning articles. Second, our measure is straightforward to interpret: a value of 1 indicates extreme pro-attitudinal news consumption, a value of 0 means both groups consume news with an identical slant, and a value of -1 indicates extreme counter-attitudinal news consumption. Third, our measure allows for an easily interpretable comparison between article-level and outlet-level polarization. Specifically, as shown in Appendix G, 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; the two coincide if and only if users select articles within outlets independently of article-level slant.

6.1.2 Polarization Measure: Estimation

The goal of this section is to estimate our headline polarization measure 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 politically affinity variable described

[29] Recall 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.

in Section 3.2. In order to estimate a polarization measure based on party registration, we need to convert affinity-based click counts into party-registration-based click counts.

From Facebook "political affinity" to party registration. In order to move from political affinity to party registration, we proceed as follows. First, we leverage Twitter (now "X") data to construct a base matrix that tabulates, for each political affinity group (liberal, moderate, conservative), how many individuals appear in each party-registration category (Not registered, Registered Democrat, Registered Republican). Second, we account for missing affinity scores by incorporating information on the share of users who do not have a political affinity measure. Third, we allocate clicks to party registration categories using the observed or imputed distribution of users across the three party registration groups. Fourth, we compute the average slant for each party-registration group by weighting the known average slants of each political-affinity group.

Step 1: Constructing the Base Matrix Let $\mathscr{G} = \{\text{L}, \text{M}, \text{C}, \text{missing}\}$ be the set of political affinity groups in the Facebook URL-level Activity Dataset: liberal, moderate, conservative, and missing. Let $\mathscr{P} = \{\text{NR, RD, RR}\}$ be the party-registration categories: non-registered, registered Democrat, and registered Republican. We start off by constructing the following matrix

$$N_{g,p} \quad \text{for } g \in \{L, M, C\},\ p \in \mathscr{P},$$

where $N_{g,p}$ is the number of users whose political affinity is g and whose party registration is p.

We cannot observe this matrix directly for our Facebook sample, because the URL-level Facebook Activity dataset does not include party registration data. Instead, we rely on the matrix constructed by Barberá et al. (2015) using Twitter (now "X") data. The underlying methodology employed to assign a political affinity score is the same across the Twitter and Facebook datasets.[30] The Twitter dataset studied by Barberá et al. (2015) contains information about party registration for a subset of users. Assuming that the mapping between affinity and party registration $N_{g,p}$ from the Twitter data in Barberá et al. (2015) is a good approximation for our Facebook population, we take the matrix from Barberá et al. (2015) as our "base matrix."

[30] In fact, Facebook's methodology in based on the Barberá et al. (2015) paper.

Summing over all g and p in this base matrix yields the total number of users in the subset who (i) have a measured political affinity and (ii) have known party-registration status:

$$N_{\text{affinity}} = \sum_{g \in \{L,M,C\}} \sum_{p \in \mathscr{P}} N_{g,p}.$$

We know that, in the URL-level Facebook activity dataset $N_{\text{affinity}} \approx 25\%$ of the overall sample of Facebook users (denoted N_{total}). Hence, the number of users without a political affinity score is

$$N_{\text{missing}} = N_{\text{total}} - N_{\text{affinity}} \approx 75\% \times N_{\text{total}}.$$

We use this number in the next step of the construction. Note that, although the share of users without a political-affinity score is large (75%), these users account for only 26% of clicks to hard-news articles on Facebook. This distinction matters for interpreting the imputation in the next step. Missing affinity scores are common at the user level, but they affect only a minority of click observations, so the imputation applies to a relatively small share of the data. In addition, because the missing-affinity group enters the average slant of both party-registration groups through the same group-specific average slant, its main effect is to attenuate differences across groups rather than mechanically induce them.

Step 2: Imputing Party Registration for Missing Affinity Scores For the N_{missing} users who lack a political affinity score in the URL-level Facebook Activity Dataset, we make an assumption about the fraction of them that would be registered Democrats, registered Republicans, or not registered with either the Democratic or Republican party. Specifically, in the absence of more precise information, we assume that their distribution matches the general distribution in the population around the time period we are studying, namely 23% registered Democrats, 16% registered Republicans, and 61% not registered with either the Democratic or the Republican party.[31] Denote the imputed number of Facebook users with a missing political affinity variable who are registered

[31] These percentages are obtained from the 2020 American National Election Studies (ANES) dataset (American National Election Studies, 2020).

Democrats, registered Republicans, or not registered with either party, respectively, by

$$\widetilde{N}_{\text{missing},p} \quad \text{for } p \in \mathscr{P}.$$

We add these $\widetilde{N}_{\text{missing},p}$ counts to the observed counts $N_{C,p}$, $N_{L,p}$, $N_{M,p}$ from the base matrix to obtain a complete matrix $\widehat{N}_{g,p}$ of user counts that covers all $g \in \mathscr{G}$ and $p \in \mathscr{P}$. Note that even though users with $g =$ missing are not assigned a political affinity measure in the the URL-level Facebook Activity Dataset, they are not missing entirely from the dataset, in the sense that we can observe their engagement counts for each URL.

Finally, we leverage the base matrix $\widehat{N}_{g,p}$ of user counts to produce a confusion matrix yielding the proportion of users with a given political affinity $g \in \mathscr{G}$ that falls into each party registration group. Specifically, we define

$$w_{g,p} \;=\; \frac{\widehat{N}_{g,p}}{\sum_{p' \in \mathscr{P}} \widehat{N}_{g,p'}}$$

We report the confusion matrix in Appendix Table A.5.

Step 3: From User Counts to Click Counts Let K_g be the total number of news clicks in our URL-level Facebook Activity Dataset attributed to political-affinity group $g \in \mathscr{G}$. We treat $\{L, M, C, \text{missing}\}$ as four mutually exclusive and exhaustive groups for the click-level analysis.

In order to allocate clicks to party registration categories, we impose the following assumption. For each group $g \in \{L, M, C, \text{missing}\}$, we allocate K_g across $\{\text{NR}, \text{RD}, \text{RR}\}$ in proportion to the share of users in g who fall under each party-registration category. Specifically, we let

$$K_{g,p} \;=\; w_{g,p} \times K_g$$

where $K_{g,p}$ denotes the number of clicks from individuals with political affinity score g imputed to individuals with party registration p.

We also define a variable that captures the total number of clicks from party-registration category p, and denote it by:

$$K_p \;=\; \sum_{g \in \{L,M,C,\text{missing}\}} K_{g,p}.$$

Step 4: Computing Average Slant by Party Registration We observe the *average slant* of clicks for each political-affinity group $g \in \{L, M, C, \text{missing}\}$. Denote these averages by S_g.

For each party-registration category p, we compute its click-weighted average slant, S_p, as the click-weighted mean of $\{S_g\}$, using the fraction of p's clicks that come from each group g:

$$S_p = \frac{1}{K_p} \sum_{g \in \{L,M,C,\text{missing}\}} \left(K_{g,p} \times S_g \right).$$

Thus, if (for instance) 50 percent of the clicks assigned to RD originated from liberal users, 30 percent from moderates, 10 percent from conservatives, and 10 percent from users without a political affinity score (with respective slants $S_L, S_M, S_C, S_{\text{missing}}$), then S_{RD} would be the weighted average of the average slant for liberals, moderates, conservatives, and users without a political affinity score, using the weights 0.5, 0.3, 0.1, and 0.1 respectively.

Step 5: Obtaining the Party-Registration-Based Polarization Measure After computing the click-weighted average slant for each party-registration category $p \in \{\text{RD}, \text{RR}\}$ (denoted S_p), we measure polarization according to Equation 1.

6.1.3 Headline Polarization Results

Applying the procedure described in the previous section, we estimate that, in 2019, the degree of polarization in news consumption on Facebook was 0.17 (see Table 1).

As discussed, our headline polarization estimate 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 measure would be higher: 0.29 (see Table A.3). The gap is consistent with Facebook's "liberal" and "conservative" labels disproportionately selecting strong partisans. Reassuringly, as shown in Section 6.3, when we estimate polarization in an independent browsing dataset where partisanship is observed directly from party registration records, we obtain an estimate of 0.16, which is almost identical to our headline estimate.

6.1.4 Benchmarks

How can we make sense of our polarization index of 0.17? This section provides a series of benchmarks to better understand index's magnitude.

Article- vs. Outlet-level Polarization Measure A central argument of this paper is that outlet-level slant measures can understate polarization in social-media news consumption, because social media might lead partisans to sort ideologically within outlets, not just across outlets. The argument suggests, as a natural benchmark, to recompute our 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 measure of polarization in news consumption of 0.11.[32] In contrast, as shown in Section 6.1, when we allow for the selection of ideologically congenial articles within outlets, we obtain a measure of polarization of 0.17. Thus, the estimate of polarization in news consumption on Facebook increases by 55% when the channel of pro-attitudinal news consumption within outlets is taken into account. This gap is similar or larger when, instead of focusing on news consumed, we focus on the articles individuals viewed or shared (Figure A.9).

These results are in line with our initial intuition that the curation process on social media, which operates at the level of individual news articles, promotes substantive pro-attitudinal news consumption within outlets.

Additional Benchmarks Our first benchmark compares the difference in click-weighted average slant between Republican and Democrat users (i.e., the numerator in Equation 1) to the standard deviation in the overall database of news articles. The mean difference in the slant of articles clicked

[32] As 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 were to consider a measure of polarization in news consumption based directly on the political affinity score, we would 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.

on by Republican and Democrat users is 0.84. The standard deviation in our overall database of news articles is 0.99. Therefore, the mean difference in slant for articles clicked on by Republican and Democrat users is 0.85 standard deviation units.

We can also interpret the difference from the previous section in standard deviation units. Specifically, the polarization of news sharing on social media measured from article-level slant is 0.27 standard deviation units higher than measured from outlet-level (average) slant.

Another benchmark for our polarization measure compares the slant difference in articles consumed by Republicans and Democrats on Facebook to the difference between the average slant of articles from ideologically contrasting outlets. We find that the slant difference between the average news articles consumed by Republicans and Democrats on Facebook equals the distance in slant between the average article from the Washington Times and the average article from the Washington Post (where the Washington Post and The Washington Times rank at the 11th and 82nd percentile of the overall distribution in slant across our outlets, respectively).

For additional context, we can compare the slant difference in articles clicked on by Republicans and Democrats on Facebook to the average slant difference between articles written by different New York Times columnists. For instance, we find that the slant difference in news consumption on Facebook between Democrats and Republicans is comparable to the slant difference between the average article written by liberal columnist Nicholas Kristof and by moderate/conservative columnist David Brooks.[33]

Our final benchmark compares the distance in slant between articles consumed by Republicans and Democrats on Facebook to the slant difference between articles shared by various politicians on Twitter (now "X"). We find that the slant distance in our analysis is slightly smaller than the average distance between articles shared by Democratic Senator Chuck Schumer and those shared by Republican Senator Marco Rubio.

6.2 Mechanisms

What drives the high degree of polarization in news consumption on social media? In this section, we address four key mechanisms. First, we study the degree to which polarization in news

[33]David Brooks left the New York Times in 2026.

consumption on Facebook is driven by the selection of ideologically congenial articles within outlets. Second, we analyze the extent to which polarization in news consumption is driven by the articles that users are exposed to on their feeds versus the articles they select conditional on the ones they are exposed to. Third, we examine how the distribution of articles that circulate widely on social media differs from the distribution of slant on the production side. Fourth, we perform a back-of-the-envelope calculation to estimate, under a host of assumptions, the degree to which echo chambers are responsible for polarized news exposure and consumption on Facebook.

6.2.1 Pro-attitudinal news consumption

In this section, we provide direct evidence that Facebook users with a political affinity score consume pro-attitudinal news within outlets. Figure 5 focuses on the top 20 outlets in terms of online visits and shows that, even within an outlet, the average article consumed by users with a liberal political affinity score is virtually always more left-leaning than the average article consumed by users with a conservative political affinity score. Table J.1 presents statistics about the slant of the average article consumed by users with a liberal and a conservative political affinity for each outlet in our dataset. The pattern is similar: partisans tend to consume pro-attitudinal content even within outlets.

6.2.2 Selection vs. Exposure

Our measure of polarization in news consumption on social media can be decomposed into two factors: exposure and selection conditional on exposure. Exposure refers to articles that appear on partisans' feeds; selection conditional on exposure refers to the way in which partisans select which articles to click on and read among the articles that appear on their feeds. In the URL-level Facebook Activity Dataset, we proxy exposure using the number of users who *viewed* a given URL on Facebook, and we proxy news consumption using the number of users who *clicked* that URL.

We find that the lion's share of polarization in news consumption can be accounted for by differential *exposure* to slanted articles. Specifically, we find that the polarization in news exposure on Facebook is 0.14. Thus, exposure alone can account for 82% of the degree of polarization in news consumption on Facebook. In other words, if Democrats and Republicans on Facebook clicked

randomly among the articles that appear in their feeds, rather than actively selecting articles among the ones that appear in their feeds, polarization in news consumption would be only 18% smaller than it currently is.[34,35,36]

6.2.3 Production vs. Distribution

How does the distribution of slant from the top 100 U.S. outlets in 2019 translate to the social media ecosystem? The Facebook environment need not reflect the distribution of slant on the production side for three main reasons. First, the homophilous network structure on social media, together with pro-attitudinal sharing patterns by Democrats and Republicans, can generate *echo chambers* in which partisans are primarily exposed to and consume ideologically congenial news (Sunstein, 2018). Second, personalized ranking algorithms might produce *filter bubbles* that promote pro-attitudinal content (Pariser, 2011). Third, social media platforms offer ample scope for individuals to customize their feeds (Negroponte, 1995). In this section, we analyze various ways in which the distribution of slant in the Facebook ecosystem differs from the distribution of slant on the production side, as well as the drivers of such differences.

Slant of the Articles Circulating on Facebook We begin by analyzing, as a function of an article's slant, the probability that the article is shared publicly on Facebook more than 100 times. This definition corresponds to the criterion used by Meta for including an article in the URL-level Facebook Activity Dataset. We measure this outcome using an indicator that equals one if a given article from our database of articles appears in the URL-level Facebook Activity Dataset and zero otherwise.

Figure 6 shows the probability that an article from our database appears in the URL-level Facebook Activity Dataset as a function of the article's slant. We find a stark V-shaped pattern: the more

[34] As in our preferred measure of polarization from Section 6.1, this measure employs, as a definition of partisanship, being a registered Democrat or a registered Republican. If we were to consider a measure of polarization in news consumption based directly on the political affinity score, we would obtain a similar result in terms of the percent of polarization in news consumption that is explained by exposure and by selection conditional on exposure.

[35] Of course, our analysis is in "partial equilibrium." A "general equilibrium" analysis would take into account the feedback patterns that exist between exposure, consumption, and sharing behavior.

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

extreme the article, the higher the probability that it gets shared on Facebook more than 100 times. The difference in the likelihood of getting shared more than 100 times across slant categories can be substantial. For instance, centrist articles with a slant of 0 have an 11% chance of being widely shared on Facebook, whereas extreme articles with an absolute slant of 2 or more have an over 39% chance of being widely shared on the platform. Thus, extreme articles are approximately 4 times more likely to be publicly shared on Facebook at least 100 times. The pattern above reflects the combined effect of outlet posting decisions, user-sharing behavior, platform curation, and other factors that shape which articles circulate widely on Facebook. We note that this analysis includes virtually all hard news articles published online by major outlets in 2019, thus offering the most comprehensive picture to date of the relationship between slant and sharing behavior on Facebook.

In light of the sharing patterns documented above, the distribution of slant of the articles that are shared on Facebook at least 100 times is quite different from the distribution of slant on the production side. Appendix Figure A.11 overlays the two distributions and shows how the distribution of slant of the articles that are shared on Facebook at least 100 times has much thicker tails than the distribution of slant on the production side. In fact, comparing the variance of the two distributions, we see that the variance of the former is 60% higher than the variance of the latter.

Over-representation of slanted news types and topics on Facebook In Section 5, we showed that opinion pieces, national news, and articles about certain topics are particularly slanted. Here we show that articles with those characteristics are much more likely to be shared on Facebook at least 100 times.

Appendix Figure A.12 shows that 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 shows that, compared to the production side, both national news and opinion pieces are over-represented among the articles shared on Facebook at least 100 times by a factor of at least 44%.

Pro-attitudinal sharing patterns What role do pro-attitudinal sharing patterns by liberals and conservatives play in promoting the circulation of slanted articles on Facebook? For this exercise, we limit our attention to articles in the URL-level Facebook Activity Dataset; i.e. articles publicly

shared on Facebook at least 100 times. Figure 7 shows the distribution of the number of shares per article, separately for individuals with liberal and conservative political affinity scores, as a function of an article's slant. The figure depicts striking differences in sharing behavior: liberals share left-leaning content around three times more often than moderate content and almost never share right-leaning articles. Conservatives display the mirror image of this behavior: they primarily share right-leaning content, and they almost never share left-leaning articles.[37]

We present the same analysis separately for articles that users clicked on before sharing and articles that users shared without first clicking on them in Appendix Figure A.16. We do not find dramatic differences between the two figures, suggesting that reading an article does not increase or decrease the tendency to share pro-attitudinal news. This pattern is consistent with article headlines conveying slant that is similar to the slant of the full article, at least in the context of sharing behavior.

We do find that the sharing patterns of opinion articles on Facebook are substantially more pro-attitudinal than the sharing patterns for non-opinion articles (Appendix Table A.6). Specifically, the sharing patterns of opinion articles on Facebook are 2.7 times more polarized than those of non-opinion articles.

In the next section, we show how our finding that both liberals and conservatives share pro-attitudinal articles much more frequently than counter-attitudinal ones can help us assess the degree to which polarized news exposure and consumption on social media is driven by echo chambers.

6.2.4 The Role of Echo Chambers

Echo chambers can generate polarized news exposure on Facebook when two features of the platform interact: politically homophilous friendship networks and pro-attitudinal sharing by users. In Section 6.2.3, we documented strong pro-attitudinal sharing patterns. We now ask how much

[37] We present the same figure after residualizing for outlet fixed effects in Appendix Figure A.14. Compared to the un-residualized version, the distributions of shares by liberal and conservative users do move closer together, indicating that outlets play a role in explaining the differential sharing patterns among liberals and conservatives. However, even after residualizing on outlet fixed effects, the distribution of shares by liberals is markedly to the left of that by conservatives, showcasing, again, the important role of within-outlet differences in sharing behavior across the political spectrum. Figure A.15, which shows stark within-outlet differences in the mean slant of articles shared by Facebook users with conservative vs. liberal political affinity, further underscores this point.

polarization in exposure would arise if feeds were shaped only by users' friend networks and by the content shared by those friends, shutting down other channels such as the personalized ranking algorithm and users' own customization decisions. For simplicity, we employ throughout this section the polarization measure based on political affinity scores from the Facebook URL-level Activity Dataset, rather than party registration.[38]

Formally, let s_f denote the slant of a random article that a user is exposed to because it was shared by one of her Facebook friends. We are interested in the degree of polarization in exposure that would prevail under this "friends-only" counterfactual:

$$\mathscr{P}^{echo}_{exposure} = \frac{E\left(s_f \mid user\ is\ conservative\right) - E\left(s_f \mid user\ is\ liberal\right)}{5}.$$

This object captures how different the average slant of friend-generated exposure would be for conservative and liberal users if feeds reflected only what their friends shared.

A convenient way to express this quantity is through the political composition of users' friendship networks. For $i \in \{lib, con\}$ and $j \in \{lib, mod, con\}$, let $x_{\{i \leftarrow j\}}$ denote the fraction of friends of type j for a user of type i. Using the law of total expectation,

$$E\left(s_f \mid user\ is\ i\right) = \sum_{j \in \{lib, mod, con\}} E\left(s_f \mid user\ is\ i, poster\ is\ j\right) P\left(poster\ is\ j \mid user\ is\ i\right).$$

where "poster" refers to the friend who shared the article. We obtain the terms $P\left(poster\ is\ j \mid user\ is\ i\right)$ from the network homophily estimates in Table S2 of Bakshy, Messing and Adamic (2015), which provide the relevant $x_{\{i \leftarrow j\}}$ statistics.

The challenge is that we do not observe $E\left(s_f \mid user\ is\ i, poster\ is\ j\right)$, that is, the average slant of the content that a type-j poster shares specifically with a type-i user. To make progress, we impose a conservative monotonicity assumption. For each poster type $j \in \{lib, mod, con\}$, we assume

$$E\left(s_f \mid user\ is\ liberal, poster\ is\ j\right) \leq E\left(s_f \mid poster\ is\ j\right) \leq E\left(s_f \mid user\ is\ conservative, poster\ is\ j\right).$$

[38] Using the polarization measure 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.

In words, conditional on poster ideology, the content shared with liberal users is weakly more left-leaning than the average content shared by that poster type, while the content shared with conservative users is weakly more right-leaning than the average content shared by that poster type. This assumption implies that replacing the unobserved conditional means with the unconditional mean for each poster type yields a lower bound on the amount of polarization generated by echo chambers. Intuitively, if liberal posters share especially left-leaning content with liberal friends and especially less left-leaning content with conservative friends, then using the poster-type average $E\left(s_f \mid poster\ is\ j\right)$ understates the true gap in exposure across users.

Let $\bar{y}_{lib}$, $\bar{y}_{con}$, and $\bar{y}_{mod}$ denote the average slant of articles shared by liberal, conservative, and moderate users, respectively. Under the assumption above, a lower bound on echo-chamber polarization is therefore

$$\underline{\hat{\mathscr{P}}}^{echo}_{exposure}(\bar{y}_{lib},\bar{y}_{con},\bar{y}_{mod},x_{\{lib\leftarrow lib\}},x_{\{lib\leftarrow con\}},x_{\{lib\leftarrow mod\}},x_{\{con\leftarrow lib\}},x_{\{con\leftarrow con\}},x_{\{con\leftarrow mod\}}) =$$

$$= \frac{1}{5}\{[\bar{y}_{lib}\cdot x_{\{con\leftarrow lib\}} + \bar{y}_{con}\cdot x_{\{con\leftarrow con\}} + \bar{y}_{mod}\cdot x_{\{con\leftarrow mod\}}]-$$

$$[\bar{y}_{lib}\cdot x_{\{lib\leftarrow lib\}} + \bar{y}_{con}\cdot x_{\{lib\leftarrow con\}} + \bar{y}_{mod}\cdot x_{\{lib\leftarrow mod\}}]\}.$$

The first bracketed term is a proxy for the average slant to which conservative users would be exposed if their feeds were governed only by 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 measure of polarization in exposure based on political affinity scores.

We estimate $\bar{y}_{lib}$, $\bar{y}_{con}$, and $\bar{y}_{mod}$ using the sharing patterns documented in Section 6.2.3, and we obtain the network homophily terms $x_{\{i\leftarrow j\}}$ from Bakshy, Messing and Adamic (2015). Following this approach, we conclude that the amount of polarization in news exposure that would result from echo chambers alone is at least 59% of the actual degree of polarization in news exposure calculated using the political affinity scores as a measure of partisanship.

This exercise is not a causal decomposition of feed polarization. Rather, it provides a back-of-the-envelope benchmark: how much polarization in exposure would arise if feeds were shaped only by network homophily and the observed sharing patterns of users, abstracting from other channels such as ranking algorithms and user customization.

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

This section investigates whether the news consumed via Facebook exhibits greater ideological polarization compared to news accessed through other online sources. Using detailed individual-level browsing history data on a panel of 1,637 registered Democrats and Republicans, we analyze the slant of news articles consumed by panel members, comparing instances where the referrer to the article was Facebook versus other online sources.

To conduct this analysis, we regress the slant of a news article consumed by an individual on a dummy variable indicating whether Facebook was the referrer, a dummy for party registration status (Republican vs. Democrat), and their interaction, while including individual fixed effects. The inclusion of individual fixed effects ensures that identification comes from comparing the *same individual*'s news consumption across channels—i.e., how the slant of articles consumed by a given person differs when they arrive via Facebook versus other sources—rather than from differences in the types of individuals who use different channels.

The results, shown in column 2 of Table A.7, reveal that news consumption on Facebook is significantly more polarized. The interaction coefficient implies that the gap in slant between articles consumed by Republicans and Democrats is 0.13 units larger for articles accessed through Facebook compared to those not accessed through Facebook—corresponding to an increase in the polarization index of 0.03 units, or 16% relative to the mean. In other words, when the same users access news via Facebook rather than through other channels, the ideological gap in the content they consume increases substantially.

To better understand the sources of this difference, column 3 adds outlet fixed effects, thereby comparing Democrats and Republicans who read content from the same outlet. The interaction coefficient declines modestly but remains large and statistically significant (0.10). This result indicates that the higher polarization observed on Facebook is not driven solely by users sorting into more extreme, pro-attitudinal outlets on Facebook, but also reflects substantially higher within-outlet sorting into more ideologically slanted articles on Facebook relative to elsewhere.

Column 4 shows that this effect is not unique to Facebook but extends more broadly to social

media platforms. Column 5 further demonstrates that the result remains robust when including outlet fixed effects, indicating that social media amplifies polarization to a large extent through within-outlet article selection. Together, these findings suggests that the algorithmic curation and mechanisms unique to social media platforms exacerbate ideological sorting relative to the broader online ecosystem.

We next compare polarization in news consumption across the ten most common referrer types by augmenting the above regression to include referrer-type indicators and their interaction with party registration. We present the results in Figure 8. We find that polarization is smallest for news accessed via news aggregators, followed by search engines; it is largest for Faceboook and for other social media platforms (driven by Twitter, now "X"), further underscoring the potential importance of social media platforms in amplifying polarization in news consumption.

In terms of magnitudes, we find that news accessed via social media is 2.4 times more polarized than the news accessed via search engines; equivalently, the gap corresponds to 0.33 standard deviations in relative slant. In other words, polarization is more than twice as large in social media environments compared to a benchmark of news visits via search engines.

Taken together, these results suggest that while ideological sorting in news consumption exists across all channels, social media platforms are associated with substantially greater polarization, driven both by differences in outlet exposure and—importantly—by differences in the specific articles consumed from the same outlets.

6.4 Robustness Checks

This section has primarily used the URL-level Facebook Activity Dataset to estimate polarization in news consumption and to study its drivers. In this section, we discuss three natural concerns about our conclusions: i) first, whether they depend on the set of outlets in our sample, ii) whether they depend on the construction of Facebook's political-affinity measure, and iii) whether they depend on the fact that the dataset only includes URLs that were publicly shared on Facebook at least 100 times. We address each of these concerns in detail in Appendix F and briefly summarize the results here.

First, Appendix Section F.1 examines the role of expanding the outlet sample beyond the top 100 U.S. outlets in our baseline analysis. Using Comscore data to rank outlets by size, we extrapolate our polarization estimates to progressively larger outlet sets. The resulting patterns suggest that extending the analysis to additional outlets would not materially alter our main conclusions.

Second, Appendix Section F.2 presents a battery of robustness checks addressing the limitations of Facebook's political-affinity measure. The appendix shows that our results are robust to excluding the outlets that overlap with the set of political pages used in the construction of Facebook's political-affinity measure, and to replacing that measure with an alternative ideology proxy based on demographic characteristics. Together with the result from the previous section that leverages a dataset containing actual party registration status, these exercises suggest that our findings are not an artifact of the particular partisanship measure available in the URL-level Facebook Activity Dataset.

Third, Appendix Section F.3 addresses the fact that the URL-level Facebook Activity Dataset only includes articles that were publicly shared on Facebook at least 100 times. Using the browsing panel, we recompute polarization for all hard-news articles reached through Facebook, including those that do not appear in the URL-level dataset. The resulting estimate is very close to our baseline estimate, indicating that restricting attention to widely shared articles has little effect on our conclusions. Appendix Section F.3 also shows that a similar conclusion is reached using an extrapolation exercise based on the Facebook URL-level data.

Overall, these robustness exercises support the main conclusion of this section: news consumption on Facebook is polarized, article-level slant matters quantitatively for measuring that polarization, and the main qualitative patterns are not sensitive to the particular sample restrictions or measurement choices imposed by the available data.

7. Conclusion

In this paper, we introduced a novel measure of article-level slant and used it to quantify the degree of polarization in news consumption on Facebook. Our guiding hypothesis was that, due to the way in which content is curated on social media platforms, an important driver of polarization in

news consumption on social media is ideological sorting within outlets. Of course, this driver of polarization can only be captured by means of a measure of slant at the article level akin to the one that we develop. Consistent with our hypothesis, we find a much higher degree of polarization when employing our fine-grained article-level slant measure than we do when employing a coarser measure of outlet-level slant that mimics the approach taken by much of the literature.

Overall, this paper provides strong evidence that the news environment on social media is more polarized than previously thought. This insight is crucial for understanding the dynamics of political polarization in the digital age and for informing both academic inquiry and policy interventions aimed at fostering a more balanced and informed public discourse. Furthermore, this paper highlights the importance of using content-based article-level slant measures when studying today's fragmented media landscape.

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."

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*, 19(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.

Datareportal. 2023. "Facebook Users, Stats, Data & Trends." https://datareportal.com/essential-facebook-stats.

Djourelova, Milena, Ruben Durante, Elliot Motte, and Eleonora Patacchini. 2023. "Experience, Narratives, and Climate Change Beliefs." *Available at SSRN 4680430*.

Downs, Anthony. 1957. "An economic theory of political action in a democracy." *Journal of Political 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. 2023. "Curation bubbles." *OSF Preprints. URL: https://osf. io/vbwer*.

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.

Levy, Roee. 2021. "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment." *American Economic Review*, 111(3).

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.

Negroponte, Nicholas. 1995. *Being Digital.* Knopf.

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, Richard Fletcher, Kirsten Eddy, Craig T Robertson, and Rasmus Kleis Nielsen. 2023. "Reuters Institute Digital News Report 2023."

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.

Pew Research Center. 2023*a*. "American Trends Panel, Wave 135 September 2023."

Pew Research Center. 2023*b*. "News Platform Fact Sheet."

Puglisi, Riccardo, and James M Snyder. 2011. "Newspaper coverage of political scandals." *The Journal of Politics*, 73(3): 931–950.

Sunstein, Cass R. 2018. *#Republic: Divided Democracy in the Age of Social Media.* Princeton University Press.

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.

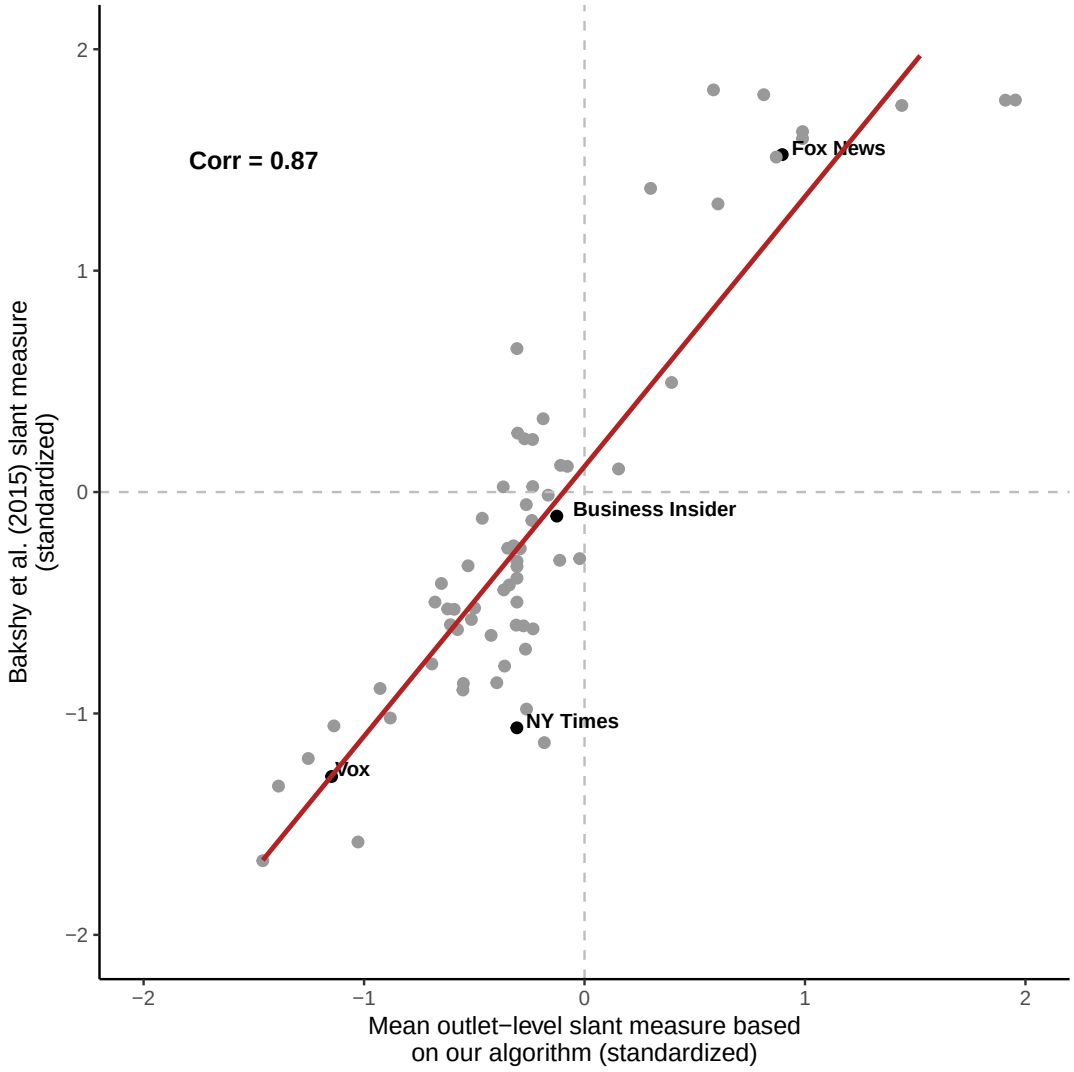
Tables and Figures

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	-0.57	0.13	0.70	0.14
Clicks	-0.57	0.27	0.84	0.17
Shares	-0.66	0.58	1.24	0.25

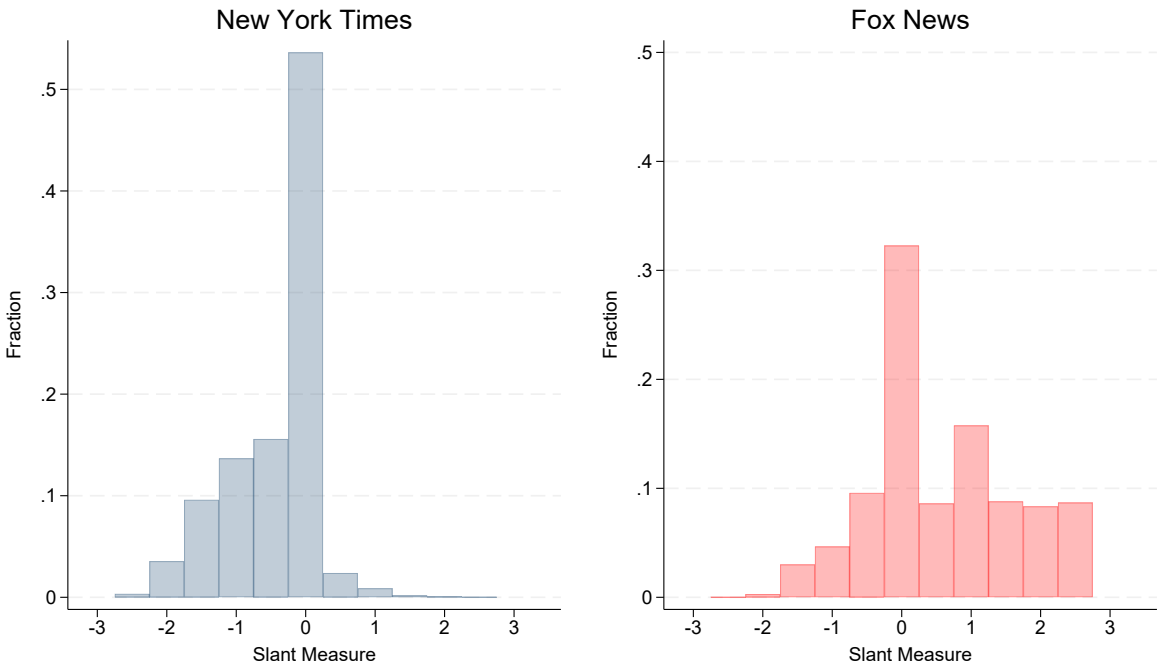
Notes: The first two columns of the table report the views/clicks/shares-weighted average slant of the news diets of registered Democrat and registered Republican Facebook users imputed according to the methodology described in Section 6.1.2. The third column shows the difference between the first two columns. The last column normalizes the difference 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 $N = 233,395$ hard 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).

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



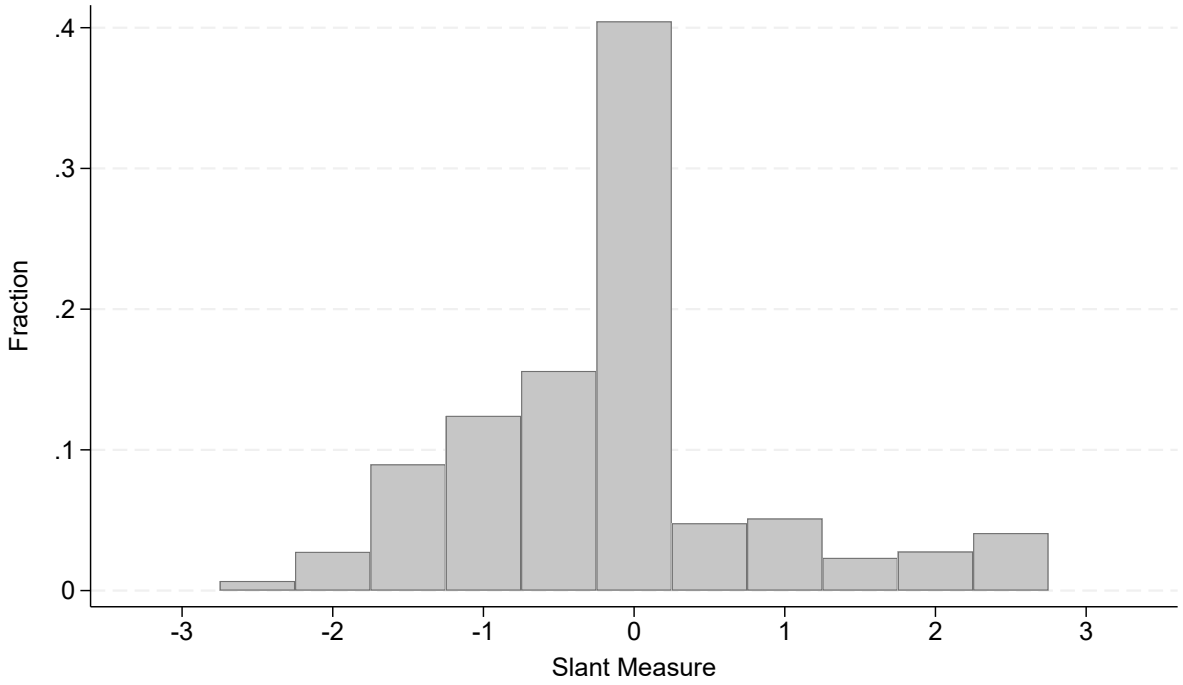
Notes: This figure presents our outlet-level slant measure on the horizontal axis, calculated as the average article-level slant across all articles published by an outlet (with article slant determined by the fine-tuned GPT-4o's predictions, as described in Section 4.2). The vertical axis represents 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



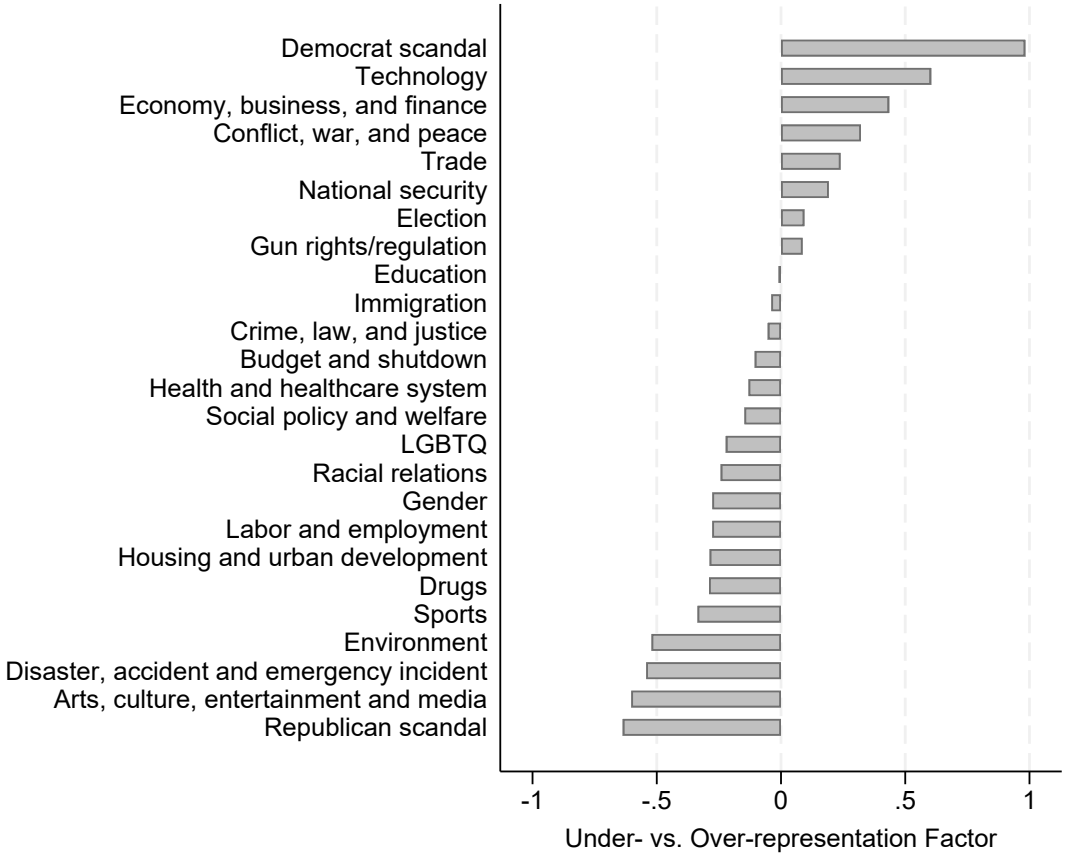
Notes: This figure shows the distribution of slant among all hard news articles ($N = 82,139$), published online by the New York Times and Fox News in 2019. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

Figure 3: Overall Distribution of Slant Across Articles



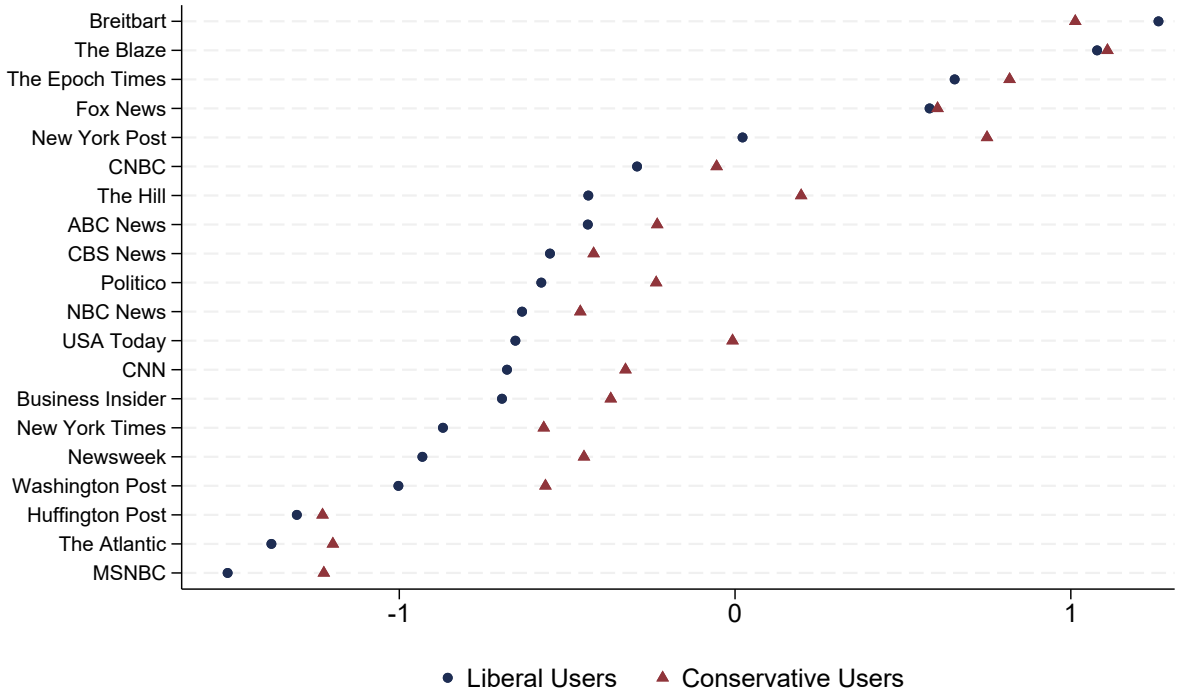
Notes: This figure shows the distribution of slant among all hard news articles ($N = 1,057,124$) published online by the top 100 U.S. news outlets in 2019. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party).

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



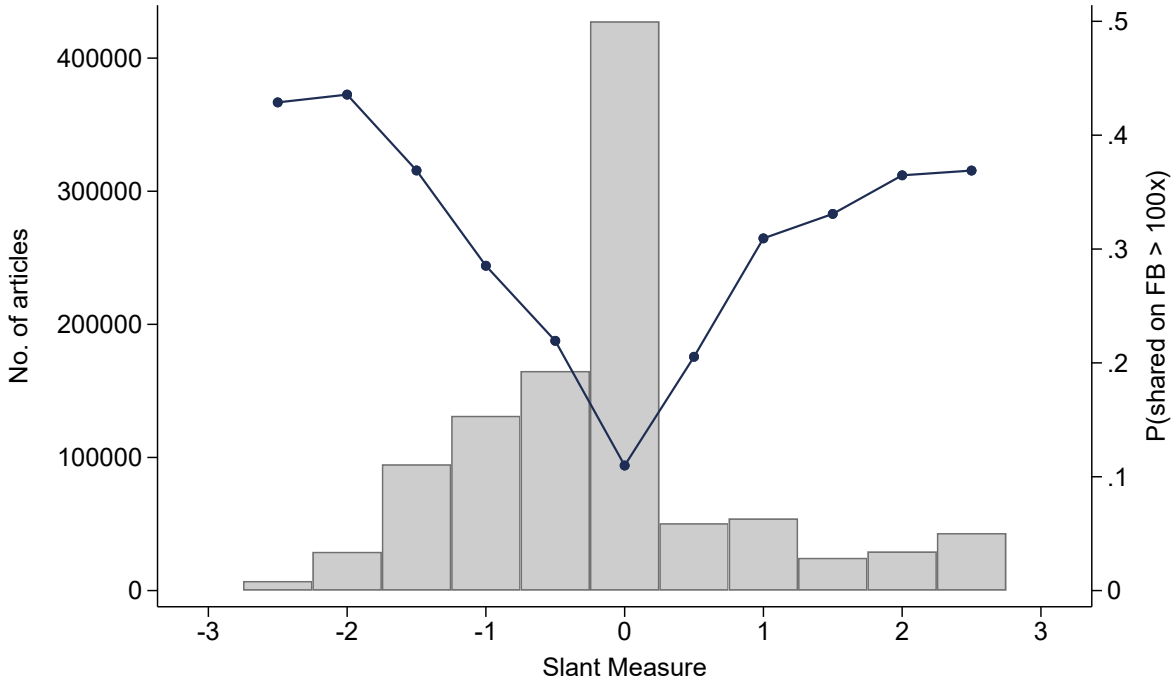
Notes: This figure shows the degree of over- vs. under-representation of topics among right-leaning articles. The figure is constructed as follows. First, all "moderate" articles with a slant between -0.5 and 0.5 are dropped. Second, the remaining articles are defined as "right-leaning" if they have 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 5: Average Slant of Articles Clicked on Facebook by Outlet and User Political Affinity



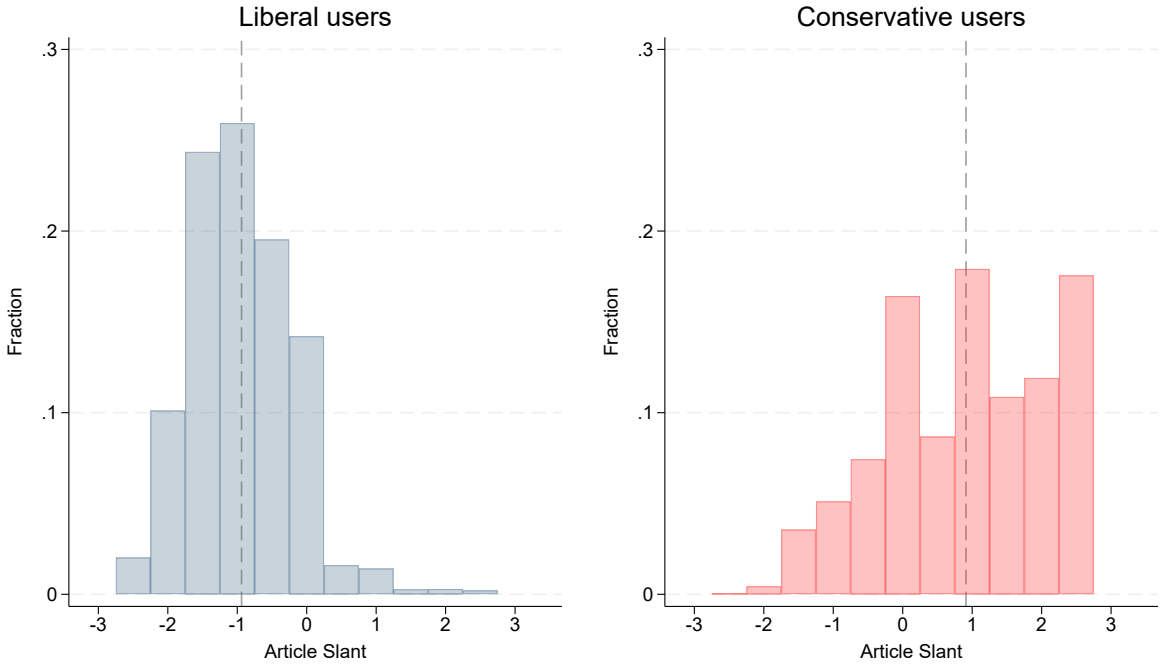
Notes: This figure shows the clicks-weighted average slant of articles by Facebook users' political affinity and outlet. The sample includes all hard news articles published online by the top 20 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e., articles shared at least 100 times on Facebook). Slant is given by the fine-tuned GPT-4o'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 Shared on Facebook $>$ 100 Times by Article Slant



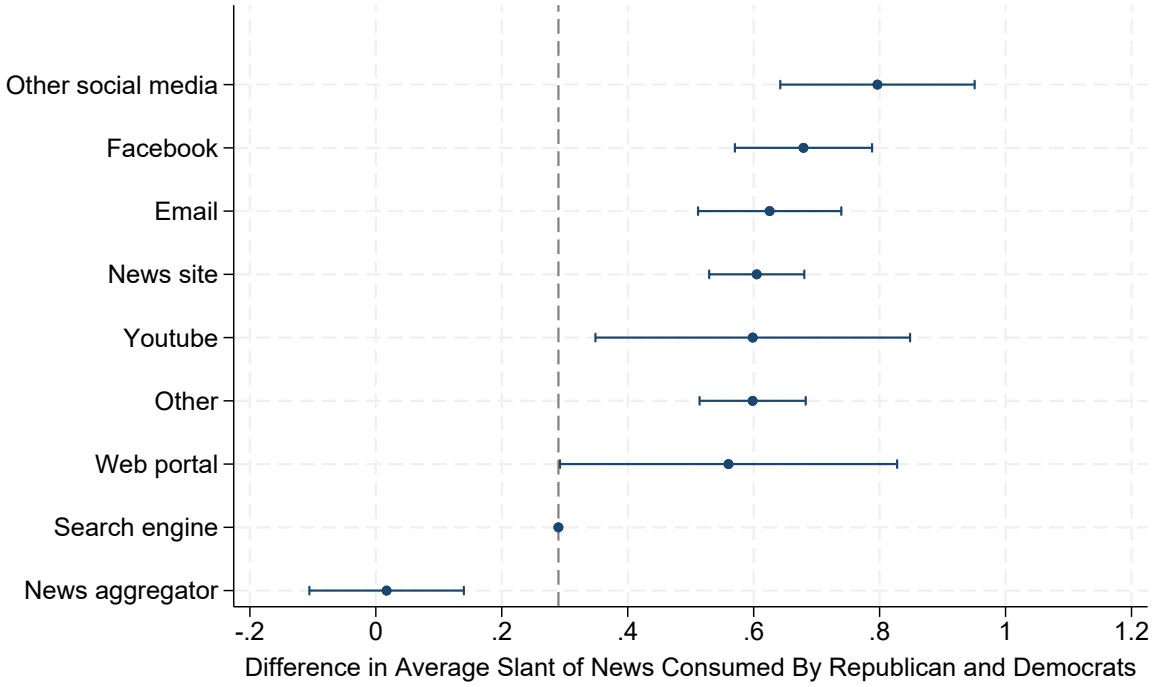
Notes: The gray bars in this figure show the distribution of slant among all hard news articles ($N = 1,057,124$) published online by the top 100 U.S. news outlets in 2019 (left y-axis). The blue line shows the fraction of articles in each bin that were shared at least 100 times on Facebook (right y-axis), as measured by whether they appear in the URL-level Facebook Activity Dataset. The sample includes all hard news articles ($N = 1,057,124$) published online by the top 100 U.S. news outlets in 2019. Slant is given by the fine-tuned GPT-4o'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 $N = 233,395$ hard 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 GPT-4o'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 8: Differences in the Average Slant of News Diets by Party Registration Status and Referrer Type



Notes: This figure shows the results from a regression of the slant (on a scale from -2.5 to +2.5) of an individual article consumed by a browsing panelist on panelist fixed effects, as well as on dummies for different referrer origins, and their interaction with a dummy that equals one if the panelist is a registered Republican. The figure plots the estimates of the interaction coefficients, to which we add the mean difference in slant observed among registered Republicans and Democrats in the left-out category (search engine referrals). "Other social media" includes referrals from Twitter, Instagram, Pinterest, TikTok, LinkedIn, Snapchat, Reddit, Nextdoor, BeReal, and Tumblr. "Web portal" includes referrals from yahoo.com, msn.com and aol.com. "News aggregator" includes referrals from google news, yahoo news, drudgereport.com, newsbreakapp.com, smartnews.com, newser.com, and fark.com. 95% confidence bars are also plotted. The sample is restricted to all hard news article visits by registered Democrats and Republicans (1,637 individuals and 48,388 observations).

Appendix For Online Publication

Table of Contents

A Additional Tables and Figures **A.2**

B Labels Construction **A.28**

C Classifying Other News Characteristics **A.29**

D Variance Decomposition **A.32**

E Correlation Ceiling Between Model Predictions and Expert Labels **A.33**

F Robustness **A.35**

F.1 Extending Results to Additional Outlets A.35

F.2 Robustness: Political Affinity Measure A.35

F.3 Robustness: Threshold in the URL-Level Facebook Activity Dataset A.36

G Polarization, Segregation, and Favorability **A.39**

G.1 Empirics A.39

G.2 Theory A.42

H Slant distribution for each outlet in the top 100 **A.51**

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	1,379 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. (2023)	Correlation between predicted ideology and slant	Content-based	Article		Subset of Twitter's 10% Decahose sample	Shares	1000 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)	Segregation	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. (2025) [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 hard-news articles of which ∼230,000 appear in the social media data

Notes: 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 go to outlets (or articles) predominantly visited by Republicans and the share of Republican visits that go to 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 Section 6.1.2. 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 6.1.2. The table only refers to measures and data used for the primary estimates (e.g., segregation or polarization). For example, Green et al. (2023) also study an audience-based measure of slant and analyze Facebook data, but this measure and data are not used to compute the principal estimand and, therefore, are not mentioned in the table. Similarly, 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.

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	0.20	0.29
	(0.41)	(0.40)	(0.45)
Slant	-0.14	-0.13	-0.20
	(0.99)	(0.90)	(1.25)
FB views (1000s)			219.40
			(769.36)
FB clicks (1000s)			9.33
			(33.40)
FB shares (1000s)			1.81
			(7.87)
N	1057124	823729	233395

Notes: This table reports summary statistics for all hard news articles published by the top 100 U.S. news outlets in column (1). See Appendix Figure J.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

	Views		Clicks		Shares	
Political Affinity	Share(%)	Mean Slant	Share(%)	Mean Slant	Share(%)	Mean Slant
Liberal	36%	-0.81	38%	-0.85	47%	-0.94
Moderate	11%	-0.44	8%	-0.39	8%	-0.42
Conservative	25%	0.46	28%	0.60	32%	0.91
Unassigned	28%	-0.34	26%	-0.29	13%	-0.23

Panel B: Polarization Based on Political Affinity

	Views	Clicks	Shares
Average Slant Difference	1.27	1.45	1.85
Polarization Index	0.25	0.29	0.37

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 $N = 233,395$ hard 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 hard 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 hard 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 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 GPT-4o'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: Most Clicked Liberal and Conservative Articles

Liberal	**Conservative**
I Wish I'd Had A 'Late-Term Abortion' Instead Of Having My Daughter*	When the villain is Obama, not Trump, news suddenly becomes not worth reporting*
Trump Is Not Well*	Mueller's report looks bad for Obama*
12 Times Ellen DeGeneres Perfectly Called People Out*	Obama Built The 'Cages' for Illegals, Not Trump, Says Obama ICE Chief
Georgia's abortion law imprisons women with miscarriages*	The speech they're trying to hide: President Trump's stellar UN speech*
A Stain on the Honor of the Navy*	Hollywood Film Depicts Trump Supporters Being Hunted for Sport by Liberals*
I'm From Alabama And Gave Birth To My Rapist's Child Because I Couldn't Get An Abortion*	CNN's Jim Acosta mocked for accidentally proving that border walls work
Should 11-year-old girls have to bear their rapists' babies? Ohio says yes.*	Trump's shutdown trap?*
I've Talked With Teenage Boys About Sexual Assault for 20 Years. This Is What They Still Don't Know*	Donald Trump was elected to break the elite. Of course they want to impeach him*
15 Reasons Why Vaccinating Your Kids Is The Worst Thing You Can Do For Them*	Rural Americans would be serfs if we abolished the Electoral College*
We should all be appalled by Donald Trump's tweet about Greta Thunberg*	Read Trump's Letter to Pelosi Protesting Impeachment*

Notes: This table lists the titles of the most clicked liberal and conservative articles with absolute slant of 2 or more, where we include up to two liberal and two conservative articles from each outlet. Articles marked with an asterisk (*) are opinion pieces. The sample covers all $N = 233,395$ hard news articles published online by the top-100 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. For clarification, the article titled "15 Reasons Why Vaccinating Your Kids is The Worst Thing You Can Do For Them" and listed in the "Liberal" column takes a fervent pro-vaccine stance and strongly relies on the rhetorical device of irony.

Table A.5: 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.6: 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	-0.76	0.73	1.49	0.30
Clicks	-0.77	0.83	1.60	0.32
Shares	-0.92	1.38	2.30	0.46
Panel B: Non-opinion pieces				
Views	-0.52	-0.03	0.49	0.10
Clicks	-0.51	0.08	0.59	0.12
Shares	-0.59	0.27	0.86	0.17

Notes: This table shows the same statistics as in Table 1 for all opinion news articles ($N = 66,898$) 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 C.

Table A.7: Differences in the Slant of News Diets by Party Registration Status Among Browsing Panelists

	(1) Article Slant	(2) Article Slant	(3) Article Slant	(4) Article Slant	(5) Article Slant
Registered Republican	0.81***				
	(0.05)				
from Facebook		-0.05*	-0.07***		
		(0.03)	(0.02)		
Republican × from Facebook		0.13***	0.10***		
		(0.05)	(0.04)		
from Social Media				-0.03	-0.05**
				(0.02)	(0.02)
Republican × from Social Media				0.16***	0.11***
				(0.04)	(0.03)
Individual FE	No	Yes	Yes	Yes	Yes
Outlet FE	No	No	Yes	No	Yes
Observations	2,118	46,934	46,934	46,934	46,934

Notes: This table shows results from a regression of the slant (on a scale from -2.5 to +2.5) of an individual article consumed by a browsing panel member on a dummy for Republican party registration in column 1; column 2 includes a dummy for whether the panelist got referred to a given article via Facebook, as well as its interaction with the Republican dummy, and individual fixed effects; column 4 includes a dummy for referral via social media (including Facebook, Twitter, Instagram, Pinterest, TikTok, LinkedIn, Snapchat, Reddit, Nextdoor, BeReal, and Tumblr) in our referrer origins, as well as its interaction with the Republican dummy, and individual fixed effect. Columns 3 and 5 are like columns 2 and 4, respectively, plus outlet fixed effects. Column 1 restricts to article clicks that i) were referred from Facebook and that ii) pertain to articles included in the URL-level Facebook Activity Dataset, while no such restrictions apply in columns 2-5. The sample is restricted to individuals who are either registered Democrats or registered Republicans. The table is based on individual-level browsing panel data from a representative sample of individuals as described in Section 3.3. An observation is an individual click on a hard news article by a browsing panelist. Robust standard errors are provided in parentheses.

Table A.8: 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	-0.44	0.26	0.70	0.14
Clicks	-0.48	0.34	0.82	0.16
Shares	-0.56	0.77	1.33	0.27

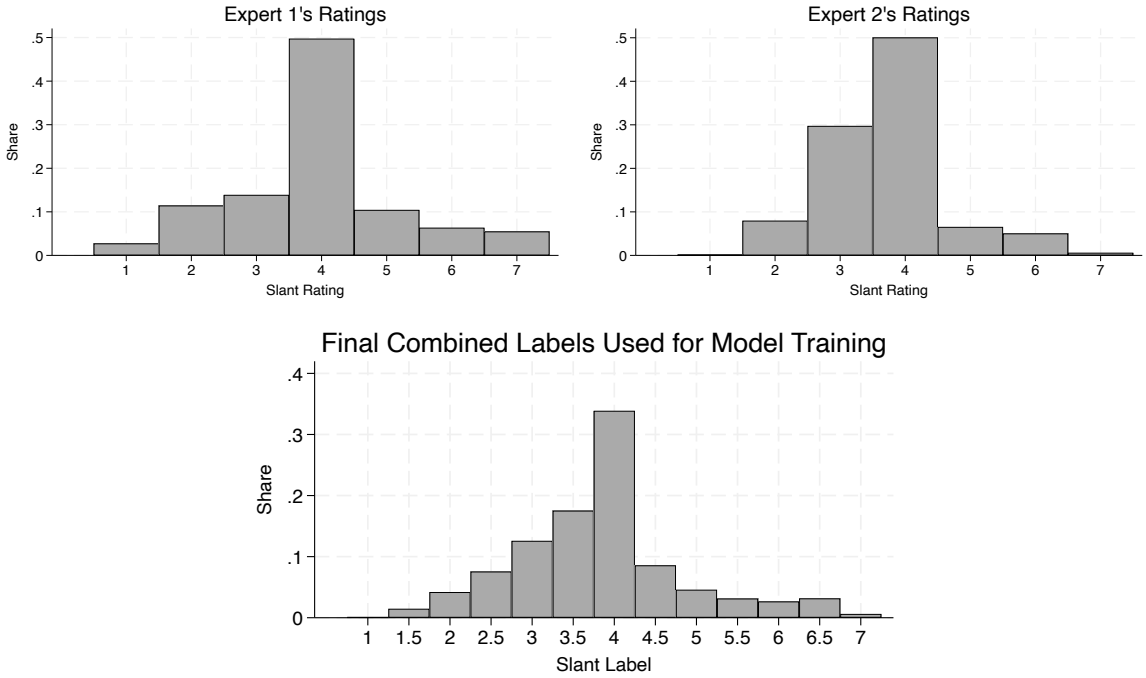
Notes: This table shows the same statistics as in Table 1, only that when constructing the statistics, we omit all articles from the 18 outlets that feature in the computation of Facebook users' political affinity scores in the URL-level Facebook Activity Dataset (see Section 3.2 for details).

Table A.9: 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	-0.26	0.13	0.39	0.08
Clicks	-0.31	0.20	0.51	0.10
Shares	-0.30	0.39	0.69	0.14

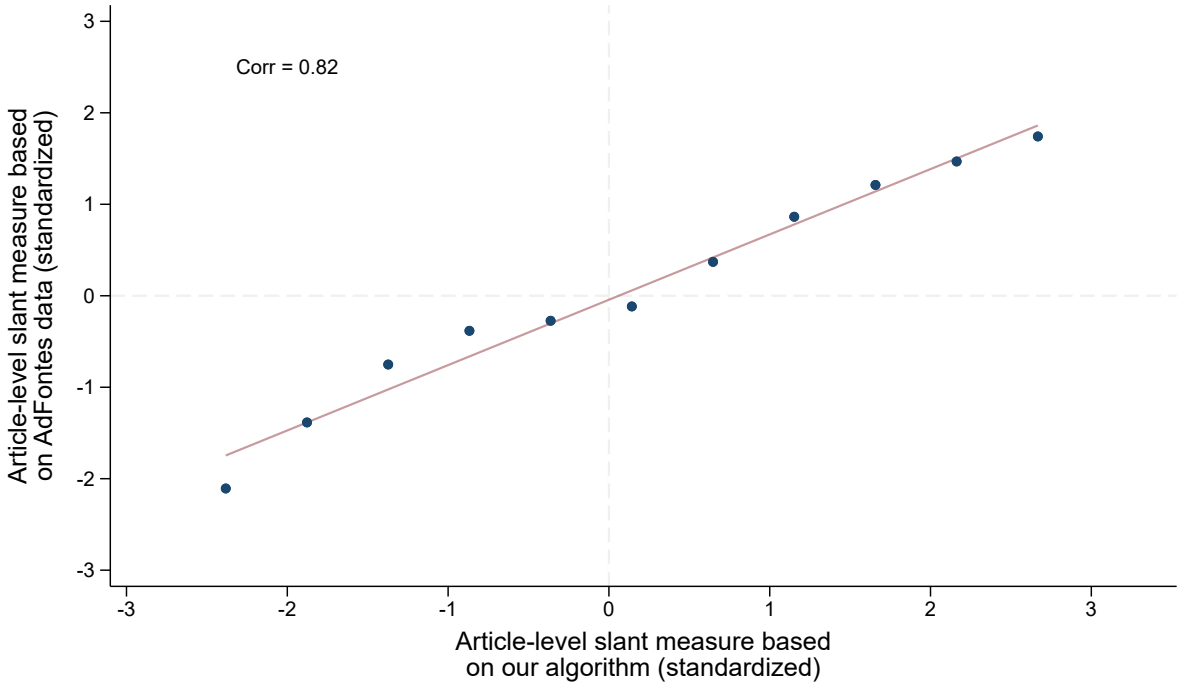
Notes: This table shows the same statistics as in Table 1 with two differences. First, when constructing the statistics, we omit all articles from the 18 outlets that feature in the computation of Facebook users' political affinity scores in the URL-level Facebook Activity Dataset (see Section 3.2 for 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



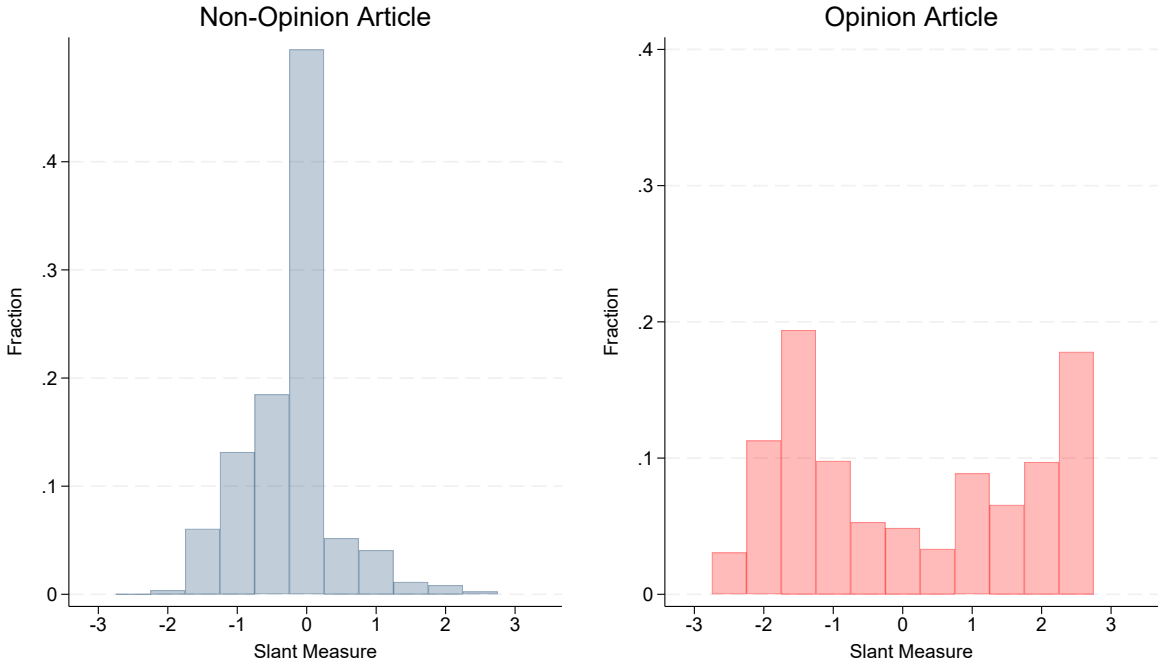
Notes: This figure shows the distribution of the raw expert slant ratings for the $N = 4,434$ articles they rated. The combined labels are created by taking the average raw score of the two raters after the RA-correction for cases where the raters did not agree on whether an article is liberal or conservative. More details are provided in Section 4.1.

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



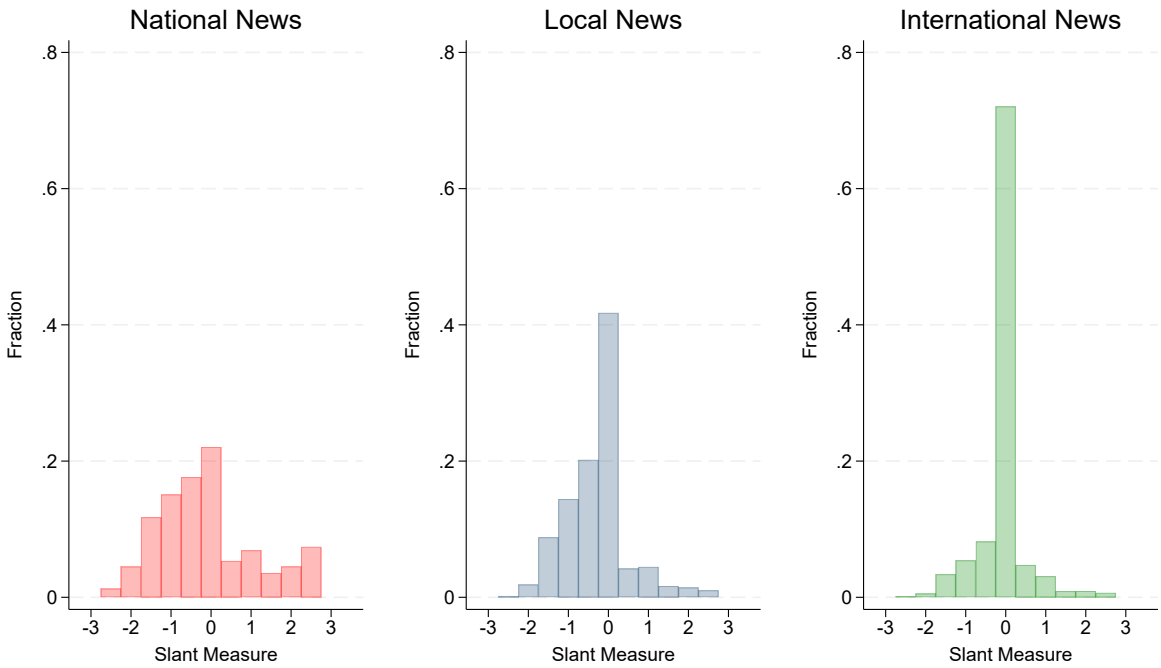
Notes: This figure shows, for each model-predicted slant bin (horizontal axis), the average third-party slant ratings (by an expert panel curated by Ad Fontes Media) across all articles in that bin. For a total of $N = 487$ articles from 2019 for which Ad Fontes Media ratings could be retrieved. To render the two scales comparable, each type of slant measure is standardized to mean 0 and standard deviation 1.

Figure A.3: Distribution of Slant: Opinion vs. Non-Opinion Articles



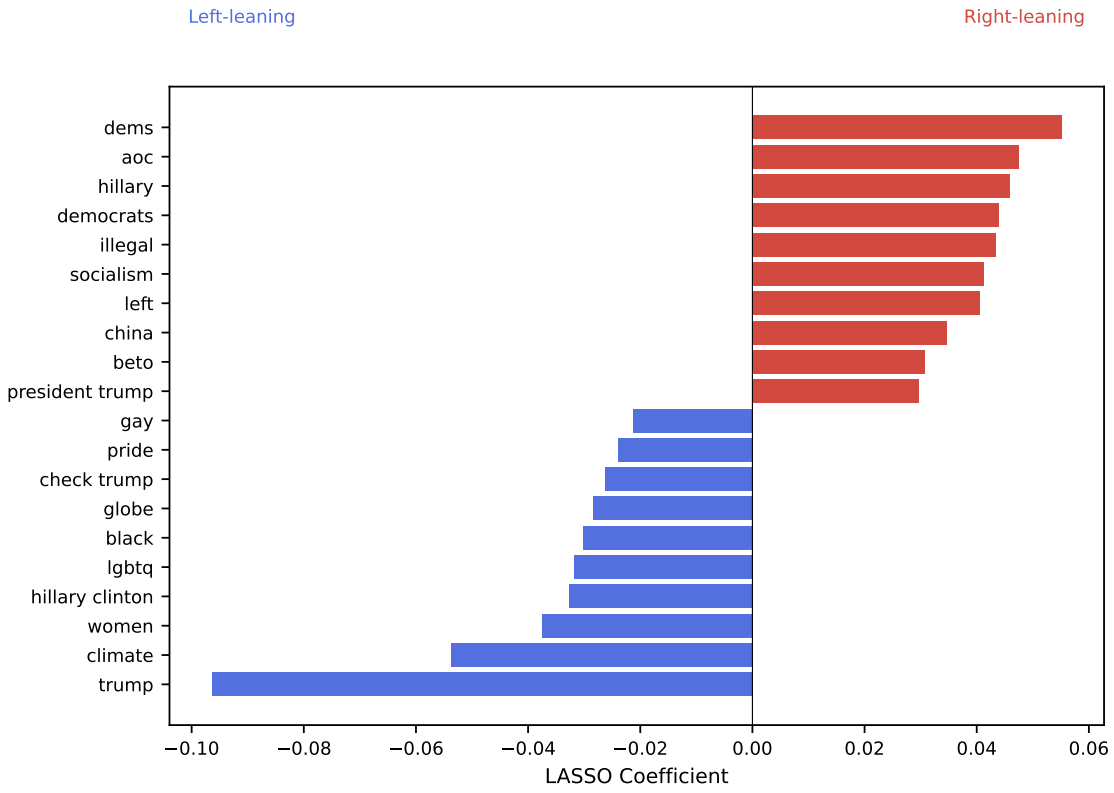
Notes: This figure shows the distribution of slant among all opinion news articles ($N = 230,744$) published online in the top-100 most-read U.S. news outlets in 2019 (right panel), and their non-opinion counterparts (left panel). Slant is given by the fine-tuned GPT-4o's prediction as described in Section 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 is described in Appendix C.

Figure A.4: Distribution of Slant: National, Local, and International News



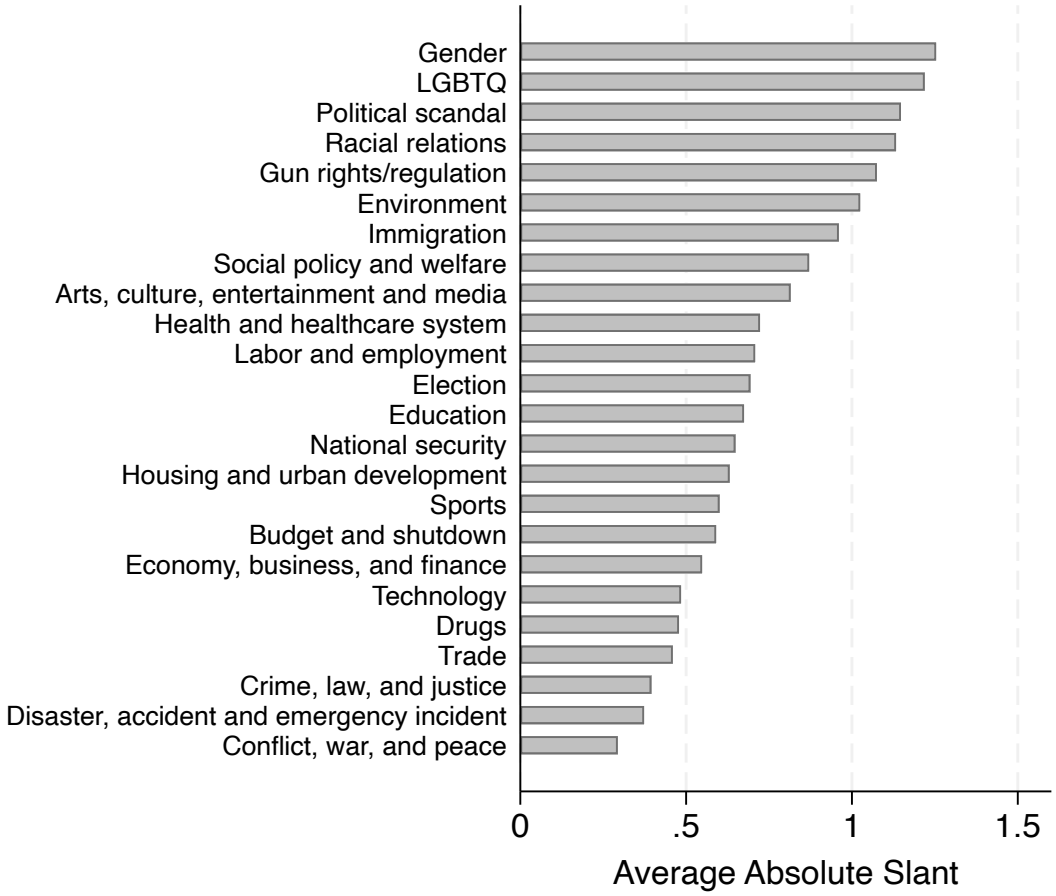
Notes: This figure shows the distribution of slant among all news articles covering national news (left panel), local news (middle panel), and international news (right panel). Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The classification into national/local/international news is described in Appendix C.

Figure A.5: Terms Most Predictive of Article Slant



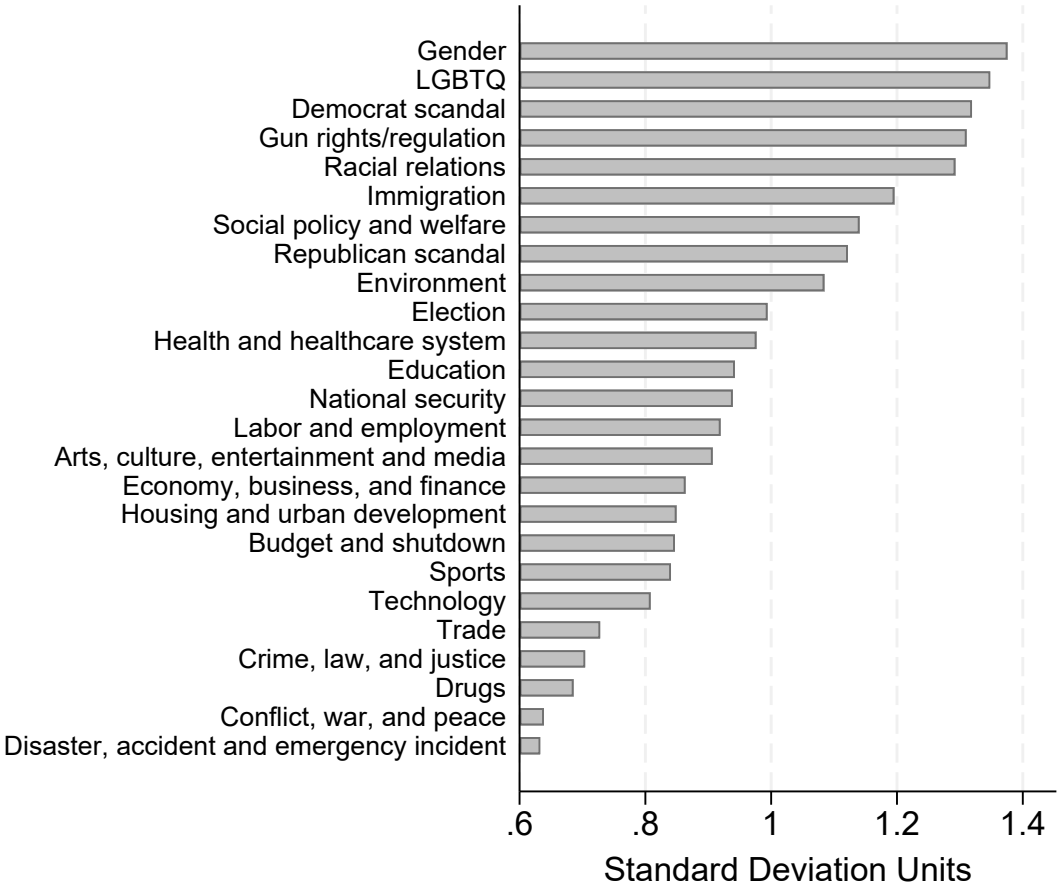
Notes: This figure presents the top ten n-grams, including both unigrams and bigrams, identified through LASSO regression as predictive of slant in news articles, based on our slant measure. These terms were extracted from article titles, with a minimum occurrence threshold of 1,000. The horizontal bar 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.6: Average Absolute Slant by Topic



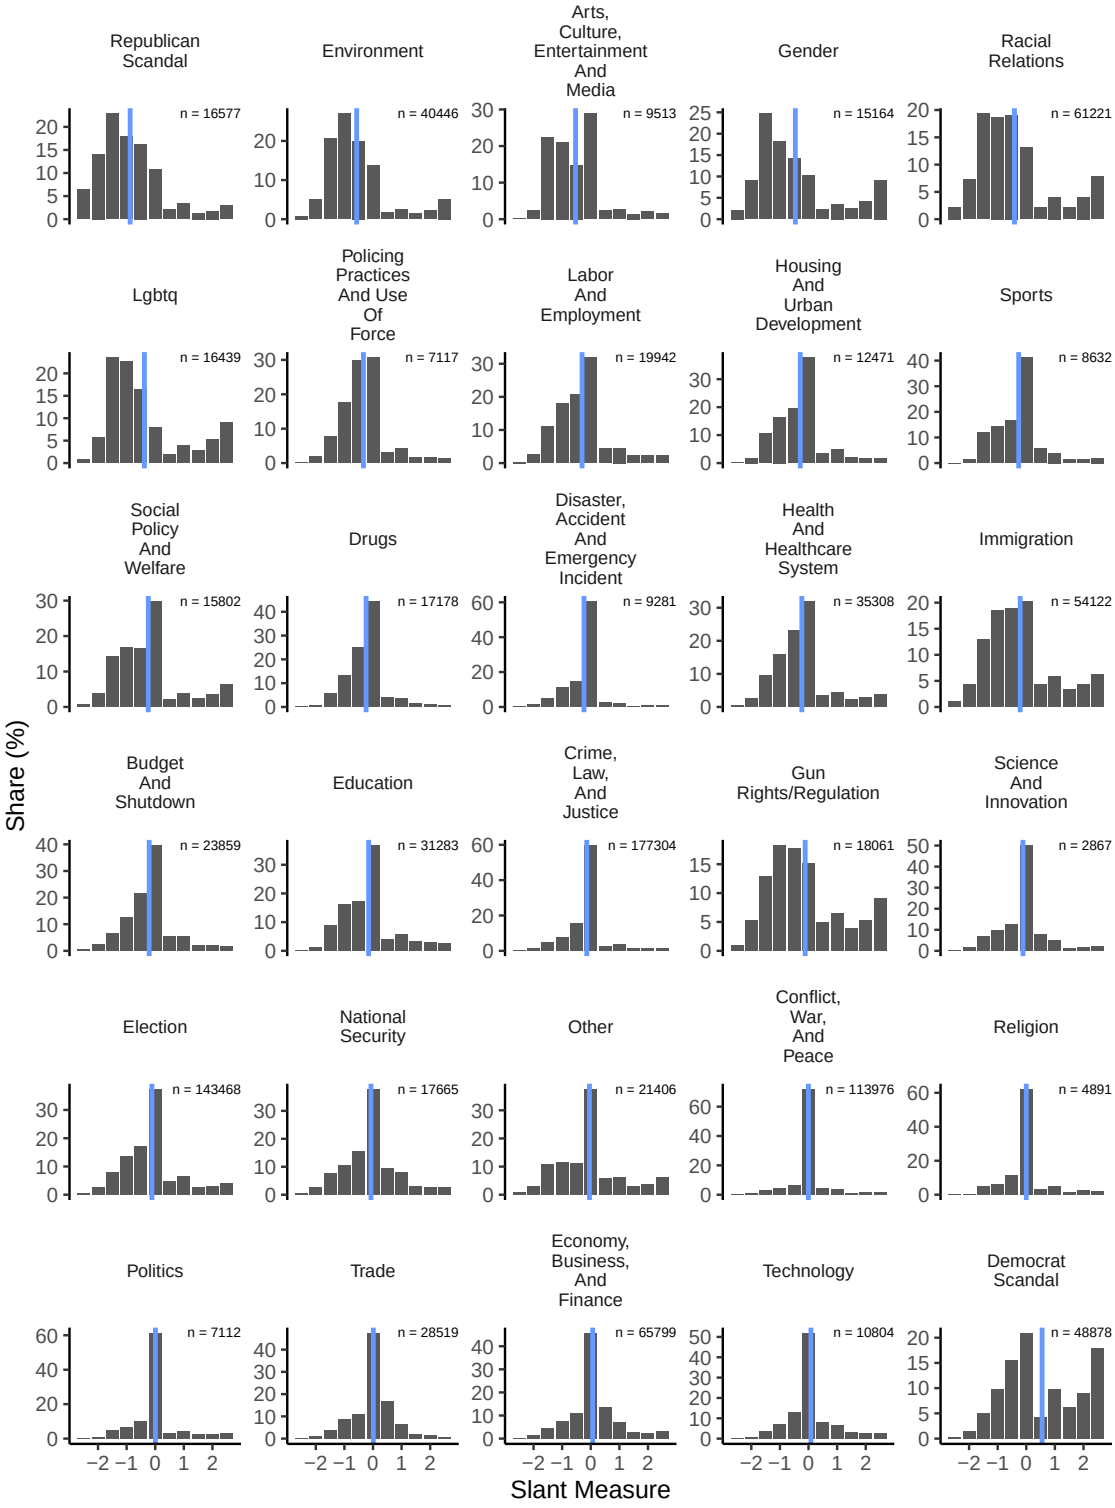
Notes: This figure shows the average absolute slant of the articles falling in each topic category listed on the vertical axis. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely 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 C.

Figure A.7: Standard Deviation of Slant by Topic



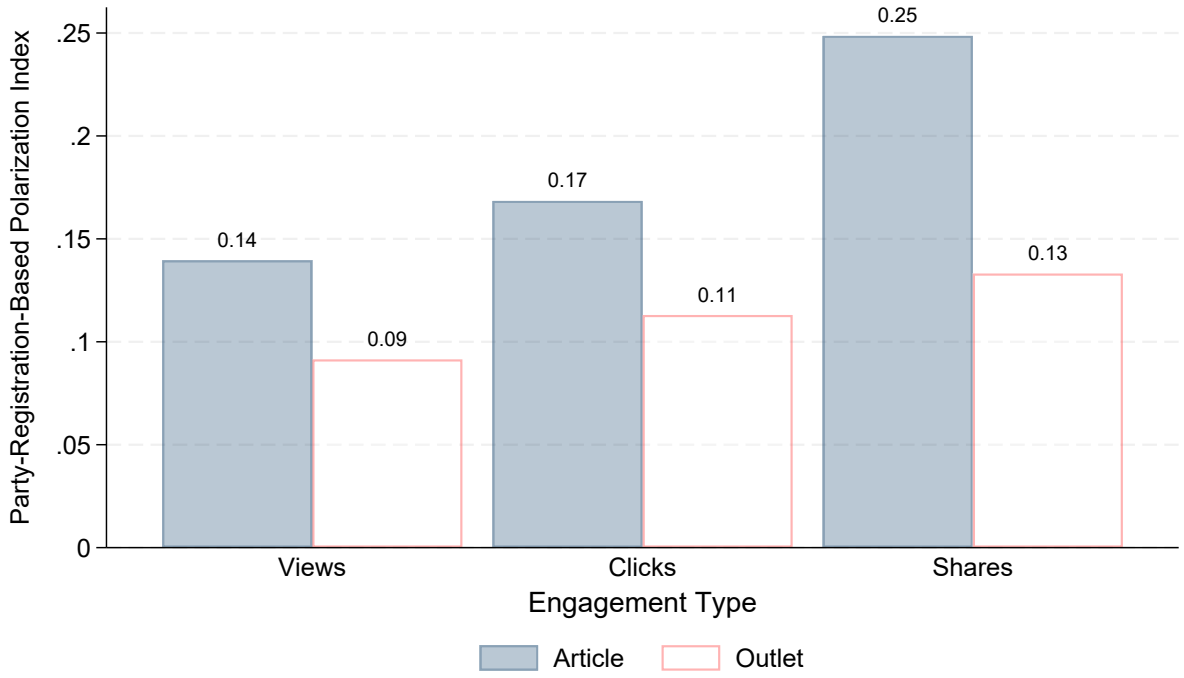
Notes: This figure shows the standard deviation of the slant of the articles falling in each topic category listed on the vertical axis. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The classification into topics is described in detail in Appendix C.

Figure A.8: Slant Distribution within Topics



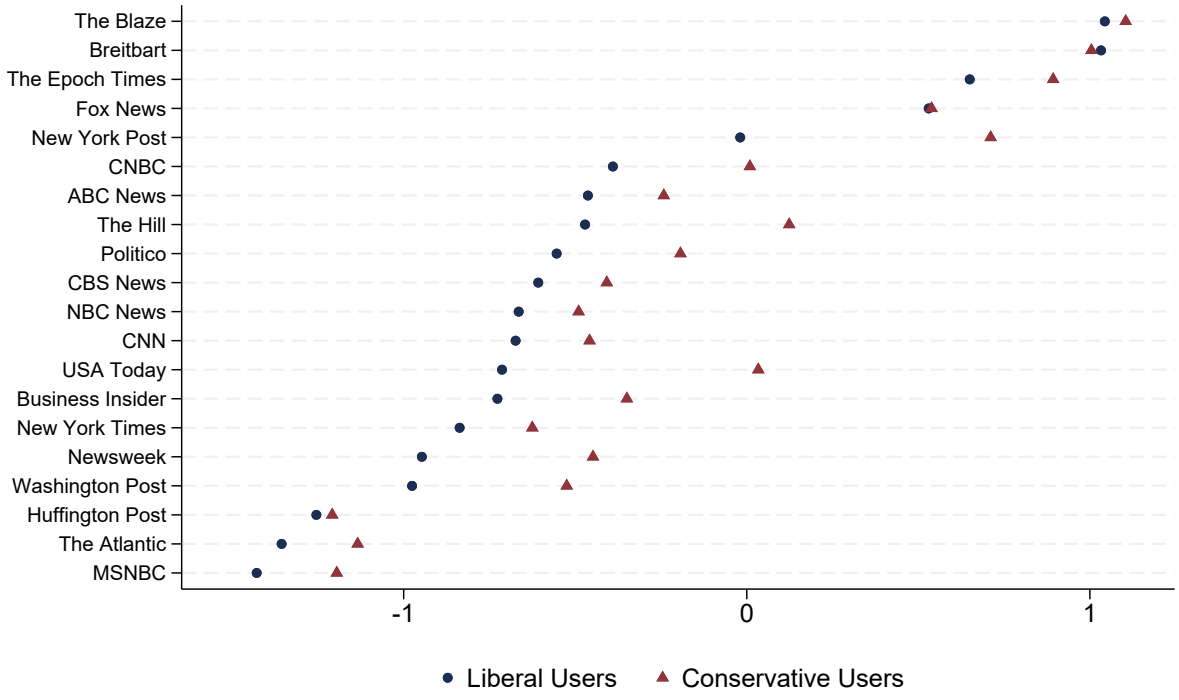
Notes: This figure shows the distribution of slant within news topics. The vertical blue line indicates the mean slant for the given topic. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from −2.5 (extremely favorable to the Democratic party) to 2.5 (extremely favorable to the Republican party). The classification into topics is described in detail in Appendix C.

Figure A.9: Polarization in News Exposure, Consumption, and Circulation on Facebook Measured at Article and Outlet Level



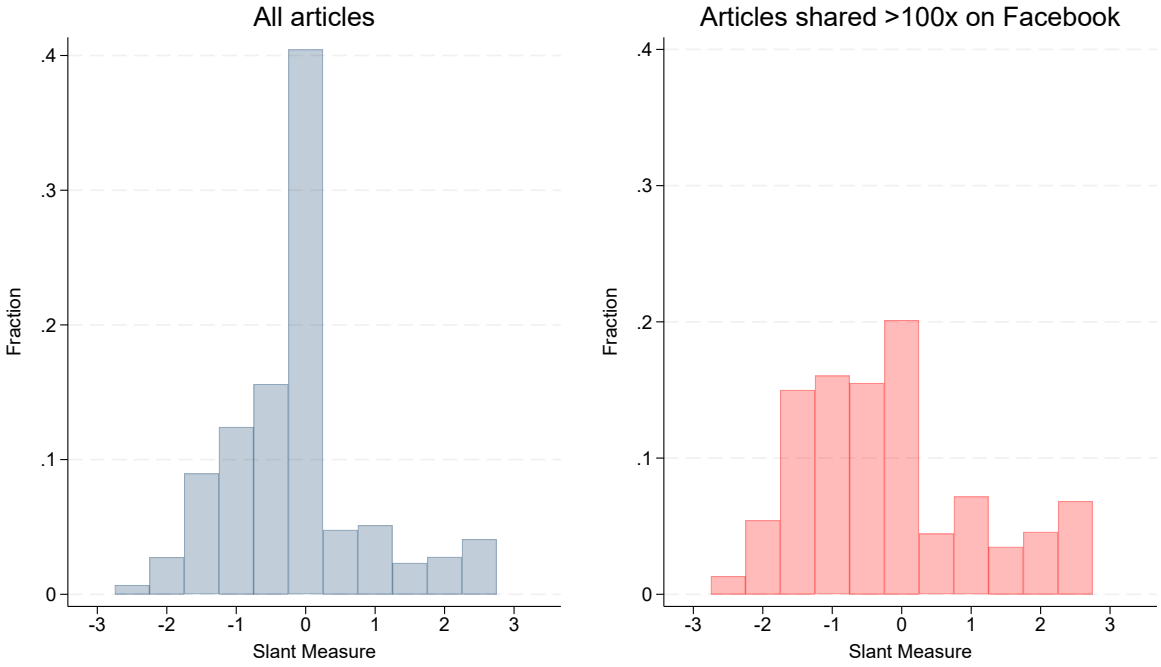
Notes: This figure shows the views-/clicks-/shares-weighted average difference in the slant of news diets by registered Republican and Democrat users on Facebook divided by 5, so that -1 indicates extreme counter-attitudinal news consumption, 0 indicates no difference in average slant across the news consumed by Democrats and Republicans, and 1 indicates extreme pro-attitudinal news consumption. The blue bars on the left show the polarization measure when slant is calculated using our main article-level slant measure (as in Table 1). The white bars on the right show the polarization measure 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 $N = 233,395$ hard 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.10: Average Slant of Articles Viewed on Facebook by Outlet and User Political Affinity



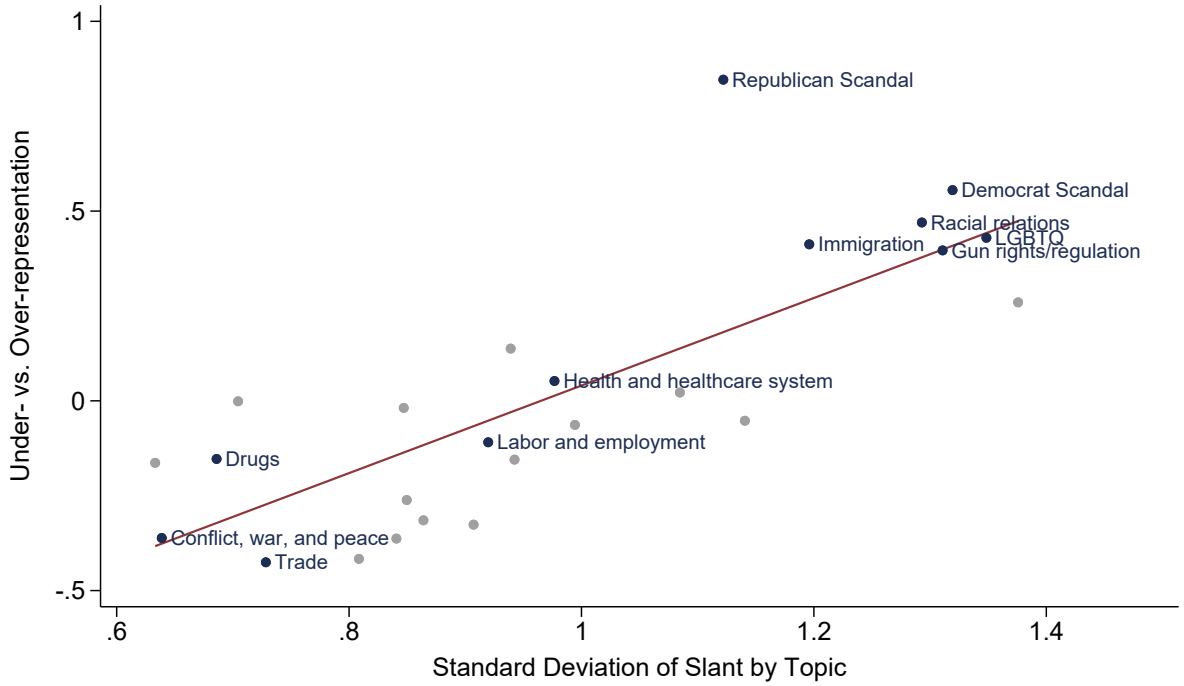
Notes: This figure shows the views-weighted average slant of articles by Facebook users' political affinity and outlet. The sample includes all hard news articles published online by the top 20 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e. shared at least 100 times on Facebook). Slant is given by the fine-tuned GPT-4o'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



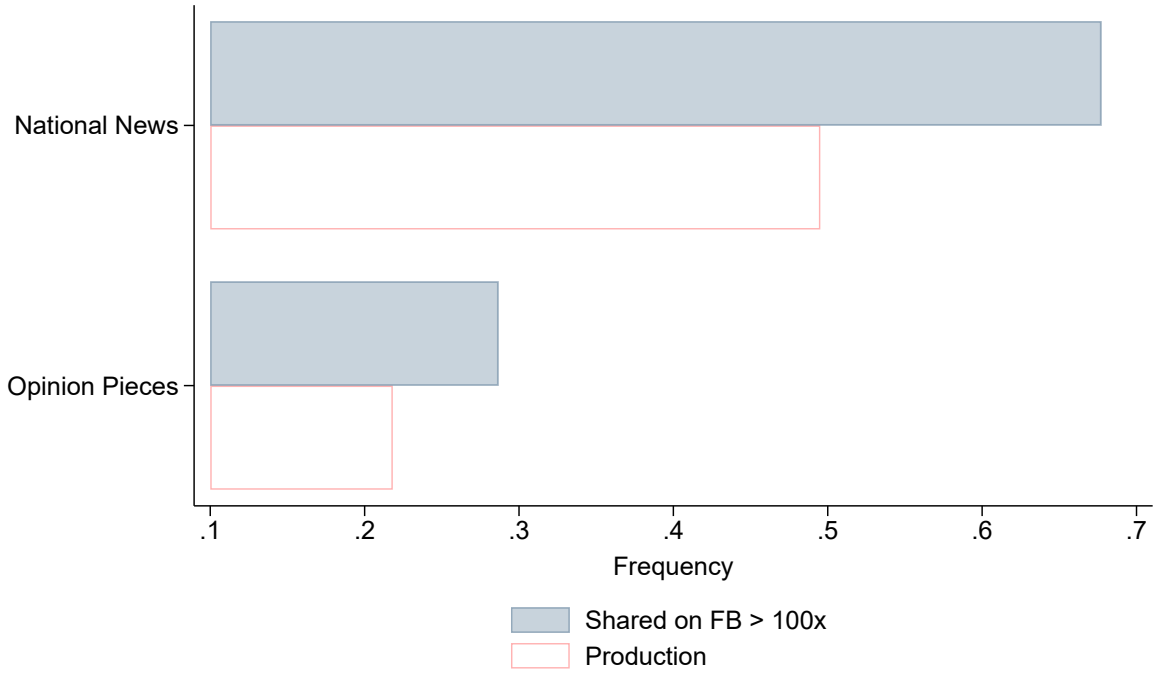
Notes: This figure shows the distribution of article-level slant for two samples: the full sample of all 1,057,124 hard news articles published online by the top-100 U.S. news outlets in 2019 (left panel) and the sub-sample of all 233,395 articles shared more than 100 times on Facebook (right panel)—i.e. the articles included in the Facebook engagement dataset (SS1). Article slant is given by the fine-tuned GPT-4o'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: Under- vs. Over-representation of Topics among Articles Shared on FB more than 100x



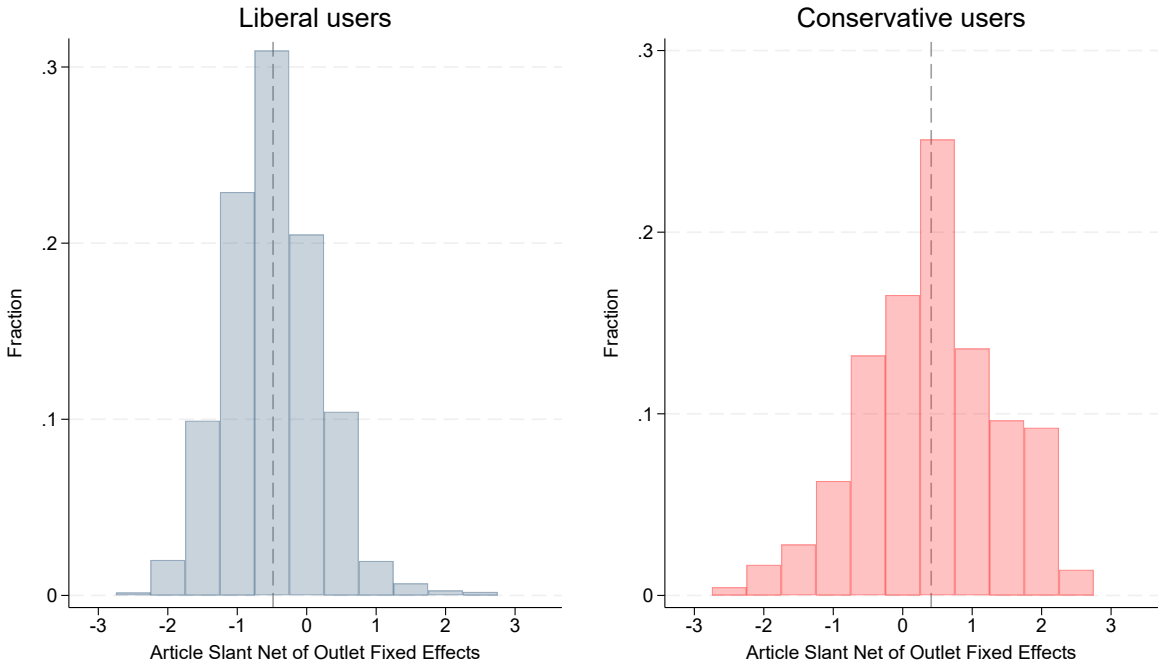
Notes: This figure shows that topics exhibiting a high standard deviation of slant tend to be over-represented in the URL-level Facebook Activity dataset. The x-axis captures the standard deviation of article-level slant for each topic. The y-axis captures the degree of under- vs over-representation. 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 C.

Figure A.13: Shares of Opinion Pieces and National News



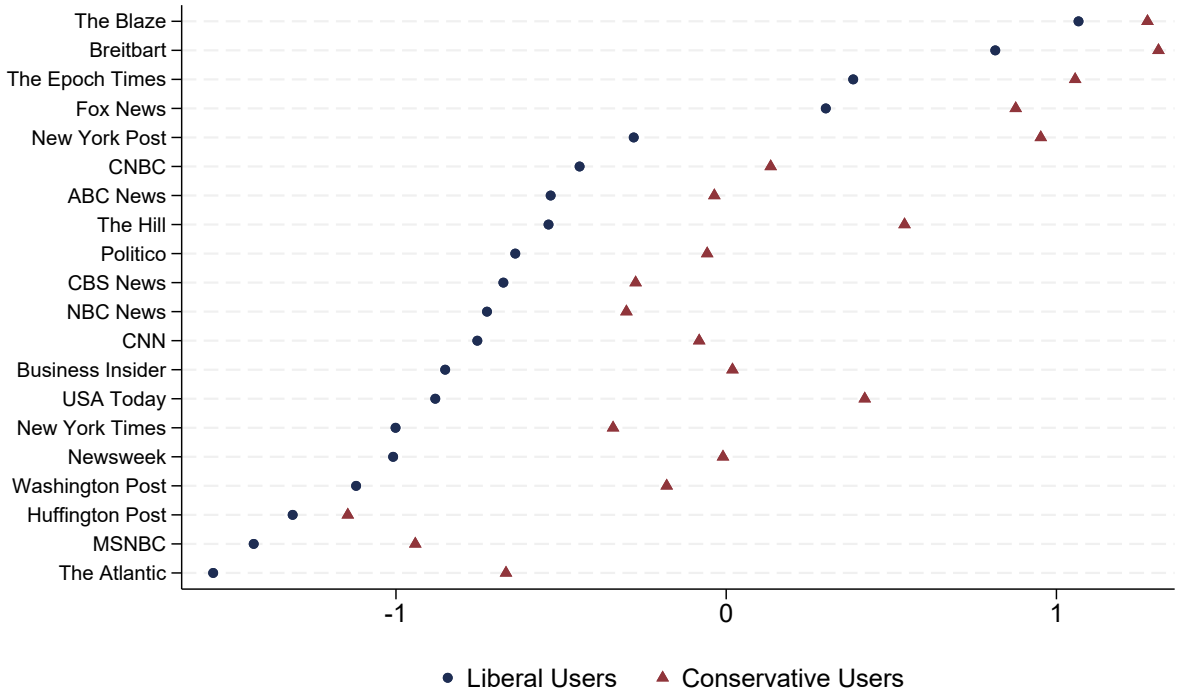
Notes: This figure presents the fraction of articles, both in our dataset ("Production") and in the URL-level Facebook Activity dataset ("Shared on FB $>$ 100x"), classified as national news and as opinion pieces. The classification into national, international, and local news, as well as the one into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

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 $N = 233,395$ hard 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 GPT-4o'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).

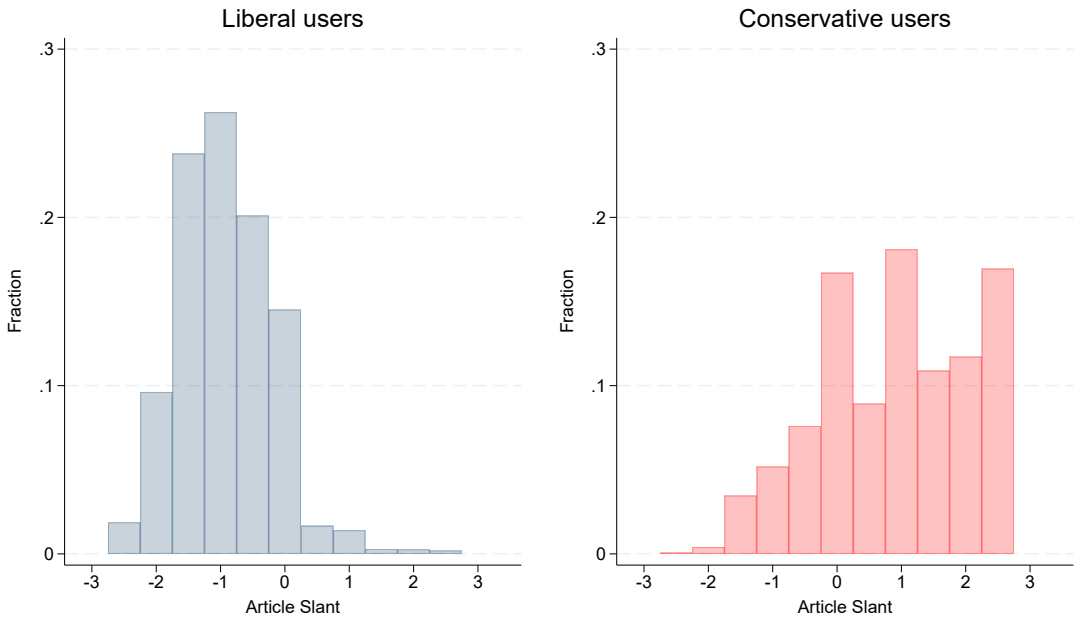
Figure A.15: Average Slant of Articles Shared on Facebook by Outlet and User Political Affinity



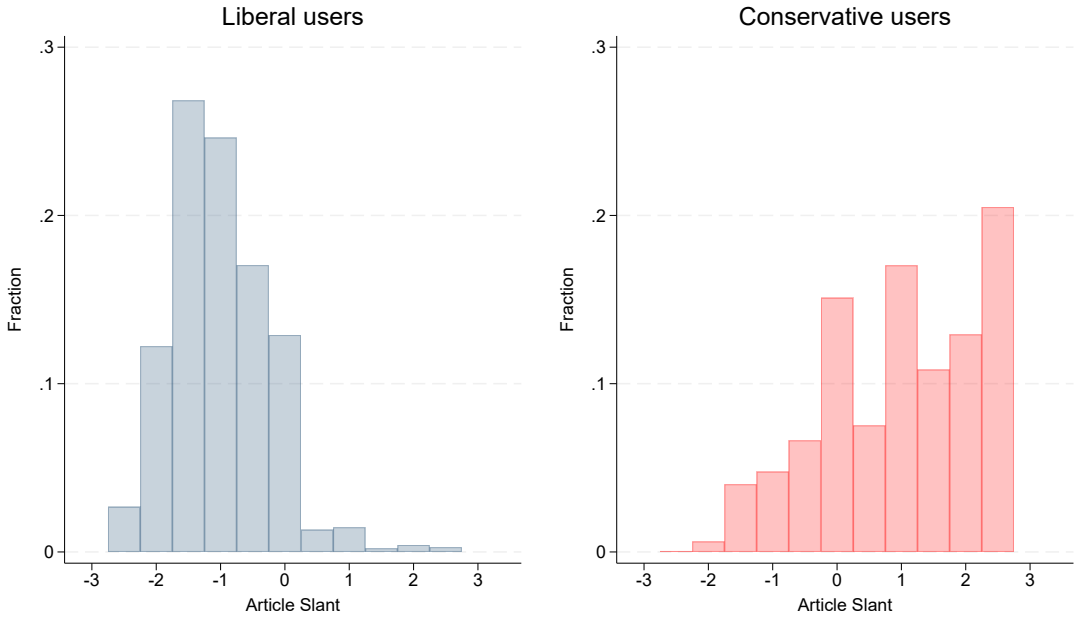
Notes: This figure shows the shares-weighted average slant of articles by Facebook users' political affinity and outlet. The sample includes all hard news articles published online by the top 20 U.S. news outlets in 2019 that are in the URL-level Facebook Activity Dataset (i.e. shared at least 100 times on Facebook). Slant is given by the fine-tuned GPT-4o'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

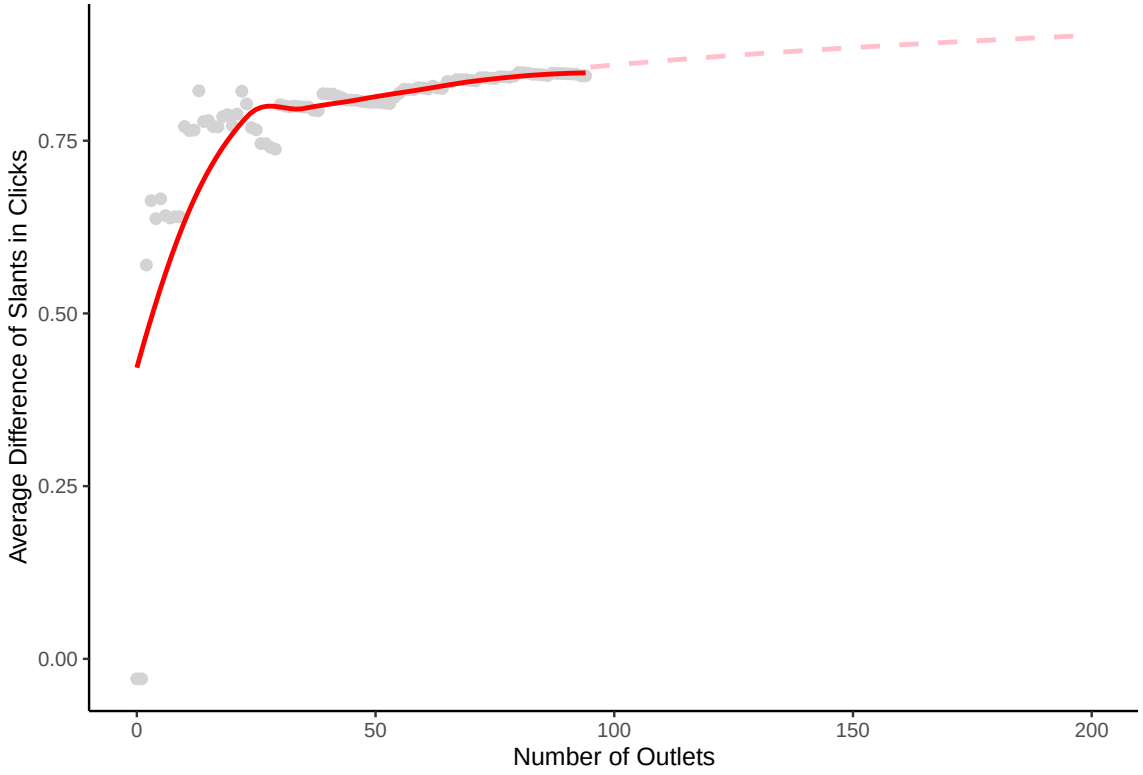


(b) Shares with Clicks



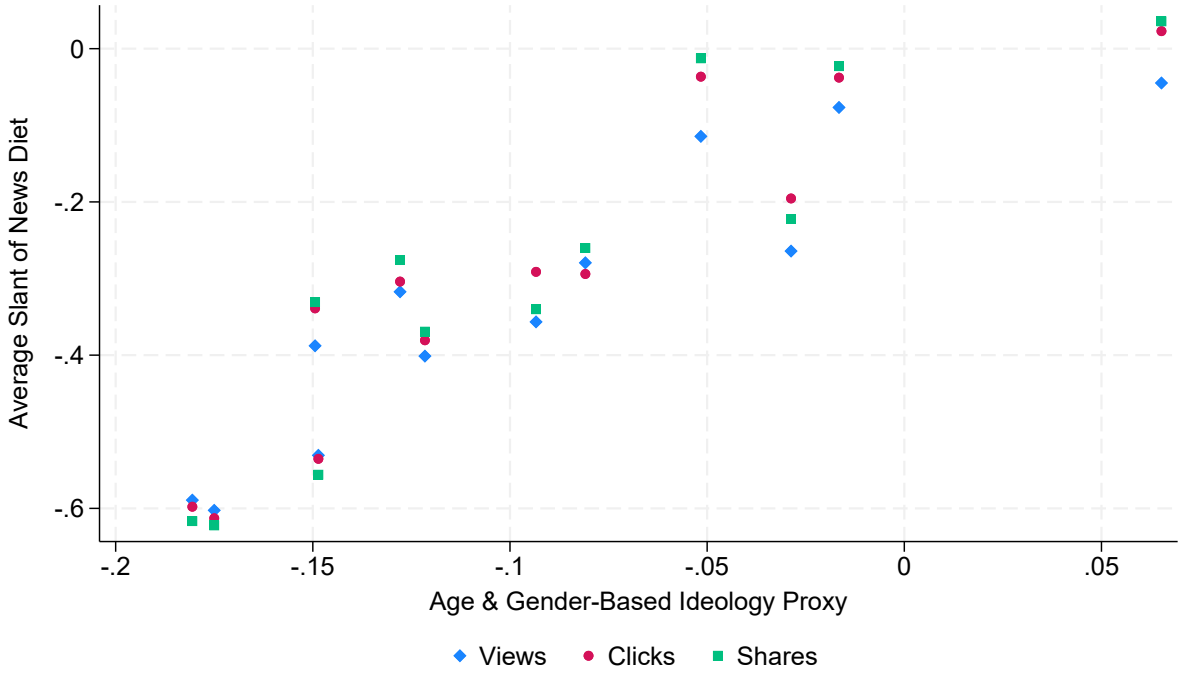
Notes: This figure shows the fraction of all shares by a given Facebook user political affinity group (liberal and conservative, respectively) that accrue to articles in a given slant bin. The figure shows the results separately for news article URLs that users clicked on before sharing on Facebook and for news article URLs that users shared without clicking. The sample includes all $N = 233,395$ hard 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).

Figure A.17: 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 Democrat users on Facebook by extending our analysis beyond top 100 U.S. news outlets as discussed in Section F.1. The news article sample includes all $N = 233,395$ hard 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).

Figure A.18: 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.

B. Labels Construction

We provided our expert human raters with the following instructions:

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

C. 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.
- Democrat 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."

D. Variance Decomposition

This appendix formalizes the variance decomposition discussed in Section 5.1.

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 $E\left(\varepsilon_{i,j,k}|O_i,A_{i,j}\right) = 0$, $E\left(A_{i,j}|O_i\right) = 0$, and $E\left(O_i\right) = 0$. Note that, as a consequence, $Var\left(\varepsilon_{i,j,k}|O_i,A_{i,j}\right) = E\left(\varepsilon_{i,j,k}^2|O_i,A_{i,j}\right)$, $Var\left(A_{i,j}|O_i\right) = E\left(A_{i,j}^2|O_i\right)$, and $Var\left(O_i\right) = E\left(O_i^2\right)$.

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

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

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

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

Note that, if we were to perform a naive variance decomposition that does not take into account measurement error, we would obtain:

$$\frac{Var\left(R_{i,j,k}\right) - E\left(Var\left(R_{i,j,k}|O_i\right)\right)}{Var\left(R_{i,j,k}\right)} =$$

$$= \frac{Var\left(O_i\right)}{Var\left(O_i\right) + E\left(Var\left(A_{i,j}|O_i\right)\right) + E\left(Var\left(\varepsilon_{i,j,k}|O_i\right)\right)}$$

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

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

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:

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

$$= \frac{Var(O_i)}{Var(O_i) + E\left(Var\left(A_{i,j}|O_i\right)\right)}$$

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.

E. 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 4.4 that our model attains at least 91% of the correlation ceiling between model predictions and expert labels.

Consider the model from the Appendix D. 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 $E\left(Z_{i,j}\right) = E\left(O_i + A_{i,j}\right) = E\left(A_{i,j}\right) = E\left(E\left(A_{i,j}|O_i\right)\right) = 0$. 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: $Var\left(\varepsilon_{i,j,1}|Z_{i,j}\right) = Var\left(\varepsilon_{i,j,2}|Z_{i,j}\right) = Var\left(\varepsilon_{i,j}|Z_{i,j}\right)$

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

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

, which, rearranging, allows us to write

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

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

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

where

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

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

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

Furthermore, it can be shown that

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

In light of the above, we can write ρ_1 as

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

By the Cauchy-Schwarz inequality, we have:

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

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

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

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 0.83.

Therefore, our model achieves at least $\frac{0.83}{0.91}=0.91$, that is 91%, of the correlation ceiling between our model prediction and our slant labels.

F. 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.

F.1 Extending Results to Additional Outlets

Our main analysis focuses on the 100 most visited U.S. outlets that cover hard news. We limited our scope to these outlets because we perceived the marginal benefit of adding more outlets to be smaller than the marginal cost of additional data collection. As discussed in section Section 3, these top 100 outlets account for 93% of all online hard-news consumption in the U.S. according to Comscore data. Thus, they provide very extensive coverage. Nonetheless, it is natural to wonder as to how our estimate of polarization in news consumption on Facebook would change if we included more outlets.

To explore this, Figure A.17 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 our polarization measure for ever-increasing subsets of these outlets.[39] Based on the trend observed, we assume a logarithmic relationship between the number of outlets included in the subset and our polarization measure. The logarithmic extrapolation suggests that expanding to 200 outlets would lead to a 7% increase over our baseline polarization measure.

F.2 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 1,637 individuals who are registered Democrats or registered Republicans. To facilitate comparability with our baseline polarization index measure, we further restrict attention to all clicks on hard news articles that i) appear in the URL-Level Facebook Activity Dataset and ii) for which

[39] We begin by only considering the most visited U.S. outlet, and then progressively expand the subset to less visited outlets.

the referrer is Facebook. We report results in column 1 of Table A.7. The table shows that news articles clicked on by registered Republicans are, on average, 0.81 units more right-slanted than those clicked on by registered Democrats. This difference gives rise to a polarization index of 0.16, which is almost identical to our preferred polarization estimate from the URL-level Facebook Activity Dataset. Thus, we replicate our headline result using a completely different dataset where partisanship is defined directly based on 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 estimate of polarization in news consumption after excluding all articles from the 18 outlets in our database that are also used in the construction of the political affinity variable. We report the results in Table A.8 and find a pattern of results virtually identical to the one described in Section 6.1. Furthermore, when omitting those 18 outlets, our measure of polarization in news consumption is 61% higher when the analysis is carried out at the article level rather than at the outlet level (see Table A.9).

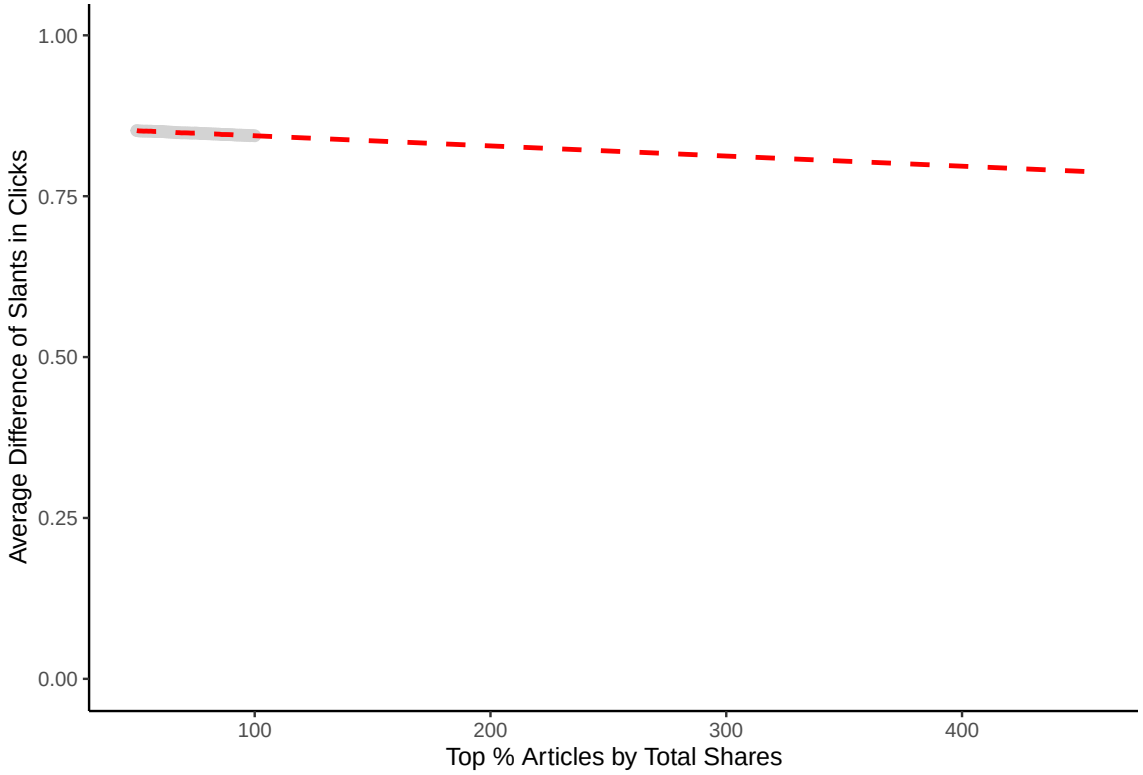
Our third robustness check involves proxying a person's political ideology using her demographic characteristics. Specifically, the URL-level Facebook Activity Dataset contains engagement numbers by age bracket and gender. Based on representative survey data from the ANES 2020-2022 Social Media Study, 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 bucket to the average slant that individuals in that bucket are exposed to and consume on Facebook. The results are shown in Figure A.18. 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.

F.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 measure derived in Section 6.1 is therefore based on the subset of hard-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 use our browsing panel data to compute polarization in news consumption for *all* hard-news articles that panelists reached through Facebook, regardless of whether those articles appear in the URL-level Facebook Activity Dataset. Using this broader set of Facebook-referred hard-news clicks, we obtain a polarization index of 0.15—a value that is only 5% smaller than the corresponding estimate that restricts attention to articles included in the URL-Level Facebook Activity Dataset. The two estimates are similar because articles shared widely on Facebook account for the large majority of Facebook hard-news consumption: in the browsing data, articles that

Figure A.19: 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 Democrat users on Facebook by total shares, as discussed in Section F.3. The news article sample includes all $N = 233,395$ hard-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).

appear in the URL-level Facebook Activity Dataset account for 81% of all hard-news article clicks through Facebook.

Second, we perform an extrapolation exercise using only the URL-level Facebook Activity Dataset. The goal of this exercise is to assess how our polarization measure 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. First, we examine how our polarization measure changes as we progressively add the least-shared articles in the URL-level Facebook Activity Dataset. Specifically, we first calculate our polarization measure after removing articles in the bottom 50% of the URL-level Facebook Activity Dataset in terms of total shares. Next, we calculate our polarization measure after excluding articles in the bottom 49% in terms of total shares. We continue in this way until we calculate our polarization measure without excluding any articles from the URL-level Facebook Activity Dataset. Figure A.19 presents the resulting pattern. The figure shows that the polarization measure calculated on articles that are shared relatively more often is generally somewhat higher, but the increase is small in magnitude. For example, excluding articles in the bottom 20% in terms of total shares increases our polarization measure by only 0.003.

In the next step, we impose assumptions that allow us to extrapolate the pattern described above and obtain an estimate of how our polarization measure would change if we 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 Figure A.19, the relationship between the percentage of articles included by total shares and our polarization measure is linear. The first assumption allows us to obtain a conservative estimate of the share of articles that are shared on Facebook fewer than 100 times; the second allows us to perform the extrapolation.

When we merge our article dataset to the URL-level Facebook Activity Dataset, we find a match rate of 22%. Thus, we know that 22% of the articles in our database are shared on Facebook at least 100 times (subject to Laplace-5 noise). We assume that the remaining articles in our database are shared on Facebook at least once, but fewer than 100 times. In light of this match rate, 100% in Figure A.19 corresponds only to 22% of articles shared on Facebook at least once. We therefore extend the figure to 455% in order to extrapolate the relationship to all articles shared at least once.

Employing the linear extrapolation described above, we find that polarization in news consumption would decrease to 0.16 if the URL-level Facebook Activity Dataset included all articles shared at least once rather than only articles shared at least 100 times. This estimate remains in the same ballpark as our main result from Section 6.1, corresponding to a 6% decrease relative to the baseline.

The intuition for this result is straightforward: articles shared fewer than 100 times account for a relatively small share of Facebook hard-news consumption. Consistent with this interpretation, in a separate analysis we study the relationship between the ranking of articles based on their shares and the number of views they receive. Using an extrapolation method similar to the one described above, we find that including all articles shared on Facebook at least once would increase the total number of views by only 26.9%.

Taken together, the browsing-panel exercise and the extrapolation based on the URL-level Facebook Activity Dataset suggest that the 100-share threshold has only a limited effect on our polarization estimates.

G. Polarization, Segregation, and Favorability

G.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 measure and the segregation (or isolation) index.[40]

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:

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

where A is the set of URLs under consideration, $\overline{C} = \sum_{\alpha \in A} C_\alpha$ indicates the total number of conservative clicks for any URL, $\overline{L} = \sum_{\alpha \in A} L_\alpha$ 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 begin by analyzing the differences between our measure of slant and the favorability scores, and then segue into the differences between our measure of polarization 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 shared 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).[41]

[40] The 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).

[41] The version of the favorability score calculated in Buntain et al. (2023) and in González-Bailón et al. (2023) are

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.[42]

Figure H.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 Figure H.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.

Figure H.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 0.82.[43] 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, Figure H.4 examines articles with favorability scores near zero (between -0.5 and 0.5) and plots the distribution of their slant values. Unlike Figure H.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.[44]

---

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.

[42]Throughout 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 measure, which are based on clicks. The results are very similar when using a shares-based favorability score.

[43]Bakshy, 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.

[44]A 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 B. 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 Brett Stephens was considered moderate by the favorability score (probably because it appeared in the New York Times), but we considered it clearly conserva-

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, a 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.[45] Our slant measure assigns to the article, independently of the title, a slant rating of zero. Conversely, the favorability score of the article varies substantially depending on the outlet that published it, as shown in the table below.

	Favorability Score	Content-based Outlet Slant
Business Insider	-0.53	-0.62
NBC News	-0.50	-0.61
ABC News	-0.09	-0.43
The New York Post	0.55	0.44

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 indeed 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 measure of polarization in news consumption and the segregation (or isolation) index. Specifically, in line with the discussion in the introduction, we can calculate both our polarization measure 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).[46] 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.

tive: 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.

[45] ABC 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."

[46] 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).

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 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 0.47. Thus, the clicks-based segregation index on moderate articles is 75% 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 fiancees, 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.

[...]"

G.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 measure of polarization 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 in news consumption at the article level will be smaller than at the outlet level. Lastly, if partisans read randomly within outlets, polarization in news consumption 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

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

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

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

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

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

Thus:

$$\begin{aligned} E[\mathbf{Y}^z] = & [2(p_1^z \cdot p_3^z + p_1^z \cdot p_4^z + p_2^z \cdot p_3^z + p_2^z \cdot p_4^z) + 4(p_1^z \cdot p_2^z) + 2(p_1^z)^2 + 2(p_2^z)^2, \\ & 2(p_1^z \cdot p_3^z + p_1^z \cdot p_4^z + p_2^z \cdot p_3^z + p_2^z \cdot p_4^z) + 4(p_3^z \cdot p_4^z) + 2(p_3^z)^2 + 2(p_4^z)^2] \end{aligned} \tag{2}$$

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

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

As a result,

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

The expression above can be simplified to:

$$E[\mathscr{P}_{article} - \mathscr{P}_{outlet}] =$$

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

To illustrate how the relationship between our measure of polarization at the outlet and at the 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 $\mathscr{P}_{article} - \mathscr{P}_{outlet} > 0$, for neutral news consumption within outlets $\mathscr{P}_{article} - \mathscr{P}_{outlet} = 0$, and for counter-attitudinal news consumption within outlets $\mathscr{P}_{article} - \mathscr{P}_{outlet} < 0$. Thus, the relationship between our measure of polarization 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

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

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

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

The segregation index at the outlet level can be written as

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

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

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

The following proposition holds.

Proposition: $E[\mathscr{S}_{article}] - E[\mathscr{S}_{outlet}] \geq 0$. Furthermore, $E[\mathscr{S}_{article}] - E[\mathscr{S}_{outlet}] = 0$ 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 $E[\mathscr{S}_{article}] - E[\mathscr{S}_{outlet}]$ as follows:

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

, where

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

and

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

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

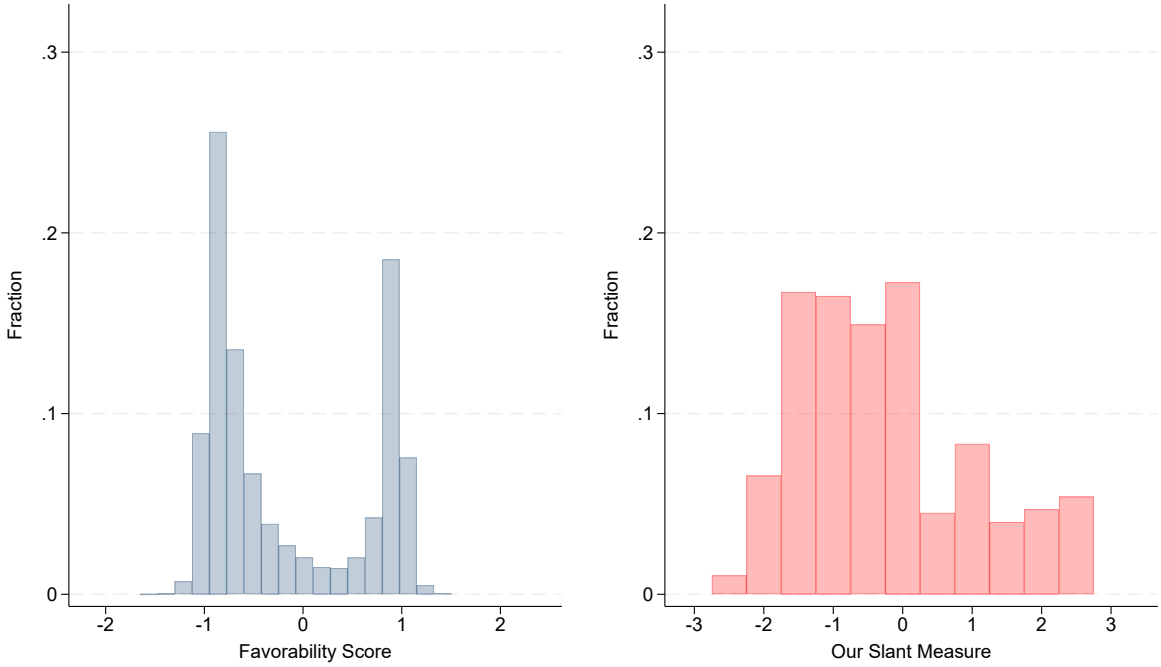
Furthermore, it is easy to see that $G_1 = 0$ if and only if $p_1^R p_2^D = p_2^R p_1^D = 0$. Similarly, it it easy to see that $G_2 = 0$ if and only if $p_4^D p_3^R = p_4^R p_3^D = 0$. Therefore, $E[\mathscr{S}_{article}] - E[\mathscr{S}_{outlet}] = 0$ if and only if $p_1^R p_2^D = p_2^R p_1^D = p_4^D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

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.[47]

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.

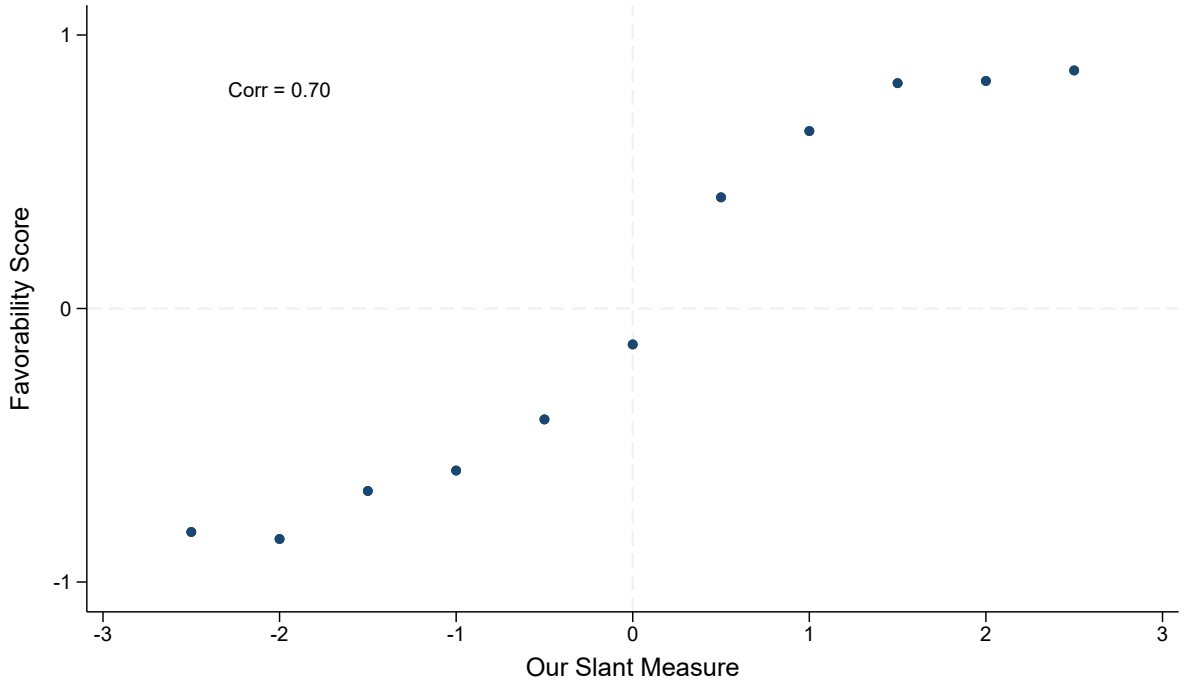
[47] The result holds generically, in the sense that the set of probability distributions $(p_1^z, p_2^z, p_3^z, p_4^z)$ for $z \in \{D, R\}$ where the result does not hold has measure zero when seen as a subset of $\mathbb{R}^8$.

Figure H.1: Distribution of Favorability Scores and Slant



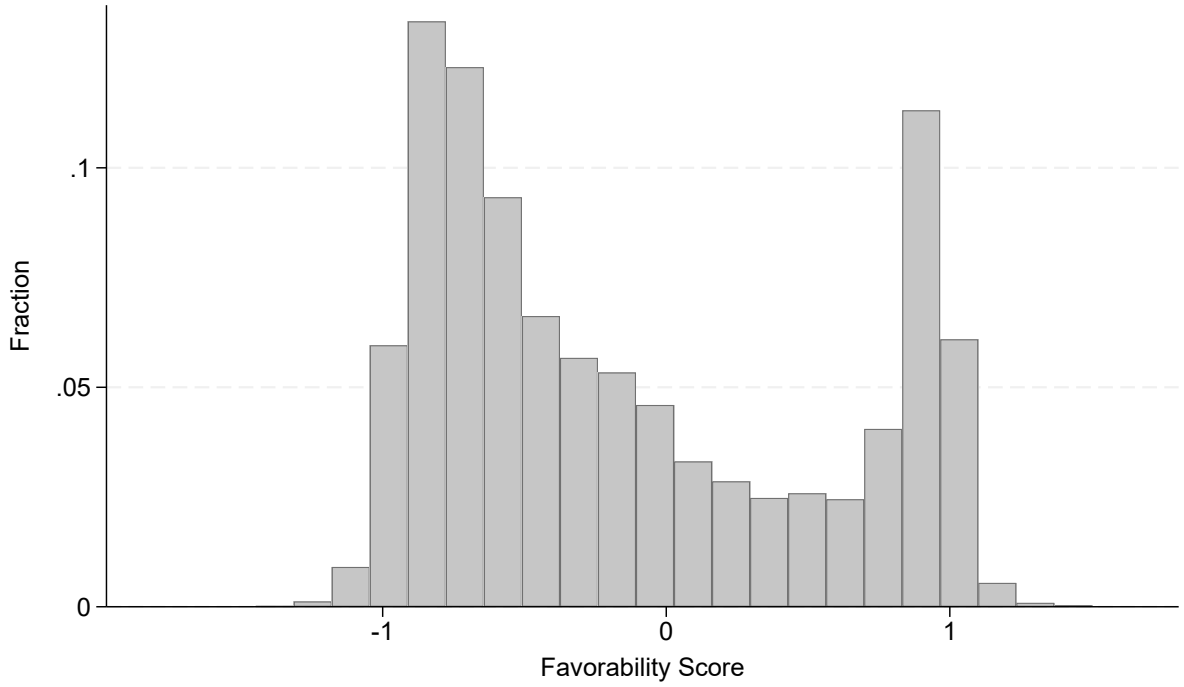
Notes: This figure shows the distribution of favorability scores (left panel) and slant (right panel) among hard news articles clicked on Facebook more than 8,962 times ($N = 38,291$) and published online by the top 100 U.S. news outlets in 2019. Our measure of slant is given by the fine-tuned GPT-4o's 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 G.

Figure H.2: Correlation between Favorability Scores and Slant



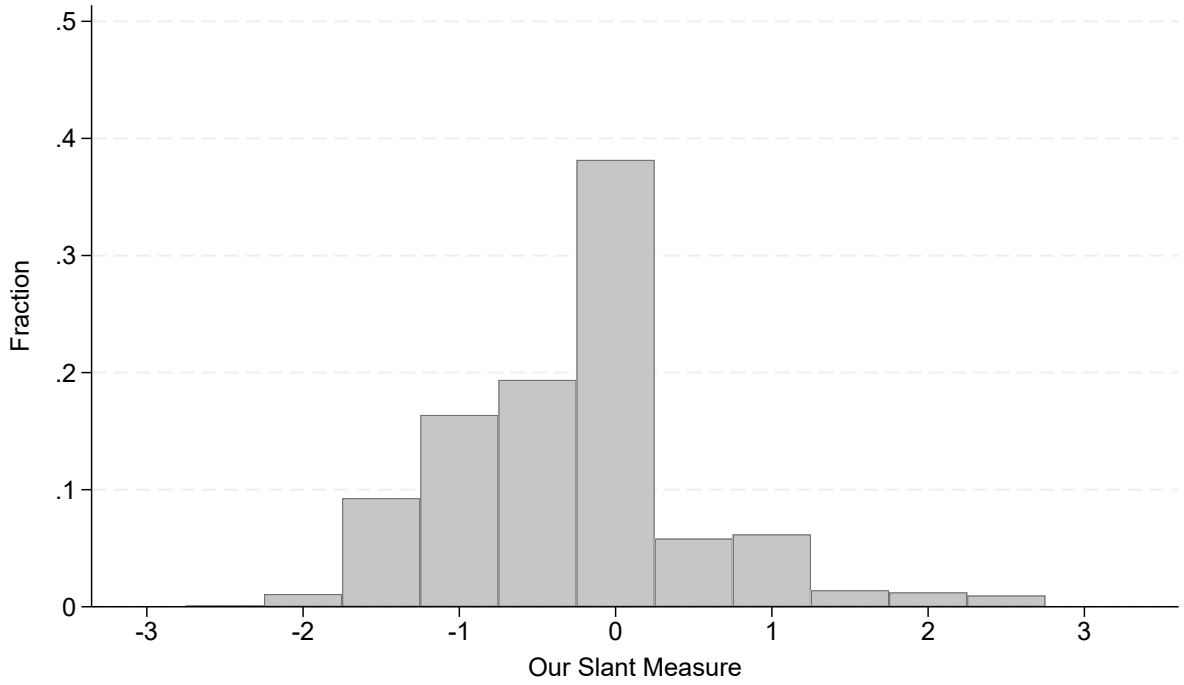
Notes: This figure shows, for each slant bin according to our slant measure (horizontal axis), the average favorability score across all articles in that bin. Our measure of slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -2.5 (extremely favorable 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 G. The sample includes all hard news articles shared on Facebook more than 8,962 times ($N = 38,291$) and published online by the top 100 U.S. news outlets in 2019.

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



Notes: This figure shows the distribution of favorability scores among hard 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 GPT-4o's prediction as described in Section 4.2. The construction of the favorability scores is described in detail in Appendix G. The sample includes all hard news articles shared on Facebook more than 8,962 times ($N = 38,291$) and published online by the top 100 U.S. news outlets in 2019.

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



Notes: This figure shows the distribution of our measure of slant among hard 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 GPT-4o'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 G. The sample includes all hard news articles shared on Facebook more than 8,962 times ($N = 38,291$) and published online by the top 100 U.S. news outlets in 2019.

H. 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 J.1: Summary Statistics within Outlets

	Obs.		Mean SS1 Slant			Sd.	% Lib
Outlet	All	SS1	All	Lib	Cons	Slant	Visits
6abc	3,630	234	-0.33	-0.33	-0.26	0.53	47.09
9news	3,168	63	-0.26	-0.45	-0.07	0.58	58.51
ABC News	21,008	2,156	-0.37	-0.46	-0.24	0.54	61.20
Abc13	3,567	305	-0.30	-0.30	-0.30	0.52	42.73
Abc7	4,754	431	-0.33	-0.34	-0.28	0.55	59.12
Ajc	11,602	868	-0.57	-0.71	-0.36	0.69	61.17
Al	3,974	1,058	-0.40	-0.55	-0.23	0.89	47.47
American Thinker	7,040	2,577	2.36	2.34	2.37	0.45	9.93
Azcentral	4,081	653	-0.69	-0.83	-0.48	0.92	63.86
Baynews9	12,636	85	-0.33	-0.33	-0.21	0.61	49.62
Bloomberg	6,343	1,415	-0.50	-0.63	-0.21	0.70	75.65
Boston	5,692	134	-0.74	-0.80	-0.73	0.67	80.91
Boston Globe	17,329	802	-0.96	-1.05	-0.80	0.86	87.58
Breitbart	40,989	15,510	1.04	1.03	1.00	1.21	2.31
Buffalonews	2,776	184	-0.39	-0.63	0.08	0.84	59.56
Business Insider	21,420	4,444	-0.58	-0.73	-0.36	0.66	69.63
Buzzfeed	1,284	372	-1.51	-1.49	-1.43	0.71	76.10
CBS News	10,474	5,196	-0.52	-0.61	-0.41	0.56	64.74
CNBC	15,060	3,616	-0.24	-0.39	0.01	0.64	63.31
CNN	24,749	14,841	-0.62	-0.67	-0.46	0.66	82.19
Cbslocal	3,785	2,446	-0.29	-0.44	-0.15	0.58	52.16
Chicago Tribune	20,643	1,569	-0.74	-0.95	-0.27	1.02	72.07
Cincinnati	2,036	262	-0.29	-0.41	0.03	0.84	60.36
Cleveland	7,978	461	-0.56	-0.69	-0.38	0.72	69.35
Click2Huston	9,611	291	-0.23	-0.32	-0.20	0.48	35.33
Clickondetroit	11,619	404	-0.32	-0.37	-0.29	0.46	56.18
Dallasnews	4,782	767	-0.44	-0.73	0.06	0.91	63.65
Detroitnews	11,019	827	-0.02	-0.42	0.64	0.97	54.99
Forbes	2,654	1,042	-0.59	-0.89	0.01	1.08	66.67
Fox News	34,865	15,467	0.57	0.53	0.54	1.10	4.92
Foxbusiness	7,827	1,298	0.61	0.62	0.60	0.88	5.99
Freep	2,728	838	-0.61	-0.64	-0.59	0.71	73.11

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

	Obs.		Mean SS1 Slant			Sd.	% Lib
Outlet	All	SS1	All	Lib	Cons	Slant	Visits
Heavy	2,390	186	-0.71	-1.24	0.25	0.89	74.10
Hot Air	20,172	1,303	1.62	1.59	1.58	1.00	4.67
Houston Chronicle	29,725	249	-0.49	-0.65	-0.20	0.67	70.60
Huffington Post	13,636	9,520	-1.25	-1.25	-1.21	0.63	91.39
Indystar	1,571	274	-0.40	-0.61	0.09	0.71	57.48
Inquirer	5,836	605	-0.91	-1.12	-0.28	0.83	74.07
Jsonline	2,049	440	-0.49	-0.56	-0.32	0.64	71.24
Khou	3,078	123	-0.23	-0.30	-0.11	0.59	45.52
Ksl	11,439	90	-0.11	-0.27	0.03	0.66	46.88
Ktla	6,024	898	-0.45	-0.47	-0.43	0.56	57.03
LA Times	14,702	3,277	-0.99	-1.12	-0.54	0.80	82.08
MSNBC	8,389	3,999	-1.38	-1.43	-1.20	0.61	94.07
Masslive	5,242	248	-0.49	-0.64	-0.33	0.59	60.92
Mediaite	8,459	708	-0.49	-0.70	-0.30	1.13	56.45
Mlive	4,658	1,013	-0.43	-0.48	-0.32	0.63	58.27
NBC News	12,173	7,667	-0.62	-0.67	-0.49	0.65	79.03
National Review	5,923	1,584	1.47	0.99	1.51	1.12	14.79
New York Post	15,714	4,869	0.38	-0.02	0.71	1.11	31.57
New York Times	47,274	14,311	-0.82	-0.84	-0.63	0.75	90.87
Newsday	9,764	157	-0.40	-0.50	-0.33	0.79	49.48
Newser	3,627	10	-1.20	-1.18	-0.80	0.75	67.31
Newsmax	10,775	228	1.49	0.99	1.50	0.99	11.98
Newsweek	7,477	5,038	-0.86	-0.96	-0.47	0.87	81.85
Nj	6,861	886	-0.45	-0.60	-0.35	0.81	57.08
Nola	3,168	373	-0.70	-0.87	-0.57	0.74	56.41
Nydailynews	12,427	2,695	-0.73	-0.84	-0.37	0.88	78.64
Omaha	3,799	72	-0.44	-0.59	0.02	0.77	58.82
Oregonlive	5,110	845	-0.53	-0.64	-0.40	0.75	68.09
Palmerreport	4,928	2,409	-2.15	-2.13	-2.39	0.41	87.98
Patch Media	23,652	1,541	-0.37	-0.41	-0.37	0.65	54.71
Pennlive	4,990	424	-0.28	-0.59	0.28	0.82	40.47
Pjmedia	5,491	2,492	2.06	1.99	2.04	0.89	8.09
Politico	12,159	5,431	-0.50	-0.55	-0.19	0.72	82.92
Raw Story	19,980	7,888	-1.27	-1.30	-0.94	0.66	93.40
Realclearpolitics	1,465	248	2.18	0.31	2.33	1.15	7.21
San Francisco Gate	35,052	280	-0.44	-0.60	0.02	0.64	75.32
Seattle Times	16,069	695	-0.77	-0.91	-0.08	0.78	85.13
Slate	4,637	1,802	-1.44	-1.47	-1.30	0.77	95.06

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

	Obs.		Mean SS1 Slant			Sd.	% Lib
Outlet	All	SS1	All	Lib	Cons	Slant	Visits
Star Tribune	26,290	480	-0.45	-0.70	0.03	0.80	62.04
Staradvertiser	2,339	128	-0.48	-0.48	-0.21	0.59	81.08
Stltoday	8,862	411	-0.53	-0.61	-0.37	0.78	68.09
Syracuse	4,557	220	-0.25	-0.38	0.02	0.68	43.29
The Atlantic	3,855	2,167	-1.35	-1.36	-1.14	0.88	90.45
The Blaze	7,473	4,418	1.10	1.04	1.10	1.09	3.72
The Epoch Times	14,578	4,017	0.76	0.65	0.89	0.85	17.35
The Hill	43,972	14,912	-0.36	-0.47	0.12	0.85	79.48
Thedailybeast	4,302	1,948	-1.04	-1.08	-0.79	0.86	91.29
Thefederalist	3,379	1,404	2.05	1.72	2.09	0.72	8.81
Time	6,393	2,105	-0.86	-0.95	-0.60	0.69	80.69
Townhall	12,394	4,358	1.75	1.67	1.74	0.95	2.73
Twitchy	9,353	722	2.20	-0.21	2.30	0.56	2.75
USA Today	13,467	5,413	-0.44	-0.71	0.04	0.88	61.68
Voanews	13,203	158	-0.84	-0.84	-0.44	0.60	76.61
Vox	5,480	2,978	-1.27	-1.29	-1.11	0.64	92.26
WND	14,344	2,045	1.86	1.72	1.86	0.83	4.12
WRAL	13,548	293	-0.32	-0.44	-0.13	0.63	43.49
WTOP	19,579	165	-0.61	-0.67	-0.34	0.54	69.86
Washington Examiner	17,829	4,056	0.86	0.33	0.92	1.03	7.98
Washington Post	38,418	14,636	-0.92	-0.98	-0.52	0.87	87.65
Washington Times	20,170	1,362	0.92	0.75	0.95	1.03	8.96
Wfaa	2,943	111	-0.17	-0.29	-0.11	0.67	47.10

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 user with a conservative political affinity score, the standard deviation of slant, and the percentage of liberal visits defined as the percentage of visits by liberal users divided by the sum of the percentage of visits by liberal and conservative users within 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 J.1: Slant Distribution within Outlets

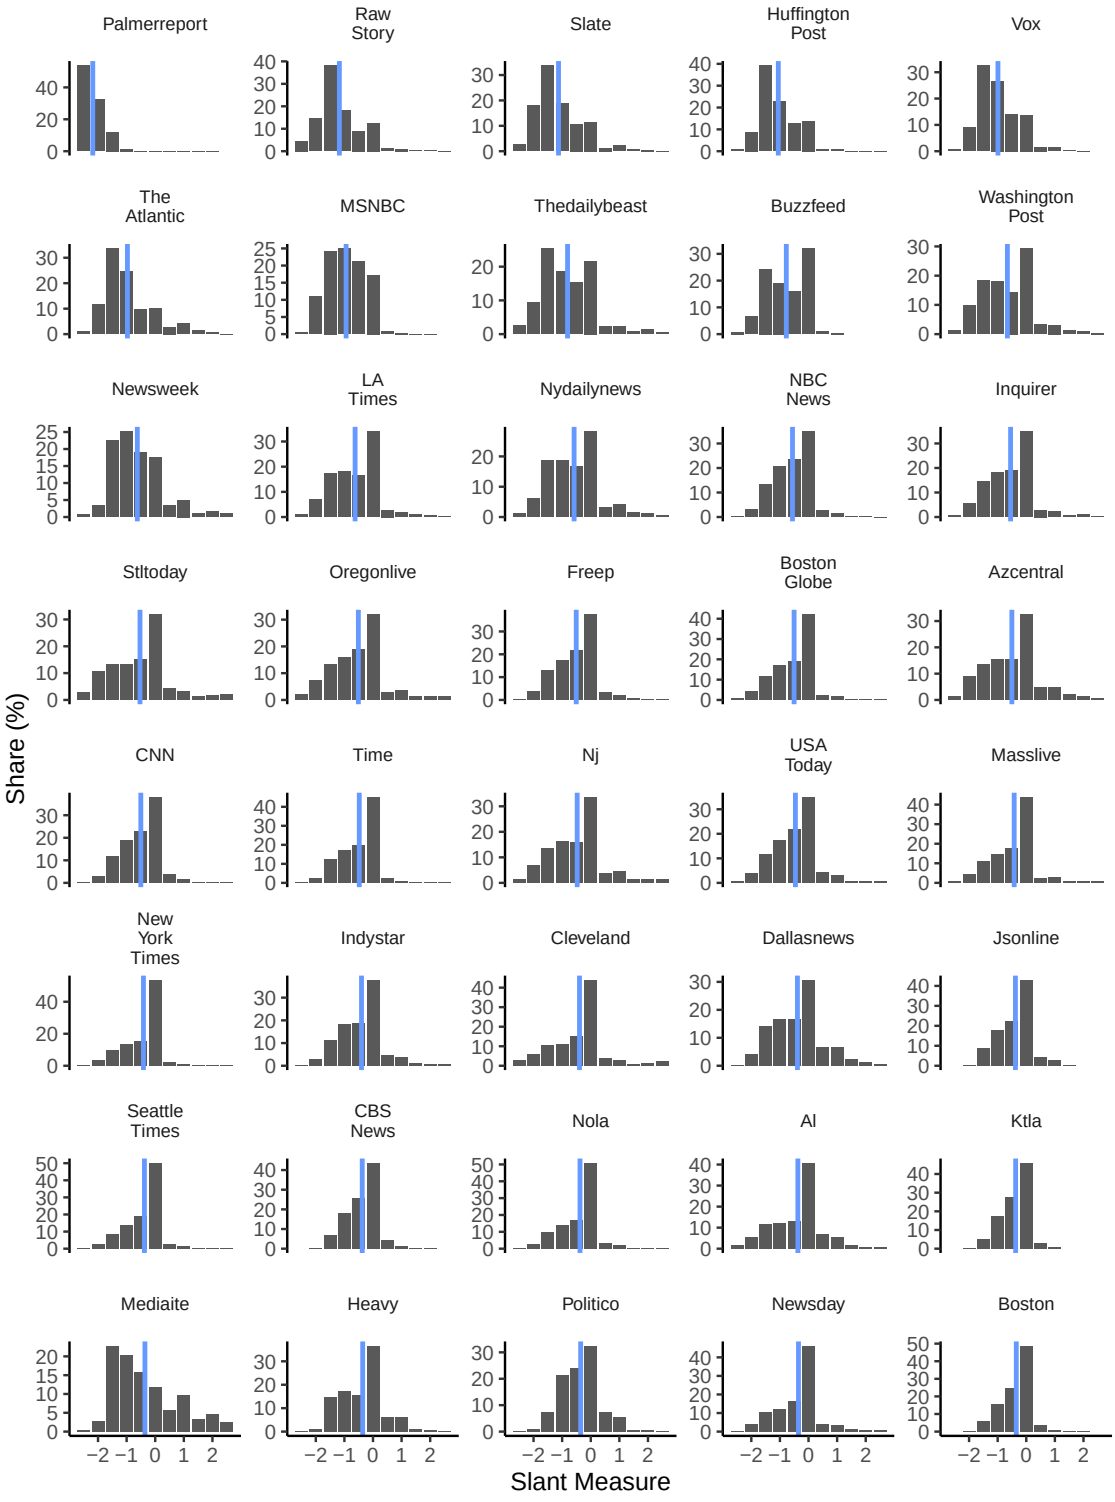


Figure J.1: Slant Distribution within Outlets (Continued)

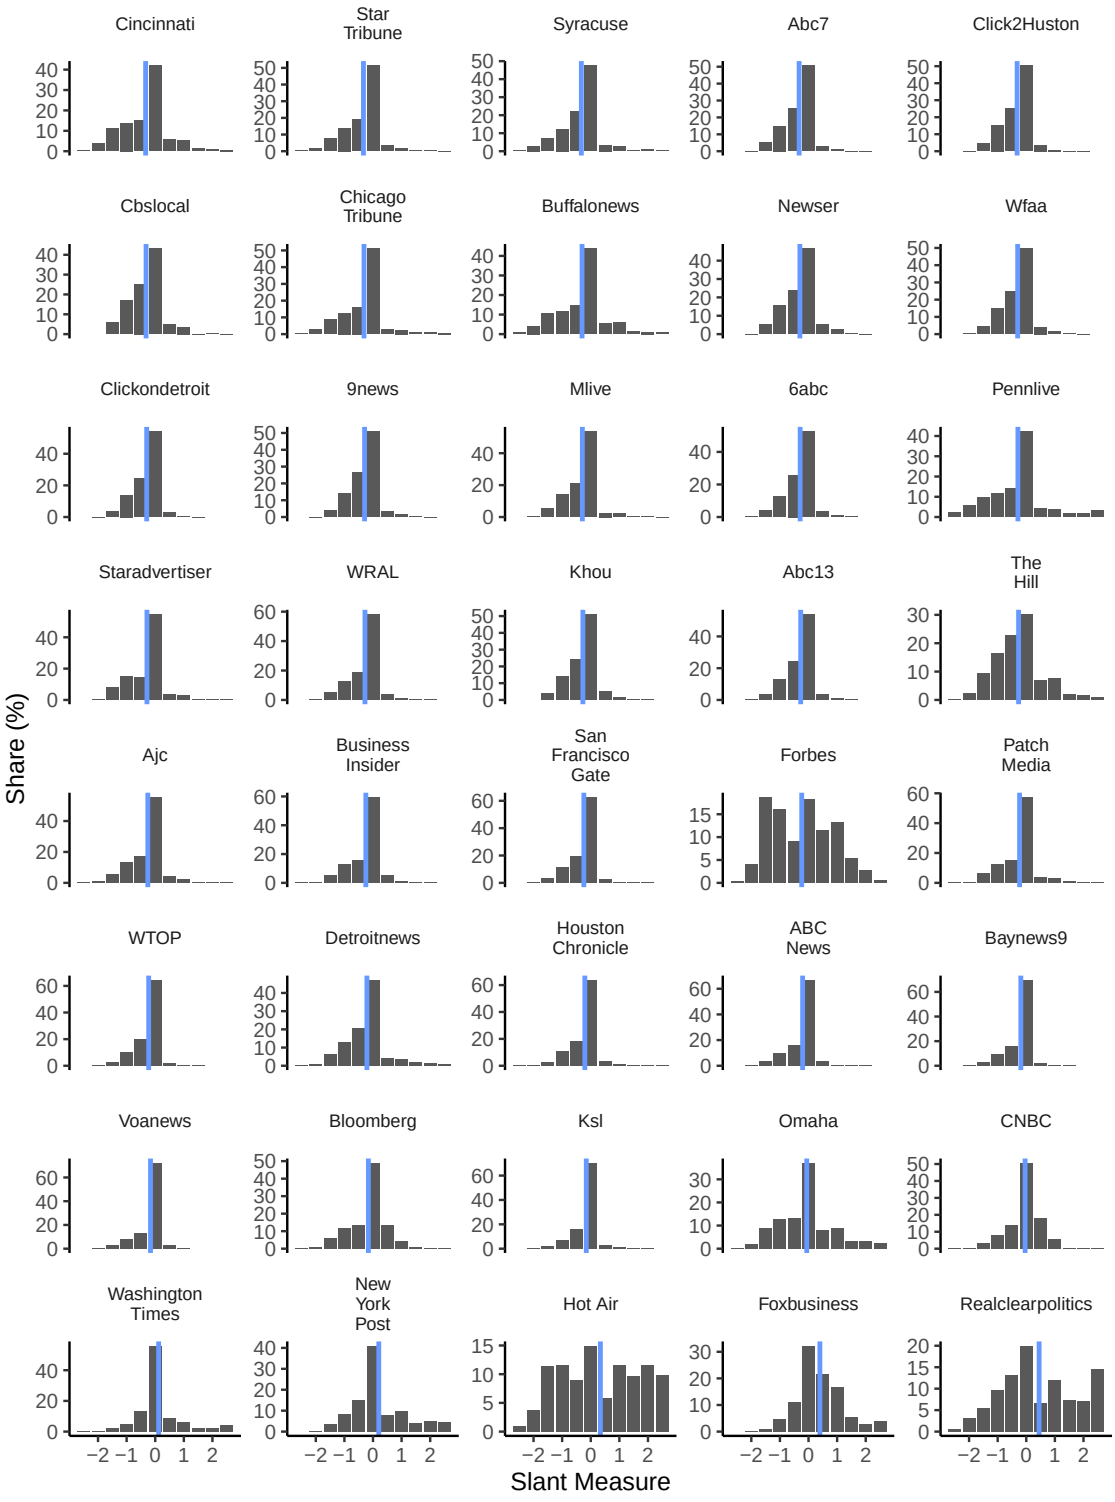
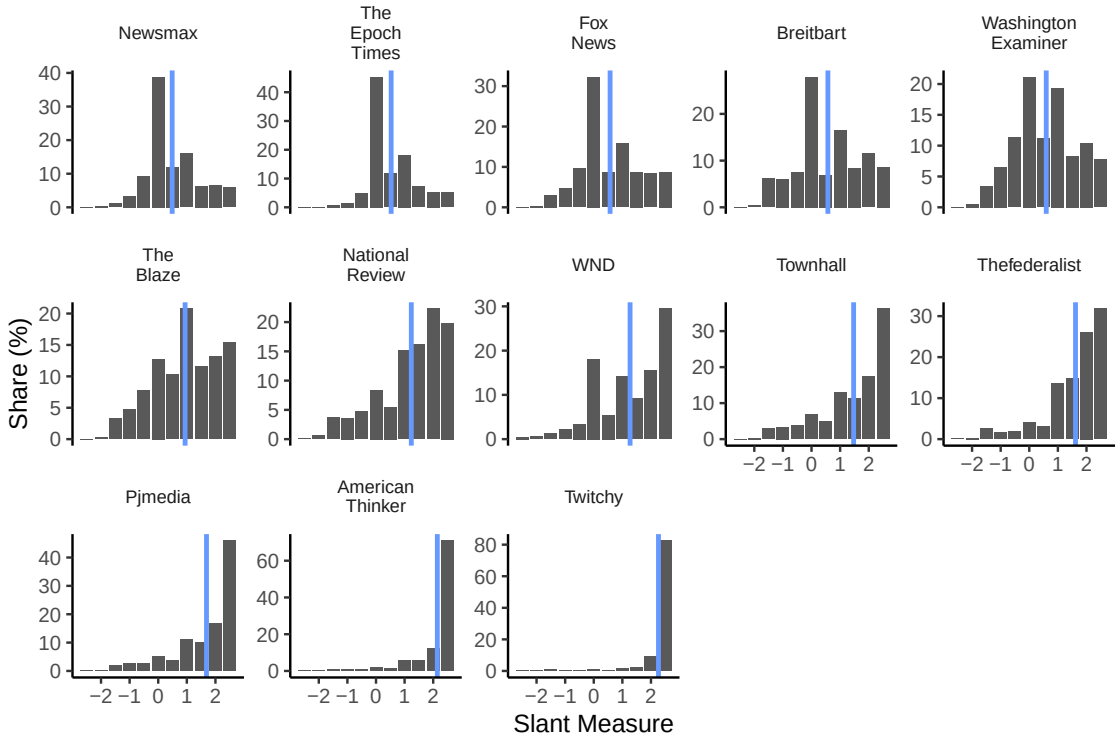


Figure J.1: Slant Distribution within Outlets (Continued)



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 GPT-4o'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).