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

December 2024

Abstract

There is widespread concern that the social media ecosystem drives users to engage with like-minded news articles, thereby fostering polarization in news consumption. Methodological limitations in estimating slant at the article level have made evaluating these claims difficult. We use data on the near universe ($\sim$ 1 million) of hard news articles published online by the top 100 U.S. news outlets in 2019, together with recent advances in natural language processing, to obtain a content-based measure of slant at the article level. Our main finding is that the degree of polarization in news consumption on social media is arguably high. Specifically, the mean slant difference between articles consumed by conservative and liberal users on Facebook is 1.5 times the ideological distance between the average New York Times and Foxnews.com article. We also show that: i) the majority (65%) of the variance in slant across articles arises within outlets, rather than across outlets, highlighting the importance of measuring slant at the article rather than the outlet level. ii) Most news produced is centrist, but the tails of the slant distribution are thick and there is substantial variation in slant across news type and topic. iii) Extreme content is much more likely to be shared widely on Facebook than moderate content. iv) There is substantial pro-attitudinal news consumption on Facebook even *within* the same outlet. v) Polarization in news *exposure* can account for the majority of polarization in news *consumption* on Facebook.

*Braghieri: Bocconi University, CEPR, CESifo, and IGIER. Email: luca.braghieri@unibocconi.it. Eichmeyer: Bocconi University, CEPR, CESifo, and IGIER. Email: sarah.eichmeyer@unibocconi.it. Levy: Tel Aviv University and CEPR. Email: roeelevy@tauex.tau.ac.il. Mobius: Microsoft Research. Email: mobius@microsoft.com. Steinhardt: UC Berkeley, BAIR, and CLIMB. Email: jsteinhardt@berkeley.edu. Zhong: UC Berkeley. Email: ruiqi-zhong@berkeley.edu. We thank Rafael Jiménez-Durán, Stefano DellaVigna, Matthew Gentzkow, Carlo Schwarz, Jesse Shapiro, and seminar participants at Bocconi University, Carnegie Mellon University, the London School of Economics, the National University of Singapore, Northwestern University, NOVA University in Lisbon, the University of Cologne, the University of Copenhagen, the University of Essex, the University of Vienna, the University of Warwick, the EJPE-CEPR Political Economy (Naples), the NYU Center for Social Media and Politics, the Polarize conference in Bergen, and the TAD annual conference for helpful comments. We thank Maor Ankri, Oliver Kiss, Heitor Sousa-Borges, Wei Pin, Uri Schwartz, and Daniel Wolkowicz for excellent research assistance. We also thank John Wihbey, Kenneth Joseph, and David Lazer for generously sharing journalist data. Funding for this project was provided by the European Research Council (ERC Grant 759885) 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' newsfeeds lead to the formation of "echo chambers" and "filter bubbles" that promote the consumption of pro-attitudinal news (Pariser, 2011; Sunstein, 2017).[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 (Bakshy, Messing and Adamic, 2015; Gentzkow and Shapiro, 2011; Guess and Nyhan, 2018; Guess, 2021; Nelson and Webster, 2017).[2] 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. 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 propose and implement a novel content-based method to assign slant to individual news articles rather

[1] Pro-attitudinal (or congenial) news refers to news whose slant matches a person's ideology.

[2] Gonzalez-Bailon et al. (2023) and Green et al. (2023) are notable exceptions. We discuss our contribution relative to those articles towards the end of the introduction.

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 two main datasets. The first dataset encompasses the near universe (~ 1 million) of hard news articles 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 articles publicly shared on Facebook at least 100 times. The two data sources are complementary: the first source contains the URL, text, and metadata of news articles in our universe. The second source contains information about various forms of interaction with each URL on Facebook (e.g., views, clicks, shares, etc.), broken down by the political ideology of users.

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 4,500 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 0.86.

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 instead of relying on slant measured at the outlet level, we perform a variance decomposition to estimate 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 ($\sim$ 65%) 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 overall distribution of slant on the production side, as well as the degree to which slant varies across news types and topics. Aggregating articles from all the outlets in our dataset, we find that the overall slant distribution is unimodal, is centered around moderate, and has relatively thick tails. We also show that opinion pieces are much more slanted than non-opinion pieces and that articles about national news are more slanted than articles about local or international news. Lastly, we find that articles about topics such as 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 of news consumption. We define polarization of news consumption as a scaled version of the difference between the average slant of the articles that liberals and conservatives consume on Facebook.[4] We find that the degree of polarization of news consumption on social media is arguably substantial: the distance on our scale between the average slant that liberals and conservatives consume on Facebook is 1.5 times the distance between the average article from the New York Times and the average article from Foxnews.com. Similarly, it corresponds to 1.45 standard deviation units in terms of the overall slant distribution of articles in our database.

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

[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]We prefer the term ”polarization in news consumption” to other common terms in the literature such as ”pro-attitudinal news consumption,” ”congenial news consumption,” and ”partisan selective exposure,” because the term ”polarization” is more suitable to describe a measure of distance like the one we are estimating.

in news consumption is approximately 50% 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 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 severe 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 liberals and conservatives are exposed to on their newsfeeds or by the articles that they select among the articles they are exposed to. We find that 88% of our measure of polarization can be explained by the articles that liberals and conservatives are exposed to on their newsfeeds, thus highlighting the role of the newsfeed curation process in driving polarization in news consumption on social media. Third, in an attempt to further unpack the role of the newsfeed 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, but not the only driver of polarization in news exposure on social media. Specifically, we find that if users' newsfeeds simply reflected the articles shared by their friends, the degree of polarization in news exposure on Facebook would be around 50% of what it currently is.

This project contributes to the literature trying to measure the differential consumption of news on social media by liberals and conservatives. 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; Flaxman, Goel and Rao, 2016; Gonzalez-Bailon et al., 2023). Specifically, an important hypothesis that researchers and policymakers are interested in assessing is whether liberals and conservatives 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, for instance, by measures of segregation or favorability.[5] The conceptual distinction between the polarization measure we introduce and measures of segregation 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.[6]

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; Gonzalez-Bailon 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).[7] 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 liberals and conservatives 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 individuals might consume pro-attitudinal articles even within outlets.[8]

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

[6]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 liberals predominantly read the article in the liberal outlet and conservatives predominantly read the article in the conservative outlet. In contrast, measures of segregation would increase if liberals accessed the article through the liberal outlet and conservatives accessed the same article through the conservative outlet.

[7]Audience-based measures of slant are those that rely on differential consumption or sharing patterns by liberals and conservatives.

[8]A relatively small literature does attempt to assign a content-based measure of slant to individual news articles

Overall, the content-based measure of slant at the article level we develop in this paper allows us to make the following qualitative contributions to the literature. First, we analyze the landscape of slant of the articles produced by the top 100 U.S. outlets in 2019, showing for the first time that most of the variance in article-level slant occurs within rather than across outlets and that there is ample variation in slant across news types and topics. Second, we introduce and estimate a novel measure of polarization in news consumption on social media, which captures the degree to which liberals and conservatives consume articles with different slant. Contrary to much of the literature, we find that polarization in news consumption on social media is arguably high.[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 contribute to 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 7 presents robustness analyses; Section 8 concludes.

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

(Garz et al., 2019). Such papers, however, do not deploy their article-level measure of slant at scale, nor do they use it to study polarization of news consumption on social media.

[9] This finding dovetails with the analysis in Gonzalez-Bailon et al. (2023) showing that liberals and conservatives tend to consume different outlets and, even more markedly, different articles.

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. 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, 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 newspaper. Conversely, the bundles of articles that social media users are exposed to are tailored to their individual tastes: news articles appear on users' newsfeeds 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 liberal 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 conservative 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 liberals and conservatives 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 two main datasets: one that includes the URL, text, and meta-data of a large collection of hard news articles, and one that details the number of Facebook views, clicks, and shares for a subset of those articles.

3.1 Near Universe of Hard News Articles

The first dataset we employ in our analysis comprises the near universe ($\sim$ 1 million) of hard news articles 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 94% of all online consumption from U.S. outlets focused on hard news, thus yielding a database of articles with very broad coverage.

To collect the news articles published online in 2019 by the top 100 U.S. news outlets, we scraped articles 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 excludes sports, entertainment, classified advertisements, weather forecasts, etc. 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 arguably 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 1,096,622 hard-news articles. Appendix Table F.1 presents the list of outlets (column 1), as well as the number of hard news articles published in 2019 by each outlet

[10] Our criteria 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 scraped English articles where the HTML 'pageType' meta field is defined as 'article'.

[12] We further excluded outlets where the scraper did not seem 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.

(column 2).

Article cleaning The news articles we collected often include chunks of text extraneous to the article text, such as advertisements or links to other articles. In our data cleaning procedure, we attempted to remove as much extraneous text as possible to be able to train our machine learning algorithm based on the content of news articles, as opposed to extraneous factors related to a specific outlet's online appearance. To clean the articles, we first tested existing packages (Domdistiller, Readability, Trafilatura) that extract relevant text from HTML files. We compared the differences across these packages and chose the most relevant package for each outlet. We also removed extra white space and removed paragraphs that contained only a short line of text (up to 100 characters) and no period, since those lines tend to be ads or captions and not 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 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, age group, and users' political affinity (discussed in detail below). [14]

For privacy protection, Meta added Gaussian noise to each data point on engagement (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 re-

[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 SS1 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.

lease of statistical summaries of datasets containing individual information (Messing et al., 2023).[15] In all our analyses based on the Facebook engagement data, we aggregate engagement counts across thousands of news articles within our fine-grained article slant bins, thereby averaging out the noise present in the engagement measures at the level of individual articles. Column (3) of Table 1 provides summary statistics for the articles in the URL-level Facebook Activity Dataset.

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. We define users in the two liberal buckets as liberals, users in the two conservative buckets as conservatives, and users in the middle bucket as moderates. The political affinity measure is derived from users' interactions with political pages on Facebook, employing a methodology inspired by Barberá et al. (2015).[16] Bond and Messing (2015) validate the 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.

The political affinity measure allows us to analyze the universe of U.S. Facebook users to whom Meta assigns a political affinity score. Since the score is only assigned to users who like at least one news page on Facebook, the political affinity measure covers approximately 25% of users.[17] Reassuringly, individuals who are assigned a political affinity score are responsible for the lion's share of engagement - 73% to 88%, depending on the engagement measure - with hard news articles on Facebook, as shown in Table A.1. Furthermore, in Section 7 we show that, qualitatively, our results go through when we use an alternative ideology proxy that relies on basic demographics observed for *all* active Facebook users in the U.S.

[15] By requiring that the probability of any outcome is at most e^{ε} different (with a small additive term δ) between datasets that differ by a single individual, differential privacy ensures that the presence or absence of any one person's data has a minimal impact on the output of any analysis. As a result, the risk of identifying or inferring information about any one person from the dataset is extremely low.

[16] 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.

[17] This number is still orders of magnitude larger than the sample size in survey-based studies that recruit members of a panel.

Besides incomplete coverage, the political affinity measure has an additional drawback: among the more than 500 political pages employed in its construction, there are 18 news outlets that overlap with the ones in our list.[18] Including news outlets from our list in the construction of the political affinity variable might introduce a degree of circularity in our estimates. In Section 7, 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, and by showing that, qualitatively, our results go through when we use an alternative ideology proxy that relies on basic demographics rather than on Facebook's political affinity measure.

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 4,500 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 political slant to news articles, it is paramount to obtain a high-quality training dataset with accurate labels. To obtain the training dataset, we employed, for 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 4,632 news articles randomly selected from our database, with the randomization stratified by week of the year. Specifically, raters were asked to assess the slant of each news article in the training set on a 7-point scale, where 1 corresponded to a very left-wing article and 7 corresponded to a very right-wing article.

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

[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.

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. The raters were told to consider primarily three criteria when assigning a measure of slant to each article: language, political position, and issue coverage. Language refers primarily to the usage of partisan terminology as in Gentzkow and Shapiro (2010). Political position refers to whether the article takes a position, if at all, that is closer to the Democratic or Republican stance on the issue. Issue coverage refers to whether the article's content disproportionately emphasized issues important to either party's base.[20] The expert raters were asked to consider all three criteria and then provide a holistic assessment for each article. The instructions we gave the expert raters can be found in Appendix B.

We decided to rely on experts, as opposed to crowd-sourced ratings (e.g., via mTurk), because we believed our experts would produce higher quality ratings than mTurk workers. In fact, previous 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 in detail below, the correlation between our labels is much higher.

Whenever the two experts disagreed about the direction in which an article was slanted (an event that occurred 6% of the time), we passed 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 rationale 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 finally provided a written justification for his rating.

The procedure above generated two labels for each of the 4,632 articles. The correlation between the original two labels is 0.56 and increases to 0.72 after including the RA's correction. Cohen's Kappa, a commonly used measure of inter-rater reliability, is 0.68, which is generally interpreted as "substantial" agreement (the second best category out of 5).[21] The disagreement in slant across the two ratings is arguably small: the median and modal absolute difference between the two labels

[20] See, for instance, the example of selective coverage in Chopra, Haaland and Roth (2024) in which the Bureau of Labor Statistics evaluates a prospective reform along two dimensions and news articles selectively report only one of the two dimensions.

[21] Cohen's Kappa is a statistical measure used to evaluate the agreement between two raters who each classify items into mutually exclusive categories. Unlike simple percent agreement, it accounts for the possibility of the agreement occurring by chance, providing a more accurate assessment of inter-rater reliability.

is 0, and the average absolute difference between them is 0.72 on our 7-point scale.

To further reduce idiosyncratic noise in the rating, we use the average slant rating of the two expert raters, after the RA correction, as the label for each article in the training data.

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 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 $\sim$3,000 articles from our training set containing the average of our two experts' labels. The remaining $\sim$1,500 labeled articles were held out for validation.

Fine-tuning GPT-4o yielded substantial improvements in the model's ability to predict the average of our two experts' labels. The Pearson correlation between the model's predictions and the average of our two experts' ratings in the hold-out portion of our training dataset is 0.72 when the model is not fine-tuned and 0.86 when the model is fine-tuned. Thus, the increase in precision due to fine-tuning is around 20%.

For reference, we list the titles of the ten most clicked articles on Facebook in 2019 that are, according to our model, far left and far right slanted, respectively, in Table A.4.

4.3 Validation Exercises

Our paper introduces the first large-scale content-based measure of the slant for individual news articles. 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 correlation of 0.86. In Appendix E, we show that, 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 of 0.91. In light of this upper bound, our model predictions attain 95% of the maximum achievable correlation, which we view as very satisfactory.

Second, 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 2 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.

Third, 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 week. For the year 2019, we found a total of 540 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.1. We find a smooth linear relationship between the two ratings. Given that the AFM raters are chosen to represent a broad range of political perspectives, the strong alignment between our model predictions and the AFM labels should also alleviate concerns about our two expert raters having the same political orientation.[22]

The comprehensive validation exercises conducted across our dataset solidify our confidence in the machine learning measure of article-level slant developed in this study.

[22]Bot our raters self-identify as weakly liberal.

5. Results: 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 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, as some of the existing literature does, 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.[23] Figure 3 shows the distribution of article-level slant separately for each outlet. The two outlets

[23] When we talk about Fox News, we refer to the articles posted on Foxnews.com.

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 1.0 point on the 7-point 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. Appendix Figure F.1 shows the distribution of slant for all outlets and demonstrates that there is substantial within-outlet variation 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 35% of the variance in article-level slant. In other words, ignoring within-outlet variation in slant induces one to miss 65% of the overall variation in article-level slant —that is, all the variation that occurs *within outlets*.

Our variance decomposition has two main implications: 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.

5.2 Overall Distribution of Slant

Figure 1 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 is unimodal and centered around moderate values. Notably, 41% of all articles have a slant of zero, classifying them as centrist.

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.

Opinion vs. Non-Opinion We employ a machine learning algorithm trained on expert-labeled data to identify opinion pieces.[24] The algorithm classifies 23% of articles as opinion pieces. We find that opinion pieces are significantly more slanted than non-opinion articles. Specifically, 43% of opinion articles have an absolute slant of 2 or more on our 0-3 scale, compared to only 2% 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.

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.

[24] Above and beyond assigning a measure of slant to each article in our training set as described in Section 4.1, the two experts also classified each article as an opinion or non-opinion piece. See Appendix C for details.

National vs. Local vs. International News Our next piece of analysis investigates the relationship between slant and whether an article covers local, national, or international news. As discussed in detail in Appendix C, the classification of articles into local, national, and international news was performed by GPT-4o.

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 (where 0 indicates no slant and 3 indicates the most extreme absolute slant 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.

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.[25] 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 Jr, 2011)

[25]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 centrist articles, namely articles with a slant between -0.5 and 0.5.

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 under-represented topics among left-leaning articles include technology, the economy, international conflicts, welfare, trade, national security, and innovation.

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.6, 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.[26]

6. Results: News Consumption on Facebook

In the previous section, we provided a broad overview of the slant landscape of the near-universe of hard news articles in our dataset. 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 introduce a new measure of polarization in news consumption. Our measure, denoted as $\mathscr{P}(\overline{S_{lib}}, \overline{S_{con}})$, calculates polarization based on the difference in click-weighted average slant between liberal and conservative users, normalized to range between -1 and 1. The measure is constructed using the following formula:

$$\mathscr{P}(\overline{S_{lib}}, \overline{S_{con}}) = \frac{\overline{S_{con}} - \overline{S_{lib}}}{6} \quad (1)$$

[26]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.

where $\overline{S_j}$ represents the click-weighted average slant for users with political affinity $j \in${liberal, conservative}. To interpret the measure on a standardized scale from -1 to 1, we normalize it by dividing it by the maximum possible distance (six points) on our scale.

We believe our measure has three desirable features that make it an important complement to other measures developed in the literature, such as measures of segregation. First, and foremost, 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, indeed, liberal users tend to consume left-leaning articles and conservative 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, in the example from Appendix F, a comparison of our polarization index calculated at the article level (henceforth, article-level polarization) and the index calculated at the outlet level (henceforth, outlet-level polarization) reveals the following intuitive results. 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. Polarization at both levels is equal if and only if individuals select articles within outlets without regard to their political slant, as would occur if users consumed random articles from within each outlet.

Many of the standard measures in the literature (e.g., measures of segregation) lack these features. First, measures of segregation rely solely on consumption patterns; they do not feature an independent measure of the slant of the content consumed. Second, the interpretation of segregation measures is arguably less intuitive. For instance, segregation measures can reach their maximum even when liberals and conservatives consume counter-attitudinal rather than pro-attitudinal news. Lastly, segregation measures suffer from a small-sample bias, as discussed in Gentzkow, Shapiro and Taddy (2019) and as shown in the illustrative example comparing segregation at the outlet and article levels analyzed in Appendix F.[27]

[27] The small-sample bias is unlikely to be relevant for a dataset as large as the URL-level Facebook Activity Dataset, but it could be relevant, for instance, for datasets of browsing behavior containing a relatively small set of users.

Using our novel article-level measure of slant, we estimate that, in 2019, the degree of polarization in news consumption on Facebook was 0.24. The next few paragraphs provide a series of benchmarks that show this level of polarization in Facebook news consumption is arguably high.

Our first benchmark compares the difference in click-weighted average slant between conservative and liberal users (i.e., the numerator in Equation 1) to the standard deviation in the overall database of news articles. The mean difference (on the 7-point scale) in the slant of articles clicked on by conservative and liberal users is 1.44. The standard deviation in our overall database of news articles is 0.99. Therefore, the mean difference in slant for articles clicked on by conservative and liberal users is 1.45 standard deviation units.

Another benchmark for our polarization measure compares the slant difference in articles consumed by conservatives and liberals 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 liberals and conservatives on Facebook is 1.44 times the distance between the average article from The New York Times and the average article from Fox News.

For additional context, we can compare the slant difference in articles clicked on by liberals and conservatives 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 liberals and conservatives is comparable to the slant difference between the average article written by the moderate-conservative columnist David Brooks and by the liberal columnist Paul Krugman. Similarly, this difference approximates the distance between the average editorial by conservative journalist Bret Stephens and one by the New York Times editorial board.

Our final benchmark compares the distance in slant between articles consumed by liberals and conservatives on Facebook to the slant difference between articles shared by various politicians on Twitter. We find that the slant distance in our analysis is similar to the average distance between articles shared by Democratic Senator Elizabeth Warren and those shared by Republican Senator Lindsey Graham.

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 consumption on Facebook is driven by the selection of ideologically congenial articles within outlets. Second, we compare how the distribution of articles that circulate widely on social media differs from the distribution of slant on the production side. Third, we decompose our measure of polarization of news consumption on social media into two components: exposure and selection conditional on exposure. 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 Within vs. Across Outlets

In order to study the degree to which polarization in news consumption on Facebook is driven by the selection of ideologically congenial articles within outlets, we proceed in two steps. First, we provide direct evidence that Facebook users consume pro-attitudinal news within outlets. Second, we study the extent to which shutting down the channel of pro-attitudinal news consumption within outlets affects our measure of polarization.

Figure 8 focuses on the top 20 outlets in terms of online visits and shows that, even within an outlet, the average article consumed by liberals is virtually always more left-leaning than the average article consumed by conservatives. Table F.1 presents statistics about the slant of the average article consumed by liberals and conservatives for each outlet in our dataset. The pattern is similar: partisans tend to consume pro-attitudinal content even within outlets.

In order to study how polarization in news consumption changes when we shut down the channel of pro-attitudinal news consumption within outlets, we re-calculate 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 some 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 shut down the channel of pro-attitudinal news consumption within outlets as described above, we obtain a measure of polarization in news consumption of 0.16. 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 0.24. Thus, the estimate of polarization in news consumption on Facebook increases by 50% when the channel of pro-attitudinal news consumption within outlets is taken into account. Figure A.14 shows that this gap is similar or larger when, instead of focusing on news consumed, we focus on the articles individuals viewed or shared. 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.

6.2.2 Production vs. Distribution

How does the distribution of slant from the top 100 U.S. outlets in 2019 translate into 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 liberals and conservatives, 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 newsfeeds (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 5 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 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, moderate articles with a slant of 0 have an 11% chance of being widely shared on Facebook, whereas extreme articles have an over 39% chance of being widely shared on the platform. Thus, extreme articles are 3.5 times more likely to be publicly shared on Facebook at least 100 times. We note that this analysis includes virtually all hard news articles published online by major outlets in 2019, thus arguably 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. Figure A.7 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 59% 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 at least 100 times on Facebook.

Figure A.8 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 40%. Figure A.9 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.

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 6 shows the distribution of the number of shares per article, separately for liberals and conservatives, 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 content.[28]

Appendix Figure A.11 presents the same analysis separately for articles that users clicked on before sharing and articles that users shared without first clicking on them. We do not find dramatic differences between the two figures, suggesting that reading an article does not decrease the tendency to share pro-attitudinal news. In Section 6.2.4, 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 the phenomenon of echo chambers.

6.2.3 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' newsfeeds; 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 newsfeeds. In this section, we show that the lion's share of polarization in news consumption can be accounted for by differential exposure to slanted articles.[29]

We find that the polarization in news exposure on Facebook is 0.21. Thus, exposure alone can account for 88% of the degree of polarization in news consumption on Facebook. In other words, if liberals and conservatives on Facebook clicked randomly among the articles that appear on their newsfeeds, rather than actively selecting articles among the ones that appear on their newsfeeds,

[28] Appendix Figure A.10 shows the same figure after residualizing for outlet fixed effects. Compared to the unresidualized 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 taking out 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.

[29] We measure the degree of exposure to a particular news article by partisanship using the number of Facebook views that the article received by liberals and conservatives.

polarization in news consumption would be only 12% smaller than it currently is.[30]

As discussed, the articles that Facebook users are exposed to on their newsfeeds depend on: i) the interaction between partisan sharing patterns and the homophily structure of the social network (echo chambers), ii) the personalized content-raking algorithm (filter bubbles), and iii) the ways in which individuals choose to customize their newsfeeds. In the next section, we perform a back-of-the-envelope calculation to assess the degree to which polarization in news exposure and consumption can be explained by echo chambers.

6.2.4 The role of echo chambers

In order to assess the degree to which the phenomenon of echo chambers can explain polarization in news exposure and consumption on Facebook, we perform a back-of-the-envelope calculation. Specifically, we estimate, under a set of assumptions, the amount of polarization in news exposure on Facebook that would prevail if users' newsfeeds were only shaped by their friend networks and by the content shared by those friends. By doing so, we mechanically shut down other channels that shape users' newsfeeds such as the personalized content-ranking algorithm and the users' own customization decisions. We then compare the amount of polarization in news exposure obtained in the first step to the actual amount of polarization in news exposure on Facebook estimated in Section 6.2.3.

Echo chambers can be thought of as arising from the interaction of two distinct features of the social media environment: i) the relatively high degree of political homophily of social media networks, and ii) pro-attitudinal sharing patterns by partisans. When both features are simultaneously present, users share pro-attitudinal content in groups of like-minded individuals, thus giving rise to echo chambers. In Section 6.2.2, we documented strong pro-attitudinal sharing patterns by Facebook users. The remaining ingredient to estimate the role of echo chambers is information about the degree of political homophily of the Facebook network.

There are various measures of network homophily (Jackson, 2008). For the purpose of our analysis, we rely on the following six statistics: the fraction of friends of liberal users who are liberal

[30] 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.

($x_{\{lib \leftarrow lib\}}$), conservative ($x_{\{lib \leftarrow con\}}$), and moderate ($x_{\{lib \leftarrow mod\}}$), and the fraction of friends of conservative users who are liberal ($x_{\{con \leftarrow lib\}}$), conservative ($x_{\{con \leftarrow con\}}$), and moderate ($x_{\{con \leftarrow mod\}}$). Under a set of assumptions, those statistics allow us to estimate the degree of polarization in news exposure that would prevail if users' newsfeeds were only shaped by the structure of the Facebook network and by the content shared by their friends.

Formally, we calculate:

$$\widehat{\mathscr{P}}_{exposure}(\bar{y}_{lib}, \bar{y}_{con}, \bar{y}_{mod}, x_{\{lib \leftarrow lib\}}, x_{\{lib \leftarrow con\}}, x_{\{lib \leftarrow mod\}}, x_{\{con \leftarrow lib\}}, x_{\{con \leftarrow con\}}, x_{\{con \leftarrow mod\}}) =$$

$$= \frac{1}{6}\{[\bar{y}_{lib} \cdot x_{\{con \leftarrow lib\}} + \bar{y}_{con} \cdot x_{\{con \leftarrow con\}} + \bar{y}_{mod} \cdot x_{\{con \leftarrow mod\}}] -$$
$$[\bar{y}_{lib} \cdot x_{\{lib \leftarrow lib\}} + \bar{y}_{con} \cdot x_{\{lib \leftarrow con\}} + \bar{y}_{mod} \cdot x_{\{lib \leftarrow mod\}}]\}$$

where $\bar{y}_{lib}$ denotes the average slant of articles shared by liberal users, $\bar{y}_{con}$ denotes the average slant of articles shared by conservative users, and $\bar{y}_{mod}$ denotes the average slant of articles shared my moderate users. The first (second) part of the numerator in the expression above calculates the average slant that conservative (liberal) Facebook users would be exposed to if their newsfeed was governed solely by the sharing behaviors of their friends. Taking the difference between the first and second components and normalizing it by six yields a measure that is comparable to the actual degree of polarization in news exposure on Facebook calculated in Section 6.2.3.

When calculating the expression above, we plug in the measures of $\bar{y}_{lib}$, $\bar{y}_{con}$, $\bar{y}_{mod}$ that we obtained from the sharing patterns documented in Section 6.2.2. Furthermore, we rely on data from Bakshy, Messing and Adamic (2015) in order to obtain estimates for $x_{\{a \leftarrow b\}}$ where $a, b \in \{lib, mod, con\}$.

Following this approach, we find that the polarization in news exposure resulting from echo chambers alone would be 0.11. Since the overall degree of polarization in news exposure is 0.21, we can conclude that the amount of polarization in news exposure that would result simply from echo chambers is 52% of the actual degree of polarization in news exposure.

7. Extensions and Robustness

7.1 Extending Results to More 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 94% of all online news consumption in the U.S. according to Comscore data. Thus, they provide very wide 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.16 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.[31] 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.

7.2 Robustness: Political Affinity Measure

In this section, we present a battery of robustness checks to address one of the limitations of the political affinity measure in the URL-level Facebook Activity Dataset. As discussed in Section 3, among the more than 500 political pages employed in the construction of the political affinity measure, there are 18 news outlets that overlap with the ones in our list. The reliance on those news outlets in the construction of the political affinity measure introduces a small degree of circularity in our polarization measure.

Our first robustness check excludes the 18 overlapping outlets from our measure of polarization in news consumption, thus severing any direct mechanical relationship between political affinity

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

and slant. We report the results in Table A.2 and find a pattern of results analogous to the one described in Section 6.1. In fact, if anything, polarization in views, clicks, and shares becomes even slightly more severe compared to the baseline case analyzed in 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.3).

Our second 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 (2020 wave), 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.15. 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.

7.3 Robustness: Engagement Data Limitation

As discussed in Section 3.2, another limitation of the SS1 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 based on the articles included in the SS1 dataset. Of course, polarization of news exposure and consumption on social media might be different if one were to consider all the articles shared on Facebook at least once, rather than at least 100 times. In this section, we impose assumptions that allow us to extend our measure of polarization to 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 SS1 dataset. Specifically, we first calculate our polarization measure after removing articles in the bottom 50% of the SS1 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 SS1 dataset. Figure A.17 presents the results. The figure shows that the polarization measure calculated on articles that are shared relatively more often is generally relatively higher. However, 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 0.009.

Next, 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 of articles is shared on Facebook at least once, and ii) starting from the 50% percentile in Figure A.17, the relationship between % of articles 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 less than 100 times; the second allows us to perform the extrapolation.

When we merge our articles dataset to the SS1 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 less than 100 times. In light of the match rate, we know that 100% in Figure A.17 corresponds to only 22% of articles shared on Facebook at least once. Therefore, we would like to extend the figure to 455% (1/0.22) to extrapolate results to all the articles shared on Facebook at least once. Employing the linear extrapolation described above, we find that polarization in news consumption would decrease to 0.21 if the SS1 included all articles shared at least once and not only articles shared at least 100 times. This is in the same ballpark as our main result from Section 6.1 (a 11% decrease).

Polarization in news consumption does not change dramatically when including many additional articles because people are exposed much less often to articles shared fewer than 100 times. Indeed, 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 the articles shared on Facebook at least once would increase the total number of views by only 26.8%.

8. 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

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 Bonnea. 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.

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2023. *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.

Downs, Anthony. 1957. "An economic theory of political action in a democracy." *Journal of political economy*, 65(2): 135–150.

Flaxman, Seth, Sharad Goel, and Justin M. Rao. 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly*, 80(S1): 298–320.

Garz, Marcel, Gaurav Sood, Daniel F Stone, and Justin Wallace. 2019. "Is There Within-Outlet Demand for Media Slant? Evidence from US Presidential Campaign News." *Unpublished Manuscript*.

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.

Gonzalez-Bailon, 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." *American Journal of Political Science*, , (Forthcoming).

Jackson, Matthew O. 2008. *Social and Economic Networks.* Princeton University Press.

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): 2056305117729314.

Newman, N., R. Fletcher, K. Eddy, C. T. Robertson, and R. 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 Jr. 2011. "Newspaper coverage of political scandals." *The journal of politics*, 73(3): 931–950.

Sunstein, Cass. 2017. *Republic: Divided Democracy in the Age of Social Media.* Princeton University Press.

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.

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: Summary Statistics

	(1) All mean/sd	(2) Not in SS1 mean/sd	(3) In SS1 mean/sd
Opinion piece (binary)	0.23	0.21	0.30
	(0.42)	(0.41)	(0.46)
Slant	-0.13	-0.11	-0.19
	(0.99)	(0.90)	(1.25)
FB views (1000s)			216.42
			(752.93)
FB clicks (1000s)			9.18
			(32.62)
FB shares (1000s)			1.78
			(7.71)
N	1132208	885082	247126

Notes: Table reports summary statistics for all hard news articles published by the top 100 U.S. news outlets in column (1). See Appendix Figure F.1 for the full list of outlets. Column (2) shows summary statistics for the subset of articles not found in Facebook's SS1 dataset; while column (3) shows summary statistics for those articles that do appear in the SS1 dataset (SS1 only includes articles shared publicly at least 100 times). The table reports means, as well as standard deviations in parentheses.

Table 2: Polarization in News Exposure, Consumption, and Circulation on Facebook

	Average Slant Liberal Users	Average Slant Conservative Users	Average Slant Difference	Polarization Index
Views	-0.80	0.46	1.26	0.21
Clicks	-0.85	0.59	1.44	0.24
Shares	-0.94	0.91	1.84	0.31

Notes: Table reports the views/clicks/shares-weighted average slant of news diets by liberal and conservative users on Facebook in the first two columns. The third column shows the difference between the first two columns. The last column normalizes the difference by dividing it by 6, such that 0 means no difference and 1 means maximal difference. The news article sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. shared at least 100 times on Facebook).

Figure 1: Overall Distribution of Slant Across Articles

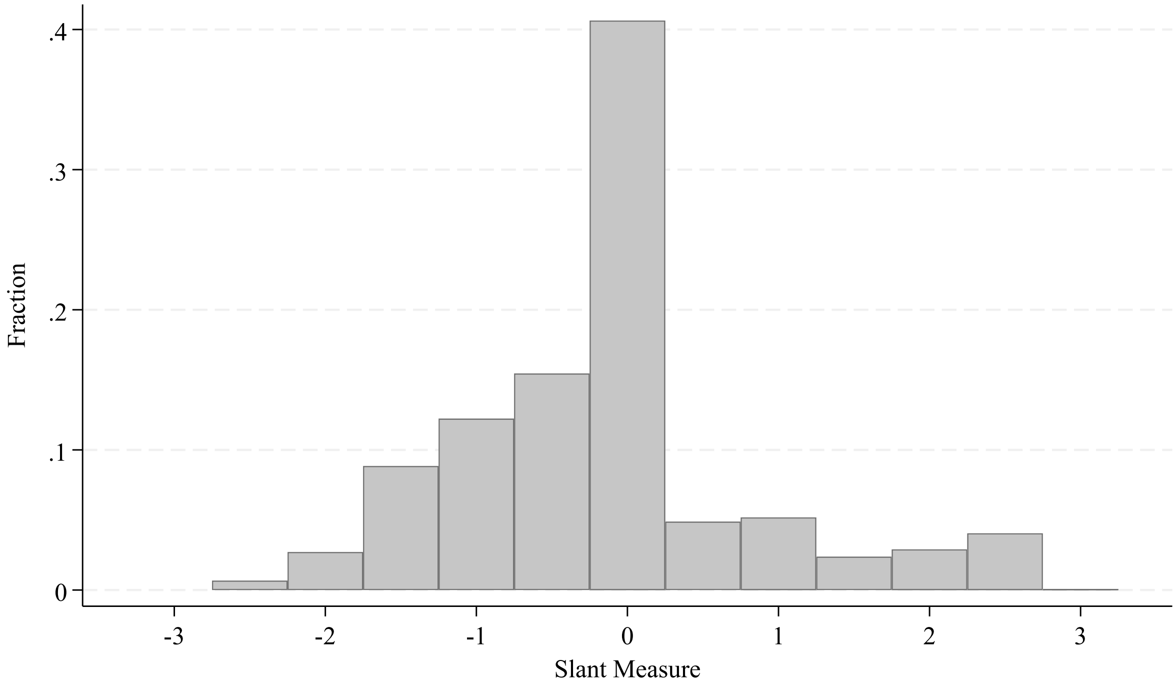

Notes: This figure shows the distribution of slant among all hard news articles ($N = 980,719$) published online by the top 100 U.S. news outlets in 2019. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure 2: Comparison of Our Outlet-Level Measure with Bakshy, Messing and Adamic (2015)

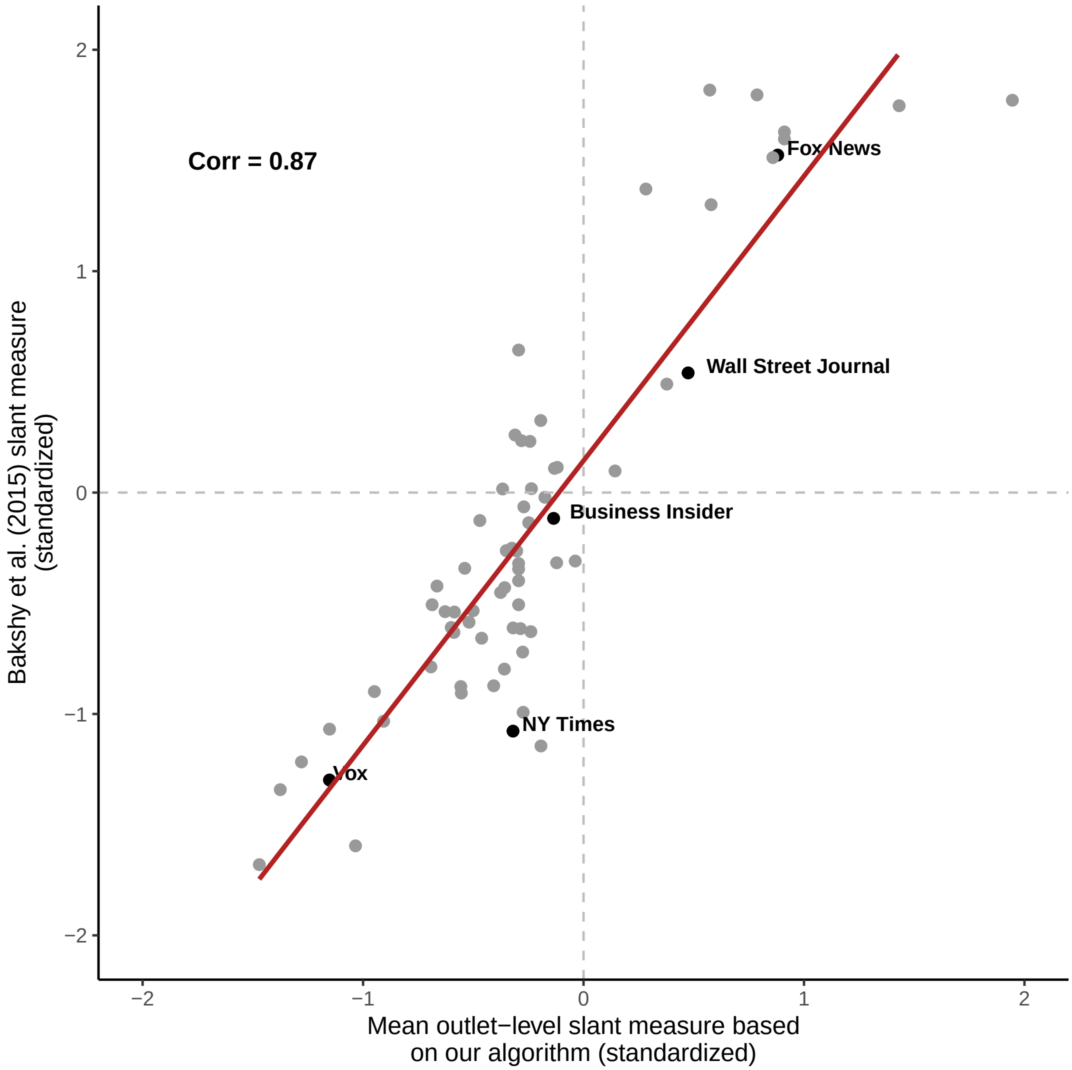

Notes: This figure shows, on the horizontal axis, our outlet-level slant measure, computed as the mean of article-level slant across all articles published by the outlet (where article slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2). The vertical axis shows an alternative outlet-level slant measure for the same outlets, given by Bakshy, Messing and Adamic (2015); the latter measures the slant of an outlet in terms of "alignment": the extent to which an outlet's articles are mostly viewed by left- vs. right-leaning users on Facebook. The figure highlights several well-known outlets.

Figure 3: Distribution of Slant: Fox News and New York Times

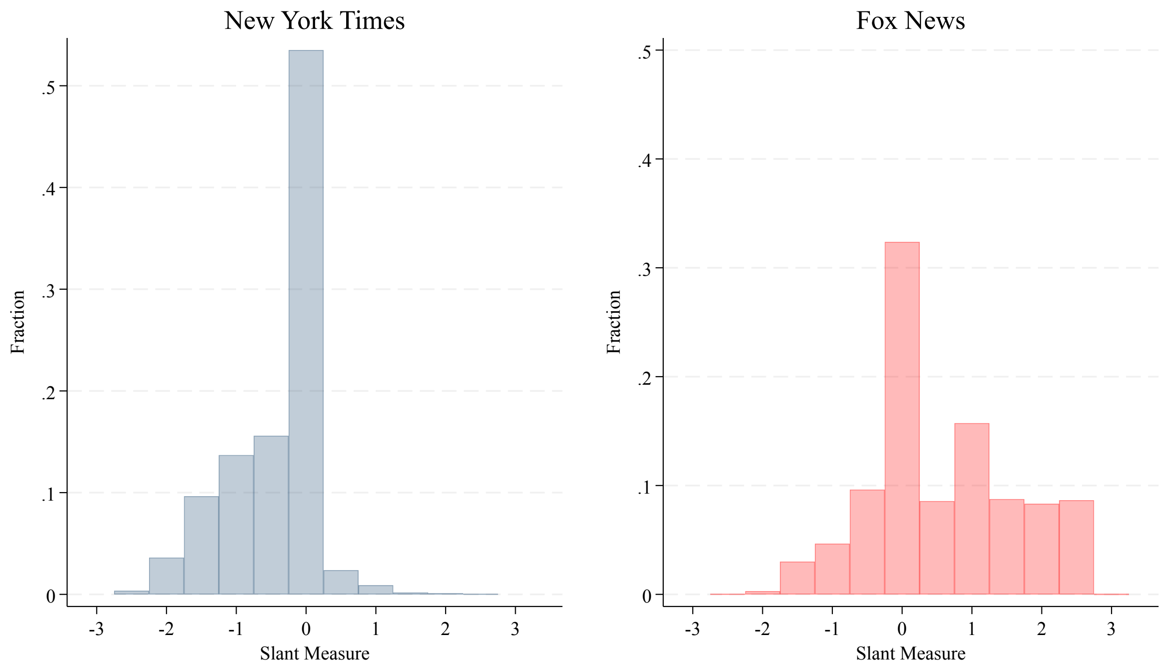

Notes: This figure shows the distribution of slant among all hard news articles ($N = 83,105$), published online by Fox News and the New York Times, respectively, in 2019. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

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

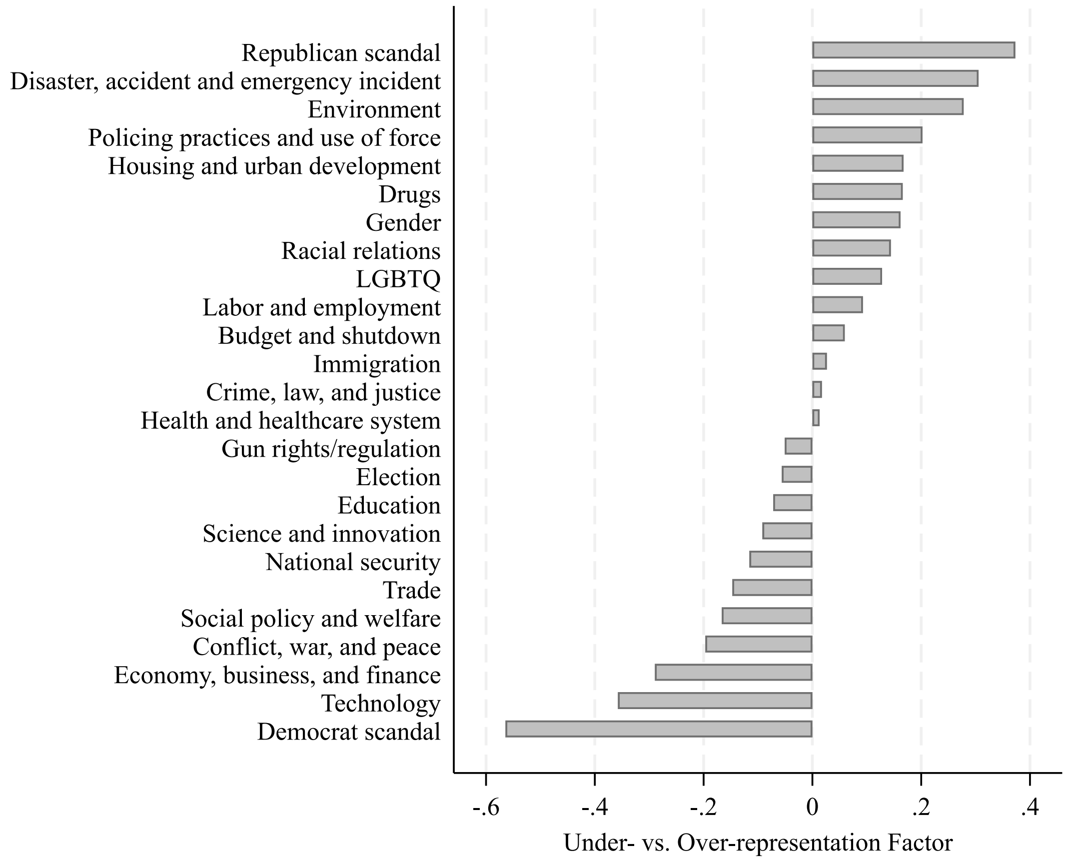

Notes: This figure shows the degree of over- vs. under-representation of topics among left-leaning articles. The figure is constructed as follows. First, all "moderate" articles with a slant between -0.5 and 0.5 are dropped. Second, the remaining articles are defined as "left-leaning" if they have a slant smaller than -0.5 and "right-leaning" otherwise. Third, we compute the fraction of left-leaning articles overall. Fourth, we compute the fraction of left-leaning articles for each topic. Fifth, we divide the fraction of left-leaning articles within a topic by the overall fraction of left-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 left-leaning articles.

Figure 5: Likelihood of Shared on Facebook $>$ 100 Times by Article Slant

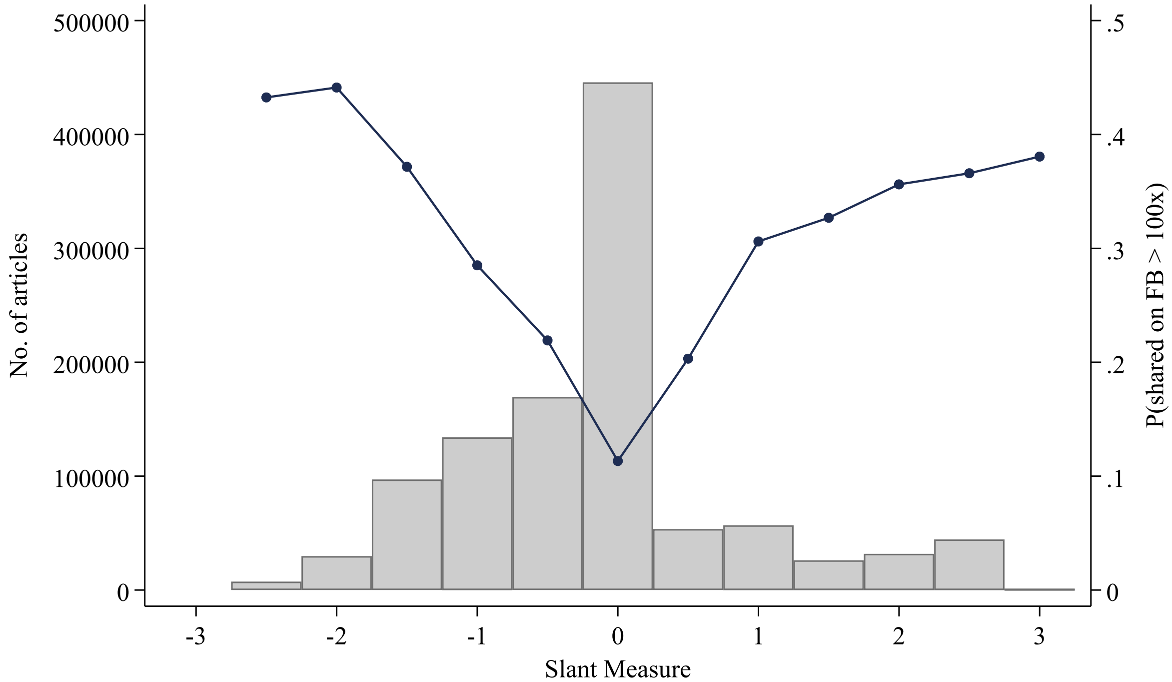

Notes: This figure shows the probability of an article being shared on Facebook at least 100 times as a function of the article's slant. Gray bars show the number of articles published on the top 100 news outlets' websites that fall into a given slant bin (left y-axis); the dotted blue line shows the 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 Facebook's SS1 dataset. The sample includes all hard news articles ($N = 980,719$) 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 -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure 6: Empirical PDF of Facebook Shares by User Ideology and Article Slant

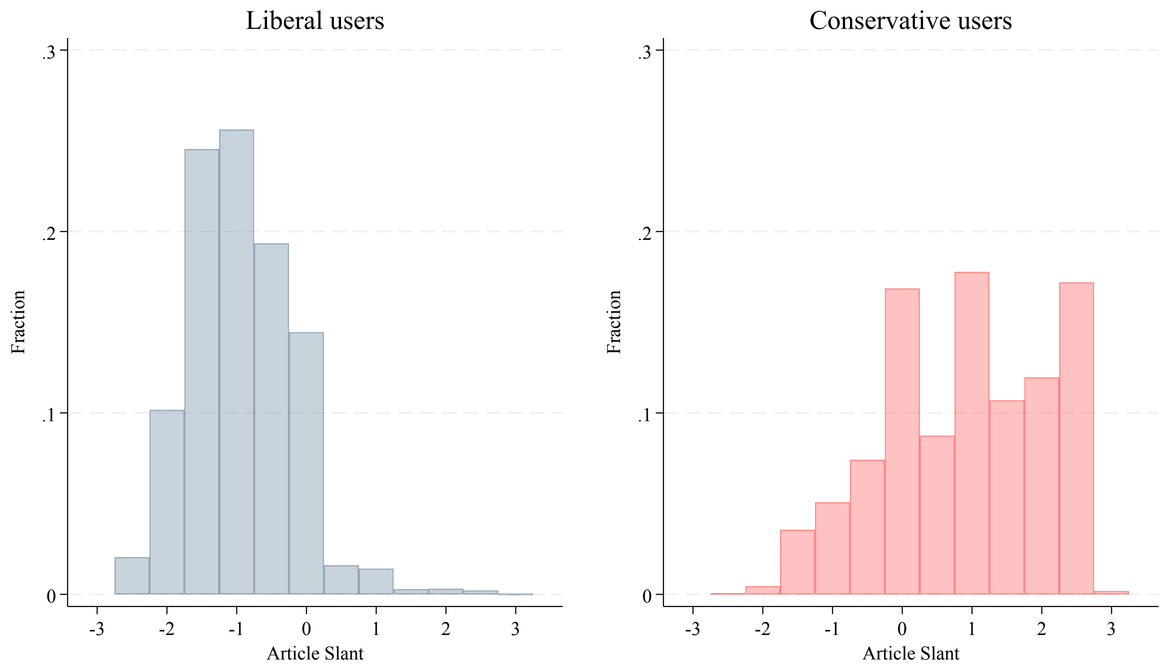

Notes: This figure shows the fraction of all shares by a given Facebook user ideology group (liberal and conservative, respectively) that accrue to articles of a given slant bin. The sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 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 -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure 7: Average Slant of Articles Viewed on Facebook by Outlet and User Ideology

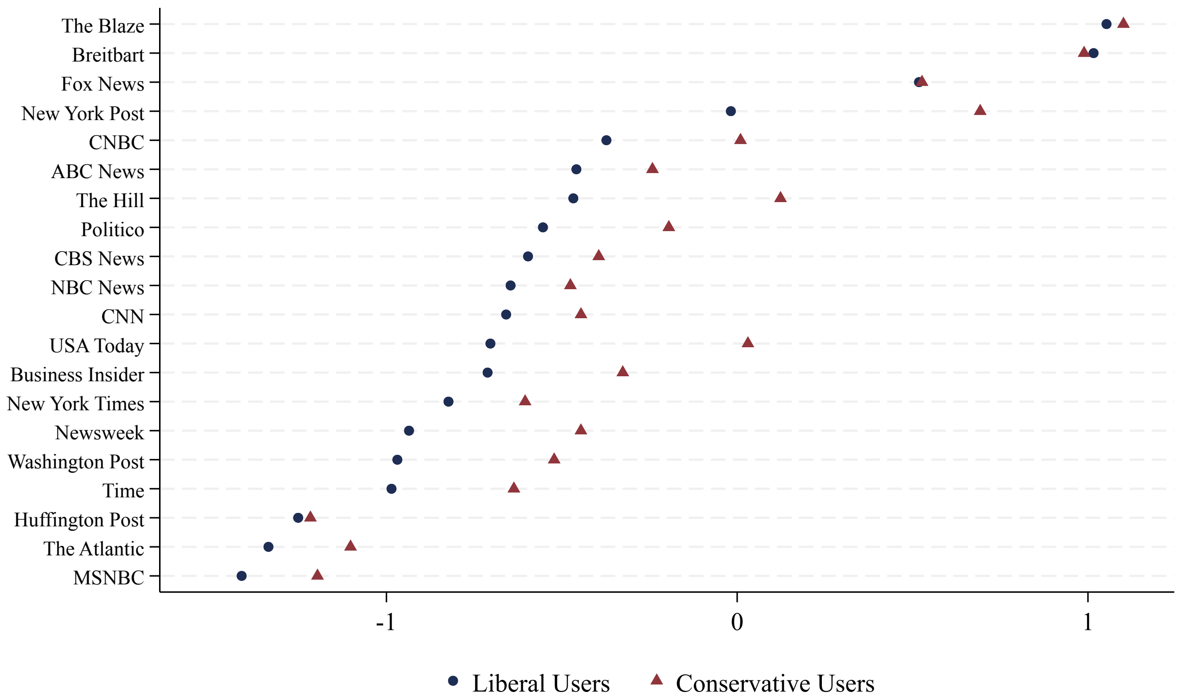

Notes: This figure shows the views-weighted average slant of articles by user ideology and outlet. The sample includes all hard news articles published online by the top 20 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. shared at least 100 times on Facebook). Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure 8: Average Slant of Articles Clicked on Facebook by Outlet and User Ideology

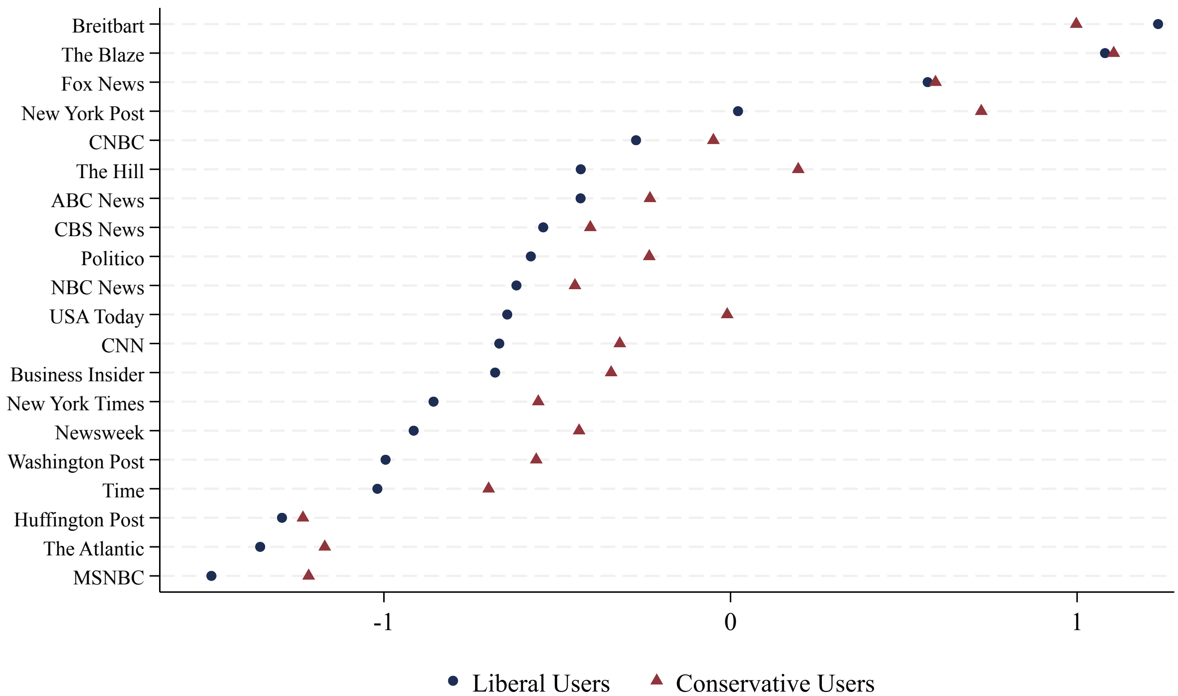

Notes: This figure shows the clicks-weighted average slant of articles by user ideology and outlet. The sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. shared at least 100 times on Facebook). Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Appendix For Online Publication

Table of Contents

A. Additional Tables and Figures

Table A.1: Activity on Facebook by Political Affinity

Political Affinity	Share of Views	Share of Clicks	Share of Shares
Very liberal	18%	20%	23%
Liberal	18%	19%	24%
Moderate	11%	8%	8%
Conservative	13%	13%	15%
Very conservative	13%	15%	18%
All assigned a political affinity	73%	75%	88%
Not assigned a political affinity	27%	25%	12%

Notes: This table shows the share of all views, clicks, and shares in the Social Science One data by the users' political affinity. For the sample of $N = 242,829$ hard news articles published online by the top-100 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. publicly shared on Facebook at least 100 times). The row *"All assigned a political affinity"* provides the share of all views/clicks/shares accruing to users who were assigned *any* political affinity (the sum of the preceding five rows).

Table A.2: Polarization in News Exposure, Consumption, and Circulation on Facebook: Excluding Outlets that Feed Into SS1's Political Affinity Measure

	Average Slant Liberal Users	Average Slant Conservative Users	Average Slant Difference	Polarization Index
Views	-0.75	0.61	1.35	0.23
Clicks	-0.79	0.70	1.48	0.25
Shares	-0.91	1.08	1.99	0.33

Notes: Table shows the same statistics as in Table 2, only that when constructing the statistics, we omit all articles from the 18 outlets that are used by SS1 to compute political affinity scores of users (see Section 3.2 for details).

Table A.3: 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 Liberal Users	Average Slant Conservative Users	Average Slant Difference	Polarization Index
Views	-0.43	0.32	0.75	0.12
Clicks	-0.51	0.41	0.92	0.15
Shares	-0.47	0.55	1.03	0.17

Notes: Table shows the same statistics as in Table 2 with two difference. First, when constructing the statistics, we omit all articles from the 18 outlets that are used by SS1 to compute political affinity scores of users (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).

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*
25 Times White Actors Played People Of Color And No One Really Gave A S**t*	Hollywood Film Depicts Trump Supporters Being Hunted for Sport by Liberals*
A Stain on the Honor of the Navy*	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 slants greater or equal to 2 in our prediction. Articles marked with an asterisk (*) are opinion pieces. The sample covers all $N = 242,829$ hard news articles 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.

Figure A.1: Correlation between Model-Predicted Slant and Third-Party Expert Slant Ratings

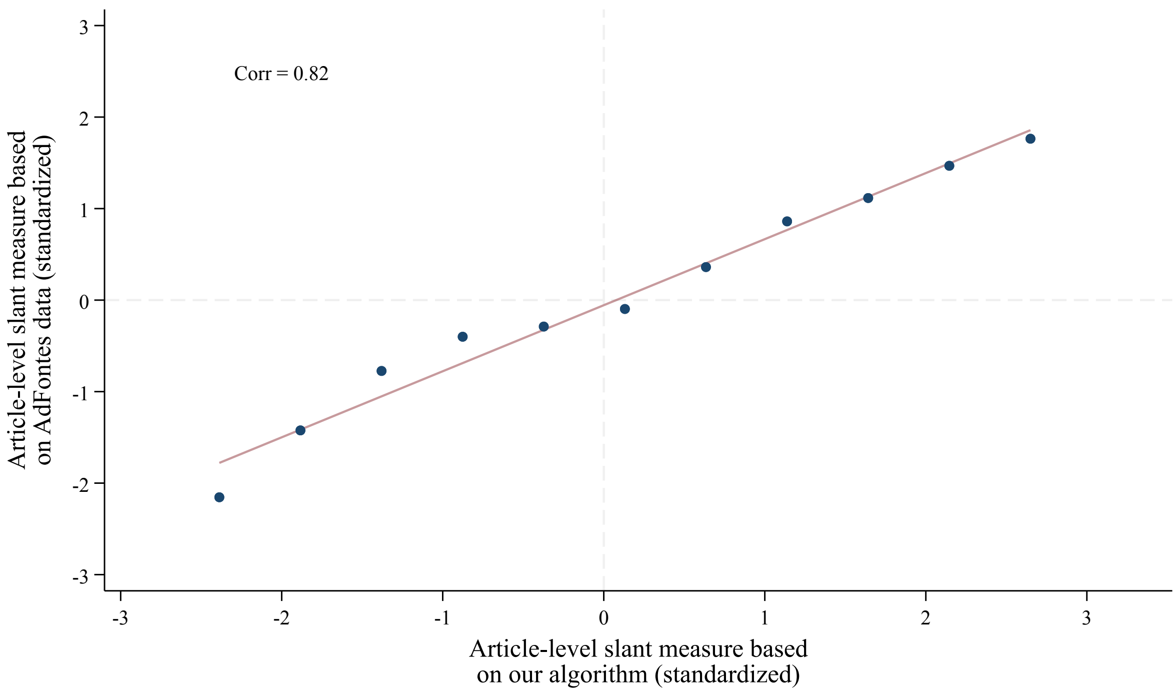

Notes: Figure shows, for each model-predicted slant bin (horizontal axis), the mean third-party slant ratings (by an expert panel curated by Ad Fontes Media) across all articles in that bin. For a total of $N = 487$ articles from 2019 for which Ad Fontes Media ratings could be retrieved. To render the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

Figure A.2: Slant Distribution within Topics

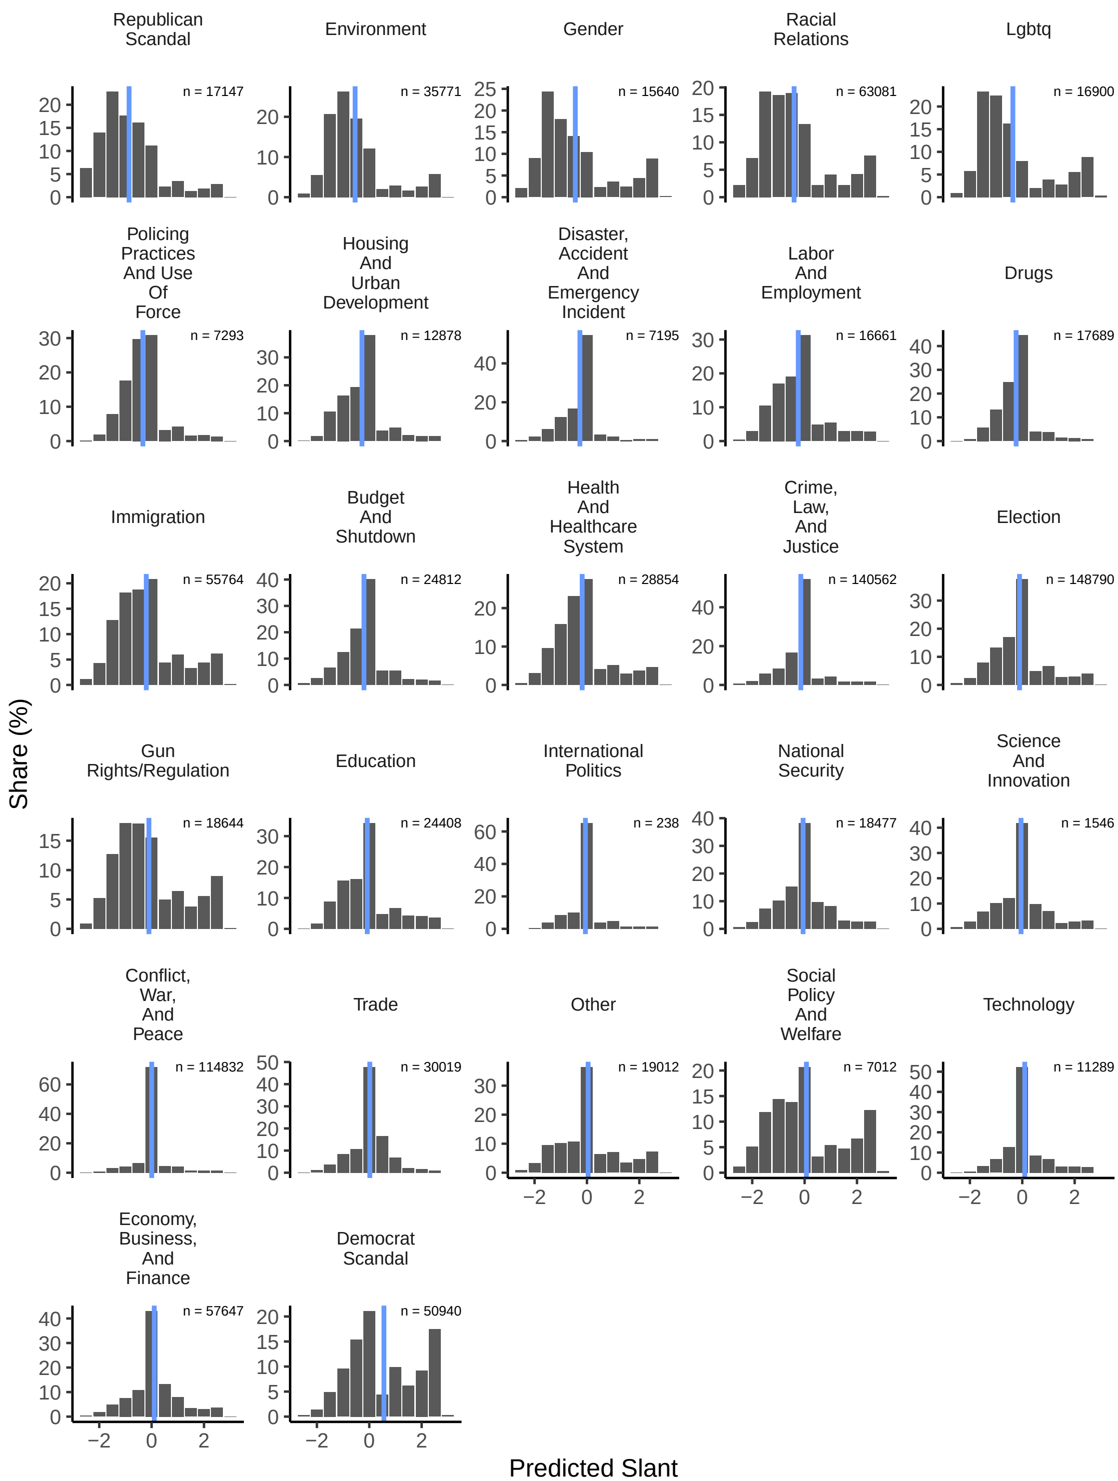

Notes: Figure shows the distribution of slant within the topics. The vertical blue line indicates the mean slant for each topic. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

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

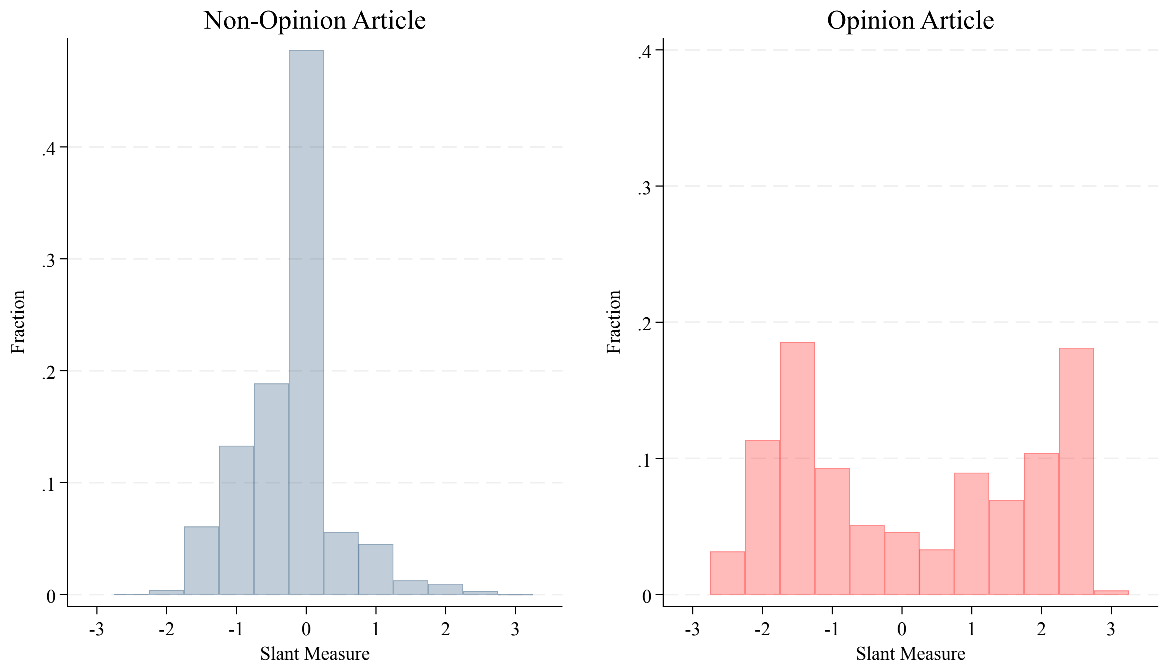

Notes: This figure shows the distribution of slant among all opinion news articles ($N = 118,053$) published online in the top-100 most-read U.S. news outlets in 2019 (red bars), and their non-opinion counterparts (blue bars). An article is classified as an opinion piece based on a machine learning algorithm trained on the labels described in Section 4.1. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

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

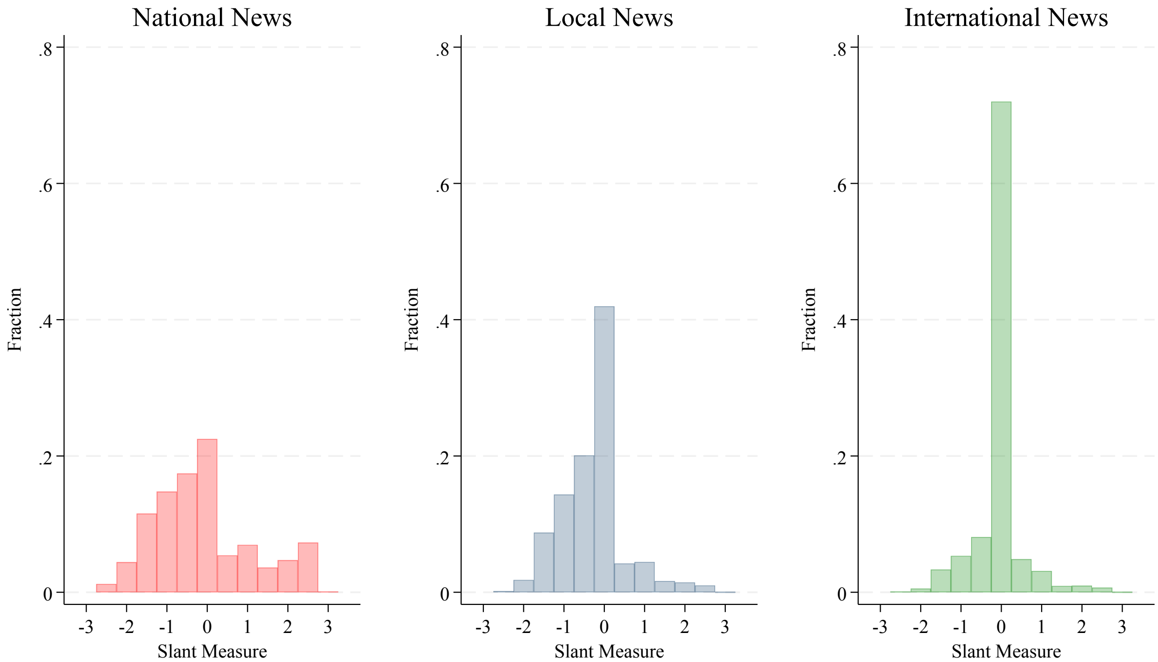

Notes: This figure shows the distribution of slant among all news articles covering national news, among those covering local news, and among those covering international news. An article is classified as national, local, and international news based on a machine learning algorithm trained on the labels described in Section 4.1. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from −3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure A.5: Average Absolute Slant by Topic

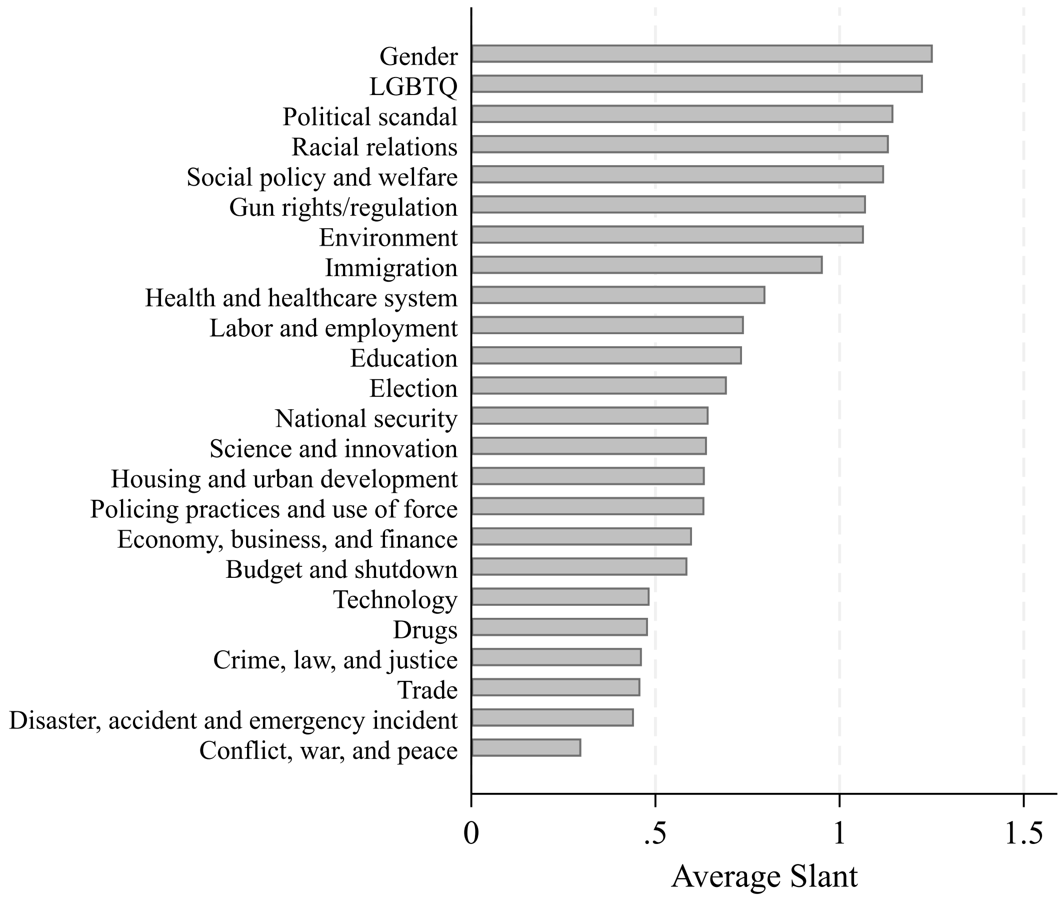

Notes: This figure shows the average absolute slant of the articles falling in each topic category. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The original slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist. The absolute value of the slant scale goes from 0 (a centrist article) to 3 (a very extreme article). The classification into topics is described in detail in Appendix C.

Figure A.6: Standard Deviation of Slant by Topic

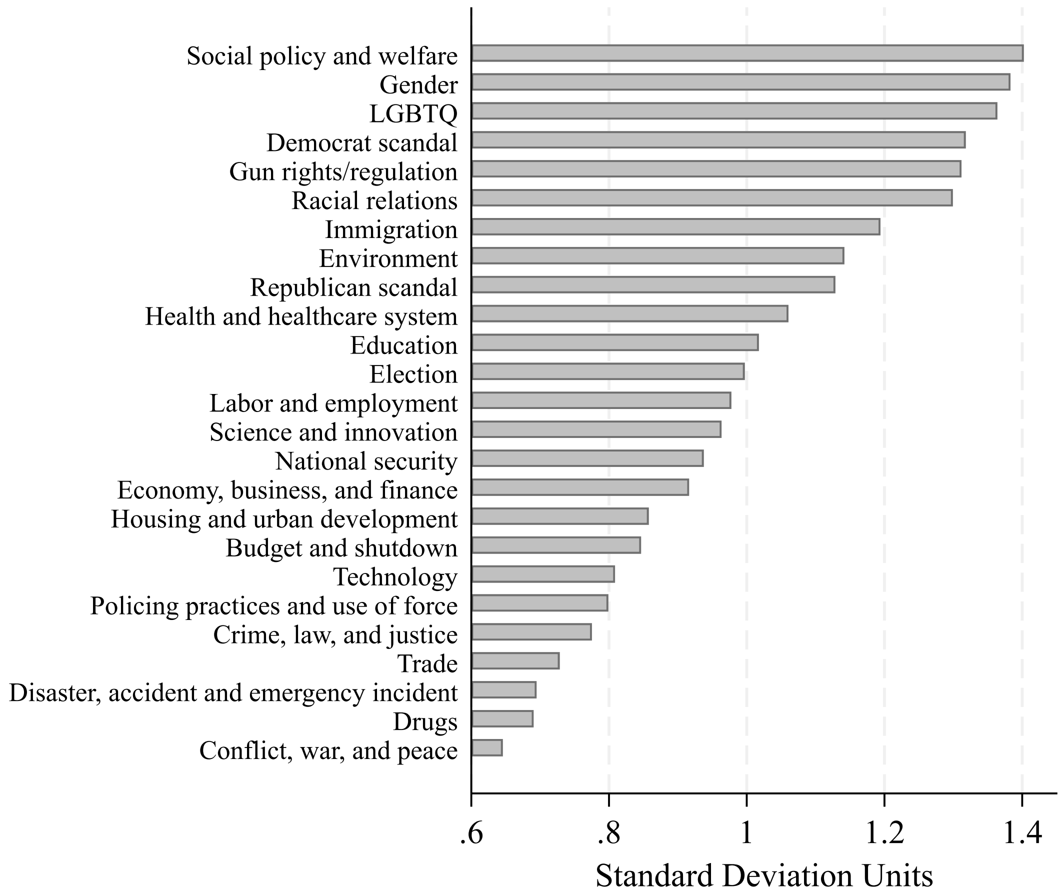

Notes: This figure shows the standard deviation of the slant of the articles falling in each topic category. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from −3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist. The classification into topics is described in detail in Appendix C.

Figure A.7: Empirical PDF of Article Slant in Full Sample vs. SS1 Sample

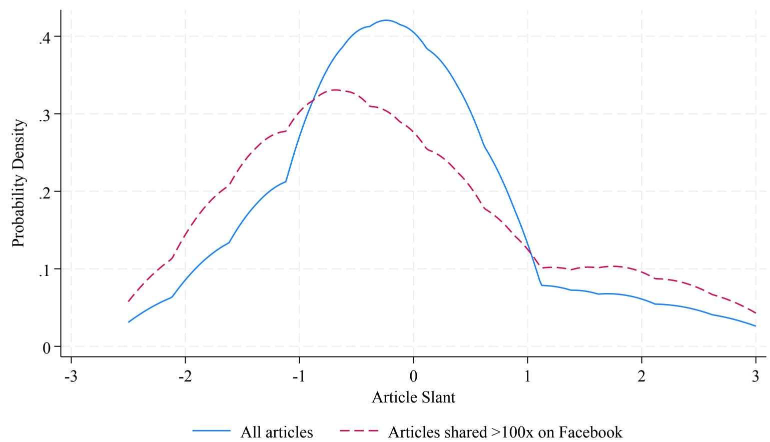

Notes: This figure shows distribution of article-level slant, for two samples: the full sample of all $N = 1,096,622$ hard news articles published online by the top-100 U.S. news outlets in 2019 (in blue) and the sub-sample of articles shared more than 100 times on Facebook (in red)—i.e. the articles included in the Facebook engagement dataset (SS1). Kernel density function estimation with a bandwidth of 0.5, using the Epanechnikov function, is used for smoothing purposes. Article slant is given by fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure A.8: Shares of Opinion Pieces and National News

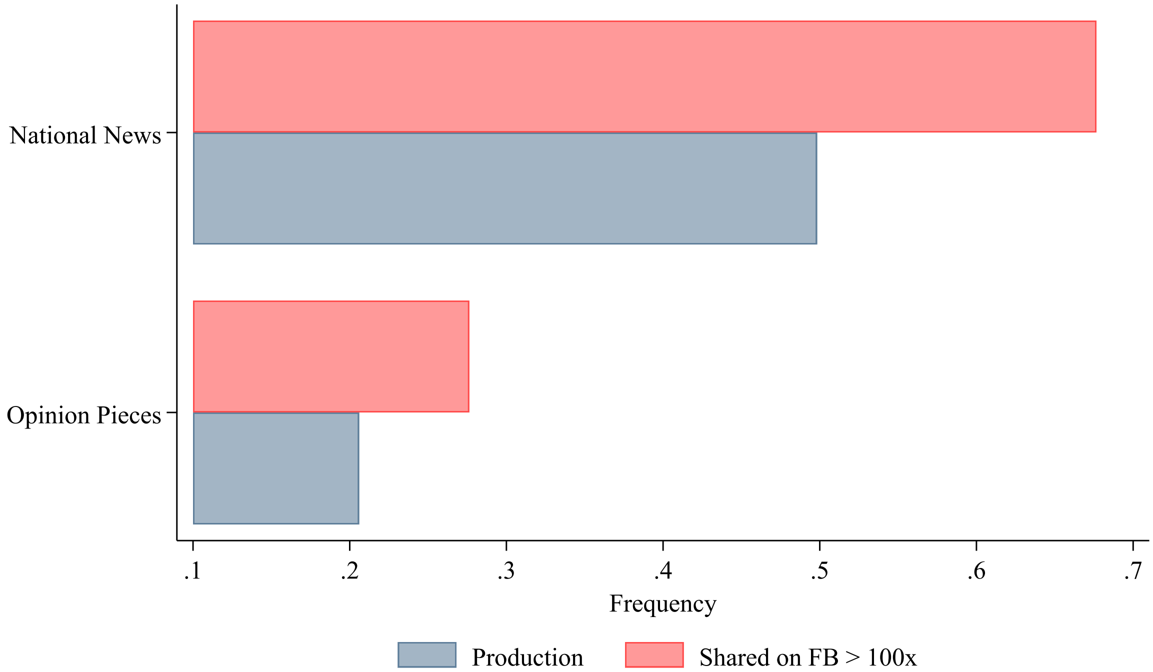

Notes: This figure presents the fraction of articles, both in our dataset and in the URL-level Facebook Activity dataset, classified as national news or as opinion pieces. As discussed in Section 3, in order for an article to appear in the URL-level Facebook Activity dataset, the article has to be shared publicly on Facebook at least 100 times. The classification into national, international, and local news, as well as the one into opinion vs. non-opinion pieces, is described in detail in Appendix C.

Figure A.9: Under- vs. Over-representation of Topics among Articles Shared on FB more than 100x

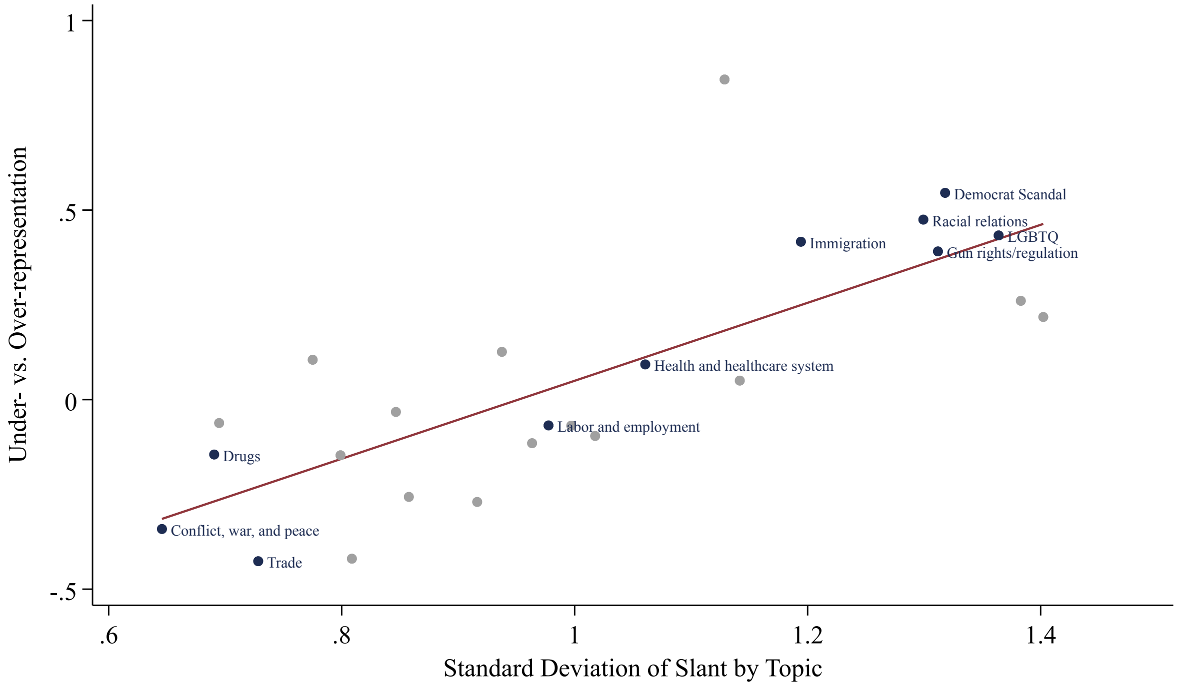

Notes: This figure shows that topics exhibiting a high standard deviation of slant tend to be over-represented in the URL-level Facebook Activity dataset. Specifically, the x-axis captures the standard deviation of article-level slant for each topic. The y-axis captures the degree of under- vs over-representation. 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 SS1 dataset compared to our overall database of articles. The classification into topics is described in detail in Appendix C.

Figure A.10: Empirical PDF of Facebook Shares by User Ideology and Article Slant Net of Outlet Fixed Effects

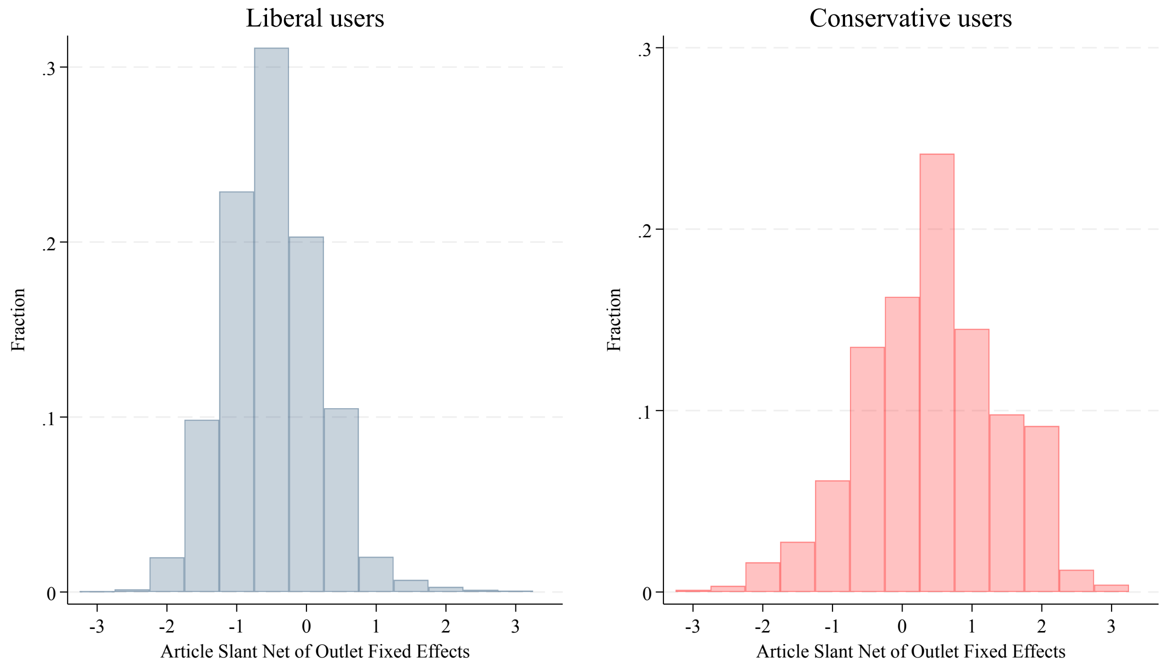

Notes: This figure shows the fraction of all shares by a given Facebook user ideology group (liberal and conservative, respectively) that accrue to articles of a given residualized slant bin. Sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 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 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 -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure A.11: Empirical PDF of Facebook Shares With and Without Clicks by User Ideology and Article Slant

(a) Shares with no Clicks

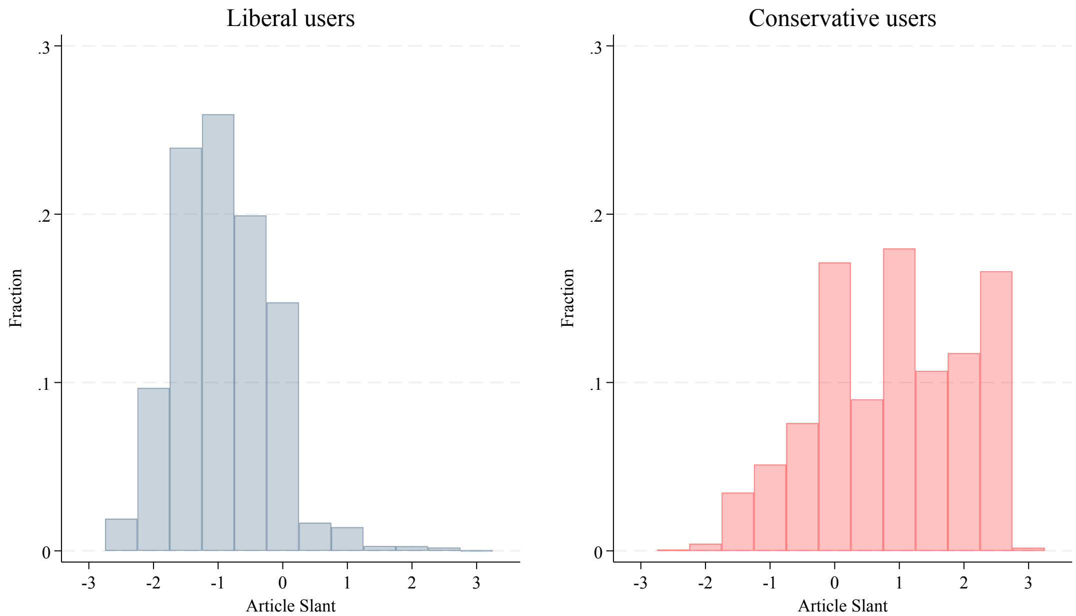

(b) Shares with Clicks

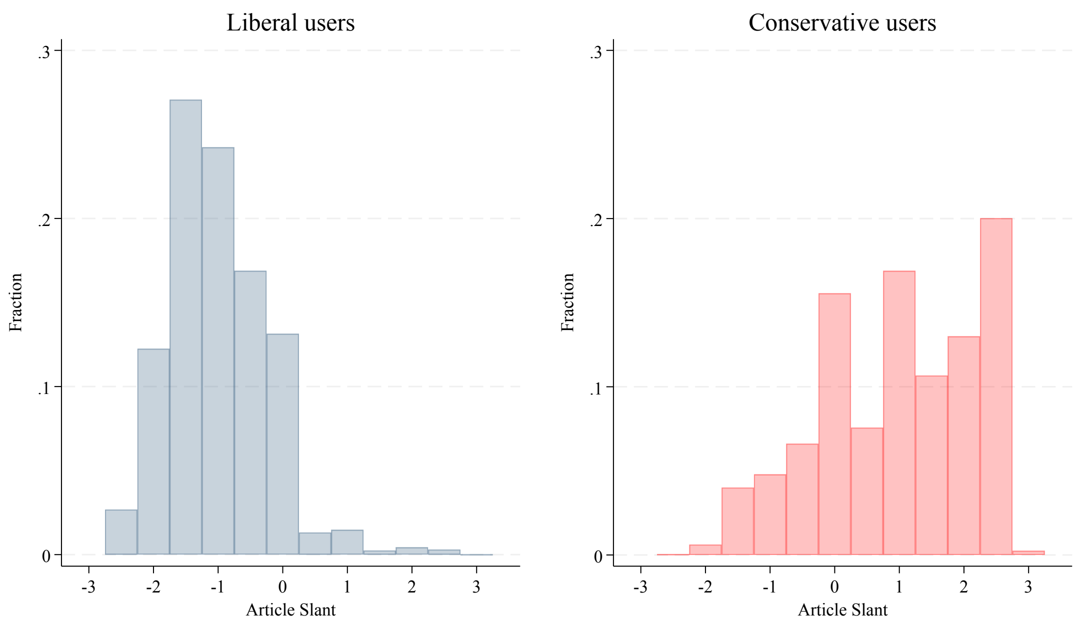

Notes: This figure shows the fraction of all shares by a given Facebook user ideology group (liberal and conservative, respectively) that accrue to articles of a given slant bin. The figure shows the results separately for article links on which users clicked first on Facebook and for those links the users shared without clicking. Sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 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 -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure A.12: Average Slant of Articles Shared on Facebook by Outlet and User Ideology

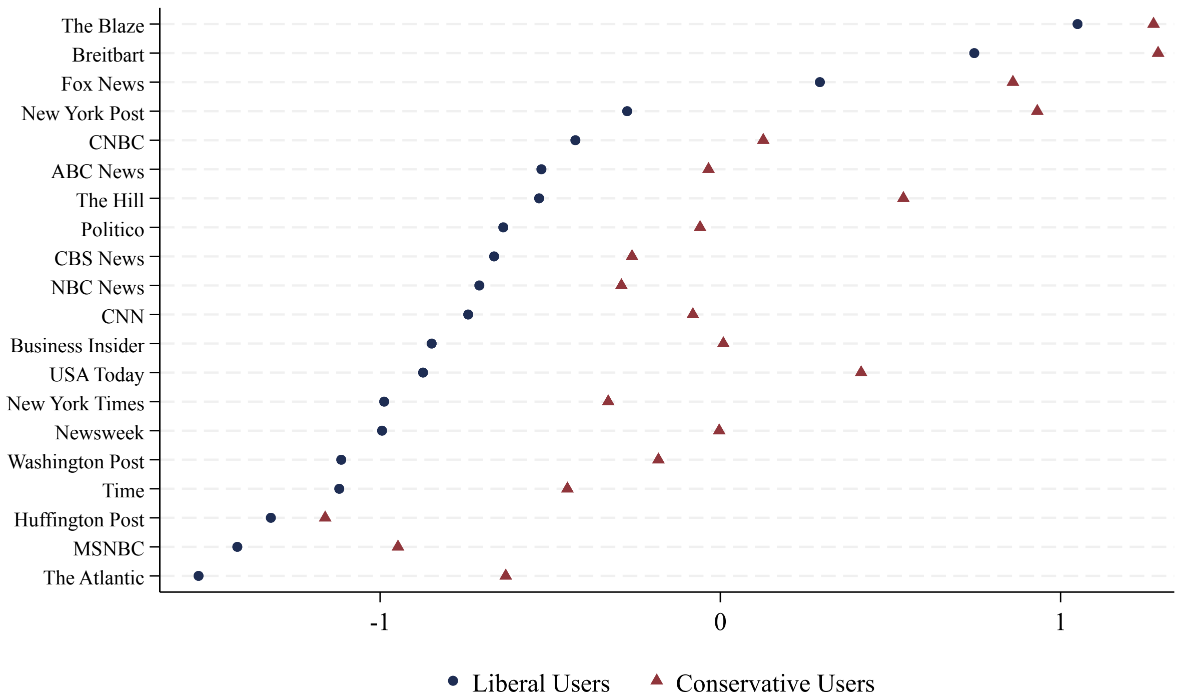

Notes: This figure shows the shares-weighted average slant of articles by user ideology and outlet. Sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. shared at least 100 times on Facebook). Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure A.13: Average Slant of Articles Shared with (no) Clicks on Facebook by Outlet and User Ideology

(a) Shares with no Clicks

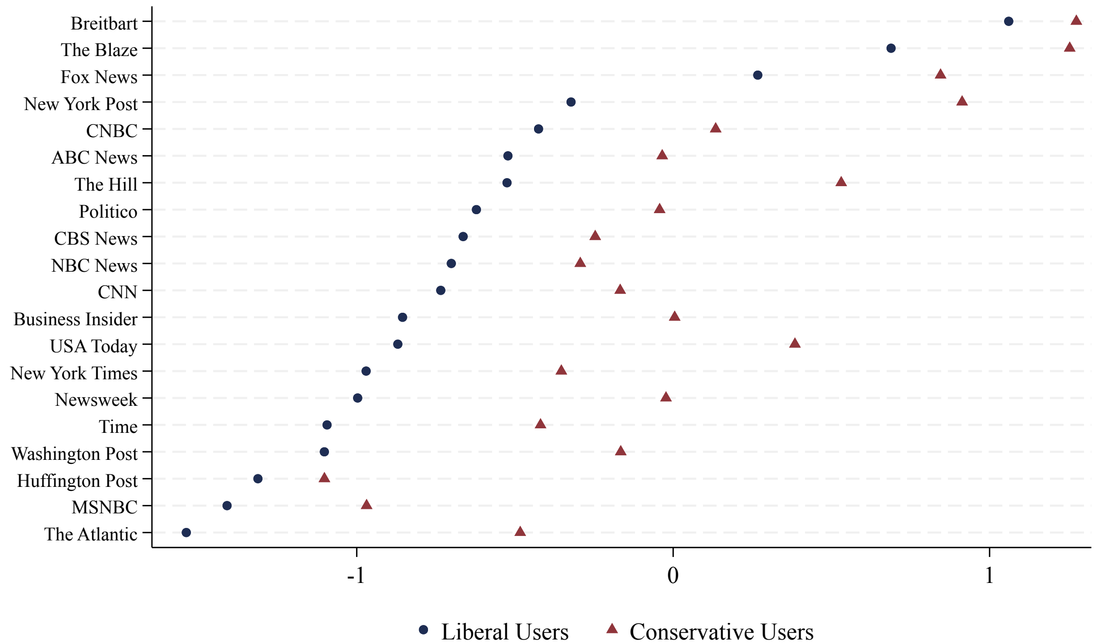

(b) Shares with Clicks

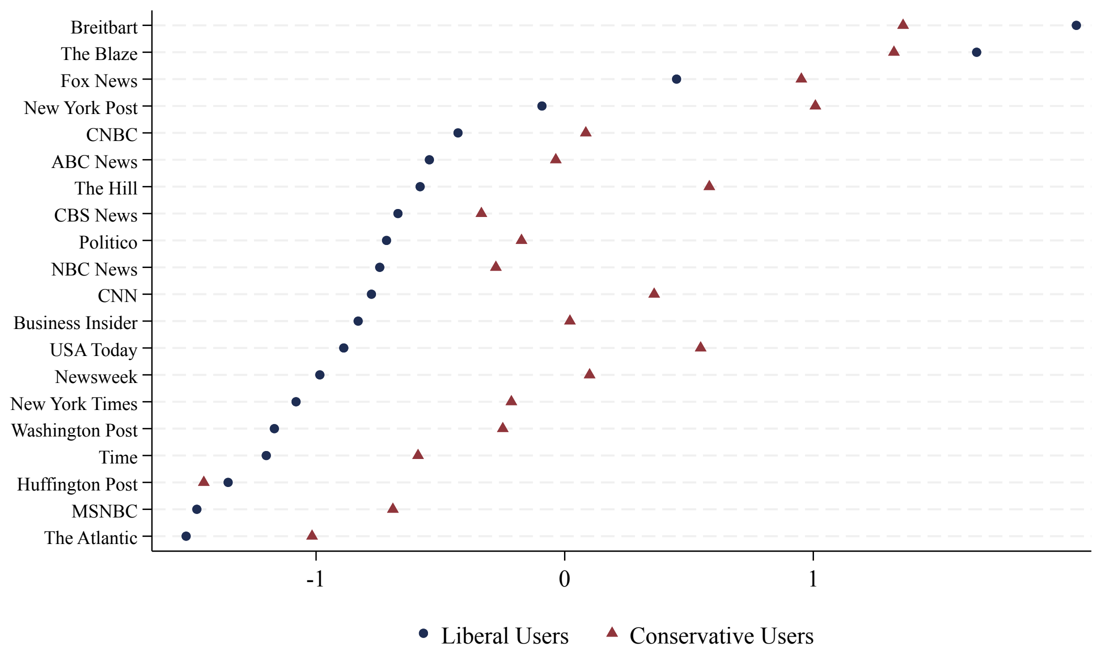

Notes: This figure shows the shares-with-no-clicks-weighted (panel a) and shares-with-clicks-weighted (panel b) average slant of articles by user ideology and outlet. Sample includes all $N = 242,829$ hard news articles published online by the top 20 U.S. news outlets in 2019 that are in the SS1 dataset (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 -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure A.14: Polarization in News Exposure, Consumption, and Circulation on Facebook Measured at Article and Outlet Level

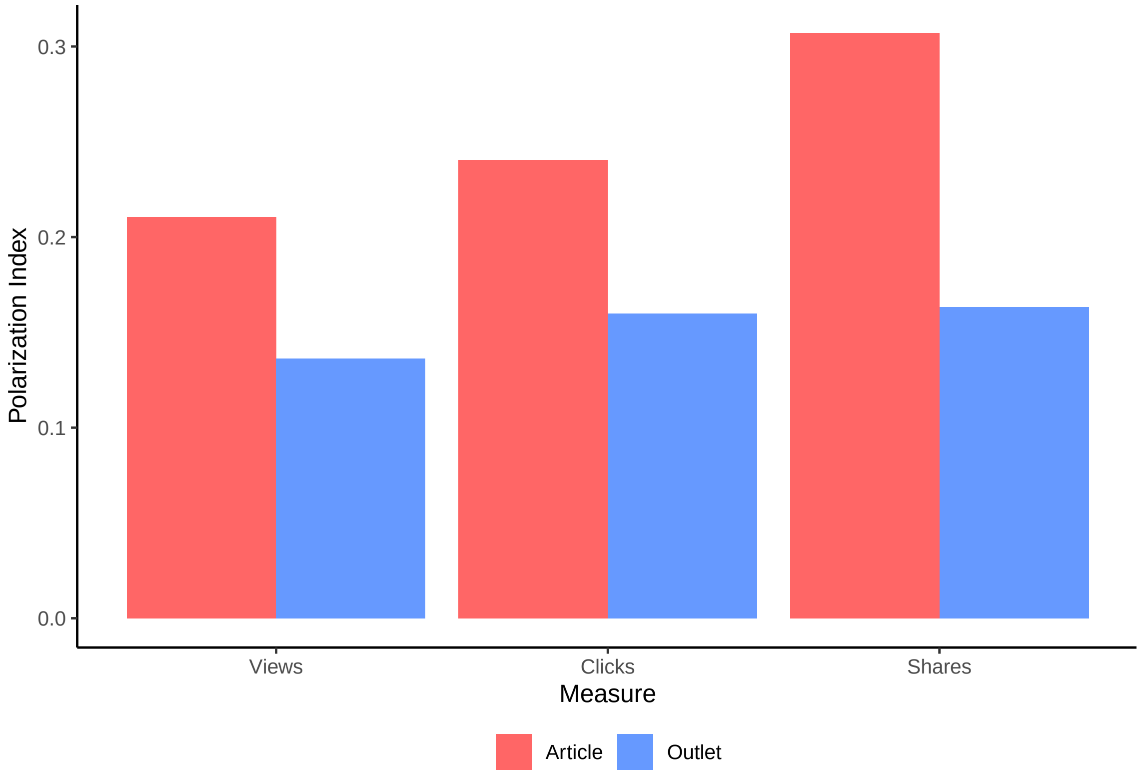

Notes: This figure shows the views/clicks/shares average difference in the slant of news diets by liberal and conservative users on Facebook divided by 6, such that 0 means no difference and 1 means maximal difference. The red bars on the left show the polarization measure when slant is calculated using our main article-level measure (as in Table 2). The blue 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 $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 dataset.

Figure A.15: News Diet Polarization Using Demographic Ideology Proxy

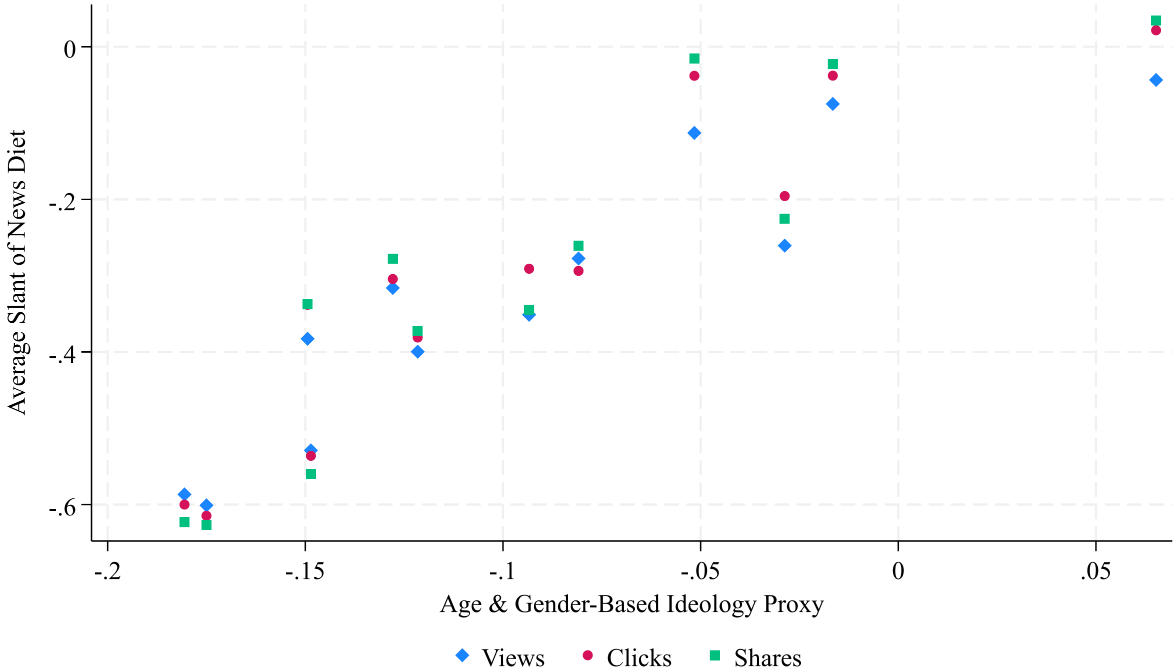

Notes: This figure shows the mean slant of a given Facebook user group's news diet (separately for views, clicks, and shares) on the vertical axis, for 12 different user groups sorted according to the group's mean ideology proxy shown on the horizontal axis. User groups are given by gender-by-age cells, with two genders and six age groups. Ideology proxy is 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), 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.

Figure A.16: Polarization by Extending Beyond Top 100 Outlets

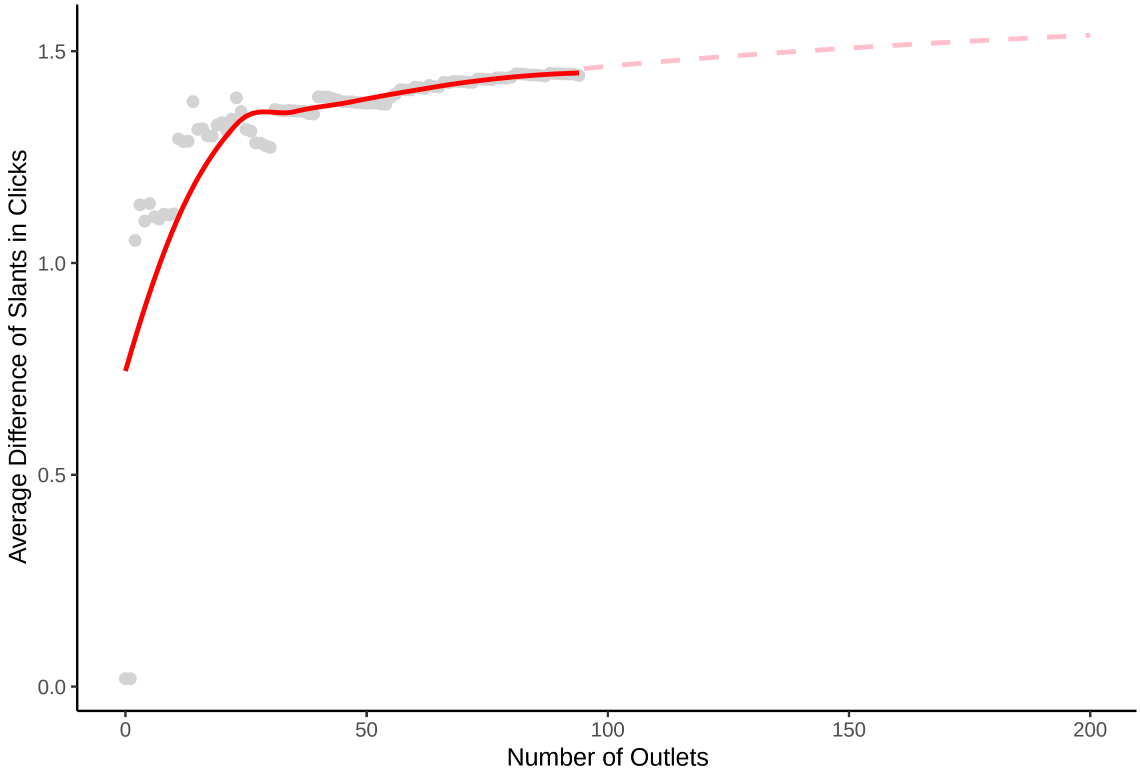

Notes: This figure presents the change of clicks-weighted average absolute difference slant of news diets by liberal and conservative users on Facebook by extending our analysis beyond top 100 U.S. news outlets as discussed in Section 7.1. The news article sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. shared at least 100 times on Facebook).

Figure A.17: Polarization by Total Shares

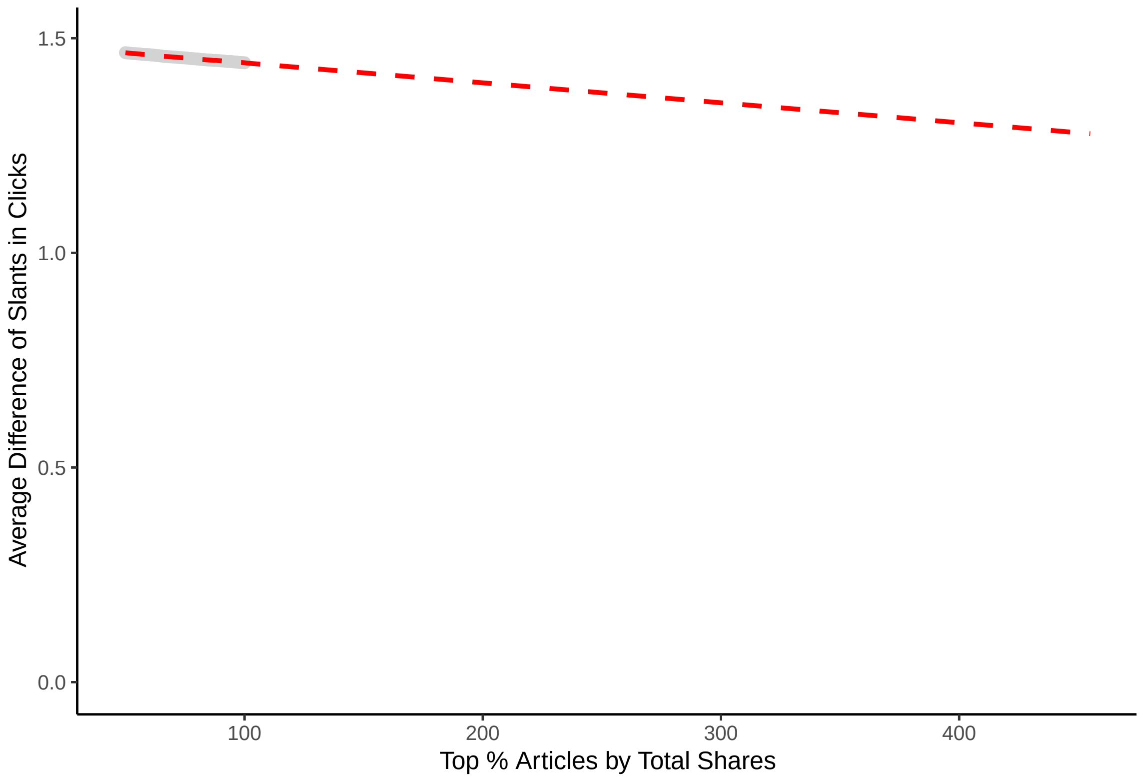

Notes: This figure shows the change of clicks-weighted average absolute difference slant of news diets by liberal and conservative users on Facebook by total shares as discussed in Section 7.3. The news article sample includes all $N = 242,829$ hard news articles published online by the top 100 U.S. news outlets in 2019 that are in the SS1 dataset (i.e. shared at least 100 times on Facebook).

B. Label Construction

We provided the 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 a political scientist called Keith Poole). 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. 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 at [date]. It was published at the following url: [URL].

We want to decide if this article is an opinion piece or not. We 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, 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, $Var\left(\varepsilon_{i,j,k}|O_i,A_{i,j}\right) = E\left(\varepsilon^2_{i,j,k}|O_i,A_{i,j}\right)$, $Var\left(A_{i,j}|O_i\right) = E\left(A^2_{i,j}|O_i\right)$, and $Var\left(O_i\right) = E\left(O^2_i\right)$.

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

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

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:

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

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

E. Maximum Achievable Correlation 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. As shown at the end of this appendix, such upper bound underlies the claim in Section 4.3 that our model attains at least 95% of the maximum achievable correlation.

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

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

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 0.86.

Therefore, our model achieves at least $\frac{0.86}{0.91} = 0.95$, that is 95%, of the maximum achievable correlation between our model prediction and our slant labels.

F. Properties of Polarization and Segregation Measures

In 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, 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

$$E[\mathscr{P}_{article}] = \frac{\frac{1}{2}\{E[\mathbf{X}^R] \cdot \mathbf{S}_a'\} - \frac{1}{2}\{E[\mathbf{X}^D] \cdot \mathbf{S}_a'\}}{6} = \frac{1}{12}\{E[\mathbf{X}^R] - E[\mathbf{X}^D]\} \cdot \mathbf{S}_a'$$

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

$$E[\mathscr{P}_{outlet}] = \frac{\frac{1}{2}\{E[\mathbf{Y}^R] \cdot \mathbf{S}'_o\} - \frac{1}{2}\{E[\mathbf{Y}^D] \cdot \mathbf{S}'_o\}}{6} = \frac{1}{12}\{E[\mathbf{Y}^R] - E[\mathbf{Y}^D]\} \cdot \mathbf{S}'_o$$

But then,

$$E[\mathscr{P}_{article} - \mathscr{P}_{outlet}] = \frac{1}{12}\{E[\mathbf{X}^R] - E[\mathbf{X}^D]\} \cdot \mathbf{S}'_a - \frac{1}{12}\{E[\mathbf{Y}^R] - E[\mathbf{Y}^D]\} \cdot \mathbf{S}'_o$$

The expression above can be simplified to:

$$E[\mathscr{P}_{article} - \mathscr{P}_{outlet}] =$$
$$\frac{1}{12}\left[\underbrace{(s_2 - s_1)}_{>0}\left[\left(p_1^D - p_2^D\right) - \left(p_1^R - p_2^R\right)\right] + \underbrace{(s_4 - s_3)}_{>0}\left[\left(p_3^D - p_4^D\right) - \left(p_3^R - p_4^R\right)\right]\right]$$

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

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

and

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

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

This shows that, in the simple example above, the segregation index at the article level is always at least as large as the segregation index at the outlet level. In fact, generically, the segregation index at the article level is always strictly larger than the segregation index at the outlet level.[32]

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.

G. 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.

[32] 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 $\mathbb{R}^8$.

Table F.1: Summary Statistics in Outlets

Outlet	Obs. All	Obs. SS1	Mean SS1 Slant All	Mean SS1 Slant Lib	Mean SS1 Slant Cons	Sd. Slant	% Lib Visits
6abc	3,744	254	-0.31	-0.32	-0.26	0.52	46.53
9news	3,181	65	-0.25	-0.44	-0.07	0.57	59.38
ABC News	21,276	2,172	-0.37	-0.47	-0.25	0.54	61.54
Abc13	3,705	336	-0.31	-0.32	-0.30	0.52	41.64
Abc7	4,937	458	-0.33	-0.35	-0.29	0.55	59.59
Ajc	11,680	873	-0.56	-0.71	-0.35	0.69	61.36
Al	4,003	1,079	-0.38	-0.54	-0.22	0.89	47.77
American Thinker	7,071	2,587	2.39	2.34	2.40	0.47	10.07
Azcentral	4,103	656	-0.69	-0.83	-0.48	0.92	63.98
Baynews9	12,753	85	-0.33	-0.33	-0.21	0.61	49.62
Bloomberg	6,384	1,415	-0.50	-0.62	-0.22	0.70	75.53
Boston	5,824	135	-0.74	-0.80	-0.73	0.67	80.91
Boston Globe	17,414	804	-0.96	-1.05	-0.80	0.86	87.58
Breitbart	41,351	15,609	1.04	1.05	1.00	1.21	2.31
Buffalonews	2,782	185	-0.37	-0.57	0.05	0.83	61.02
Business Insider	21,500	4,477	-0.57	-0.73	-0.34	0.66	68.92
Buzzfeed	1,347	426	-1.52	-1.49	-1.45	0.73	75.43
CBS News	10,750	5,359	-0.52	-0.61	-0.41	0.56	64.74
CNBC	15,321	3,629	-0.24	-0.39	0.01	0.64	63.24
CNN	24,897	14,956	-0.62	-0.67	-0.46	0.66	82.19
Cbslocal	3,878	2,618	-0.32	-0.47	-0.15	0.58	53.12
Chicago Tribune	21,007	1,616	-0.73	-0.93	-0.28	1.02	72.46
Cincinnati	2,040	262	-0.29	-0.41	0.03	0.84	60.36
Cleveland	7,994	468	-0.57	-0.70	-0.36	0.72	69.80
Click2Huston	9,699	299	-0.23	-0.33	-0.21	0.48	35.60
Clickondetroit	11,846	425	-0.32	-0.37	-0.29	0.46	56.65
Dallasnews	4,886	784	-0.48	-0.77	0.02	0.91	64.76
Detroitnews	11,079	828	-0.02	-0.42	0.64	0.97	55.10
Forbes	2,654	1,040	-0.59	-0.89	0.01	1.08	66.67
Fox News	35,210	15,603	0.57	0.54	0.54	1.10	4.93
Foxbusiness	7,908	1,304	0.61	0.61	0.60	0.88	5.98
Freep	2,739	841	-0.62	-0.64	-0.60	0.71	72.91

Table F.1: Summary Statistics in Outlets (Continued)

Outlet	Obs. All	Obs. SS1	Mean SS1 Slant All	Mean SS1 Slant Lib	Mean SS1 Slant Cons	Sd. Slant	% Lib Visits
Heavy	2,395	189	-0.71	-1.26	0.23	0.89	74.04
Hot Air	20,183	1,308	1.62	1.59	1.58	1.00	4.67
Houston Chronicle	29,961	250	-0.49	-0.65	-0.20	0.67	70.60
Huffington Post	14,004	9,801	-1.27	-1.27	-1.23	0.64	91.10
Indystar	1,572	275	-0.40	-0.61	0.11	0.71	57.81
Inquirer	5,862	618	-0.91	-1.12	-0.26	0.83	74.46
Jsonline	2,047	438	-0.48	-0.55	-0.31	0.64	71.45
Khou	3,084	123	-0.23	-0.30	-0.11	0.59	45.52
Ksl	12,216	94	-0.08	-0.26	0.06	0.67	47.26
Ktla	6,092	909	-0.45	-0.47	-0.42	0.56	56.87
LA Times	14,845	3,332	-0.99	-1.12	-0.53	0.80	82.05
MSNBC	9,113	4,453	-1.37	-1.41	-1.20	0.61	93.89
Masslive	5,264	254	-0.49	-0.63	-0.33	0.59	61.74
Mediaite	8,464	709	-0.49	-0.70	-0.30	1.13	56.45
Mlive	4,679	1,029	-0.43	-0.48	-0.32	0.62	58.27
NBC News	12,847	8,134	-0.61	-0.66	-0.48	0.65	79.03
National Review	8,847	1,990	1.57	1.16	1.62	1.07	14.59
New York Post	15,736	4,869	0.38	-0.02	0.71	1.11	31.57
New York Times	47,895	14,716	-0.82	-0.84	-0.63	0.75	90.87
Newsday	9,916	162	-0.40	-0.49	-0.32	0.78	51.14
Newser	3,627	10	-1.20	-1.18	-0.80	0.75	67.31
Newsmax	10,879	230	1.49	0.98	1.50	0.99	11.82
Newsweek	7,341	4,896	-0.83	-0.94	-0.45	0.87	81.08
Nj	6,902	901	-0.45	-0.61	-0.36	0.81	57.85
Nola	3,209	386	-0.70	-0.88	-0.55	0.74	56.12
Nydailynews	12,481	2,729	-0.72	-0.84	-0.36	0.88	78.68
Omaha	4,294	86	-0.49	-0.73	0.04	0.75	63.64
Oregonlive	5,134	859	-0.54	-0.64	-0.40	0.74	67.67
Palmerreport	4,927	2,408	-2.15	-2.13	-2.40	0.41	88.20
Patch Media	25,374	1,588	-0.37	-0.41	-0.35	0.65	53.93
Pennlive	5,028	426	-0.28	-0.57	0.29	0.84	40.61
Pjmedia	5,534	2,493	2.06	1.99	2.04	0.89	8.09
Politico	12,284	5,467	-0.50	-0.55	-0.20	0.72	82.92
Raw Story	20,852	8,378	-1.26	-1.29	-0.96	0.67	93.13
Realclearpolitics	1,477	249	2.18	0.33	2.33	1.15	7.34
San Francisco Gate	35,333	280	-0.44	-0.60	0.02	0.64	75.32
Seattle Times	16,148	696	-0.77	-0.91	-0.08	0.78	85.13
Slate	4,670	1,829	-1.44	-1.47	-1.32	0.78	95.07

Table F.1: Summary Statistics in Outlets (Continued)

	Obs.		Mean SS1 Slant			Sd.	% Lib
Outlet	All	SS1	All	Lib	Cons	Slant	Visits
Star Tribune	26,528	481	-0.45	-0.69	0.02	0.80	62.19
Staradvertiser	2,349	131	-0.48	-0.48	-0.25	0.59	78.50
Stltoday	9,234	454	-0.54	-0.61	-0.34	0.76	68.41
Syracuse	4,573	223	-0.25	-0.38	0.02	0.67	43.80
The Atlantic	3,971	2,199	-1.35	-1.36	-1.12	0.88	90.45
The Blaze	7,508	4,450	1.10	1.05	1.11	1.09	3.72
The Epoch Times	14,597	4,026	0.75	0.65	0.89	0.86	17.54
The Hill	44,844	14,940	-0.36	-0.47	0.13	0.85	79.48
Thedailybeast	4,295	1,939	-1.05	-1.09	-0.80	0.87	91.39
Thefederalist	3,393	1,410	2.05	1.73	2.09	0.72	8.96
Time	6,662	2,249	-0.91	-1.00	-0.65	0.70	80.66
Townhall	12,427	4,365	1.76	1.68	1.75	0.95	2.59
Twitchy	9,336	703	2.22	0.19	2.30	0.57	2.91
USA Today	13,568	5,447	-0.44	-0.71	0.03	0.88	61.60
Voanews	13,931	703	-0.47	-0.76	-0.20	0.37	56.80
Vox	5,578	3,035	-1.27	-1.29	-1.12	0.64	92.26
WND	14,450	2,140	1.88	1.86	1.90	0.82	3.99
WRAL	13,662	311	-0.32	-0.46	-0.14	0.62	43.88
WTOP	19,830	166	-0.61	-0.67	-0.34	0.55	69.86
Wallstreet Journal	14,194	2,815	0.44	0.23	0.77	1.00	54.92
Washington Examiner	24,278	5,226	0.79	0.31	0.85	0.97	7.87
Washington Post	38,559	14,688	-0.92	-0.98	-0.53	0.87	87.63
Washington Times	20,420	1,404	0.90	0.73	0.93	1.04	9.12
Wfaa	2,956	112	-0.17	-0.29	-0.11	0.67	46.95

Notes: This table shows the number of articles, the clicks-weighted average slant by all, liberal and conservative users, the standard deviation of slant, and the percentage of liberal visits defined by 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. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.

Figure F.1: Slant Distribution within Outlets

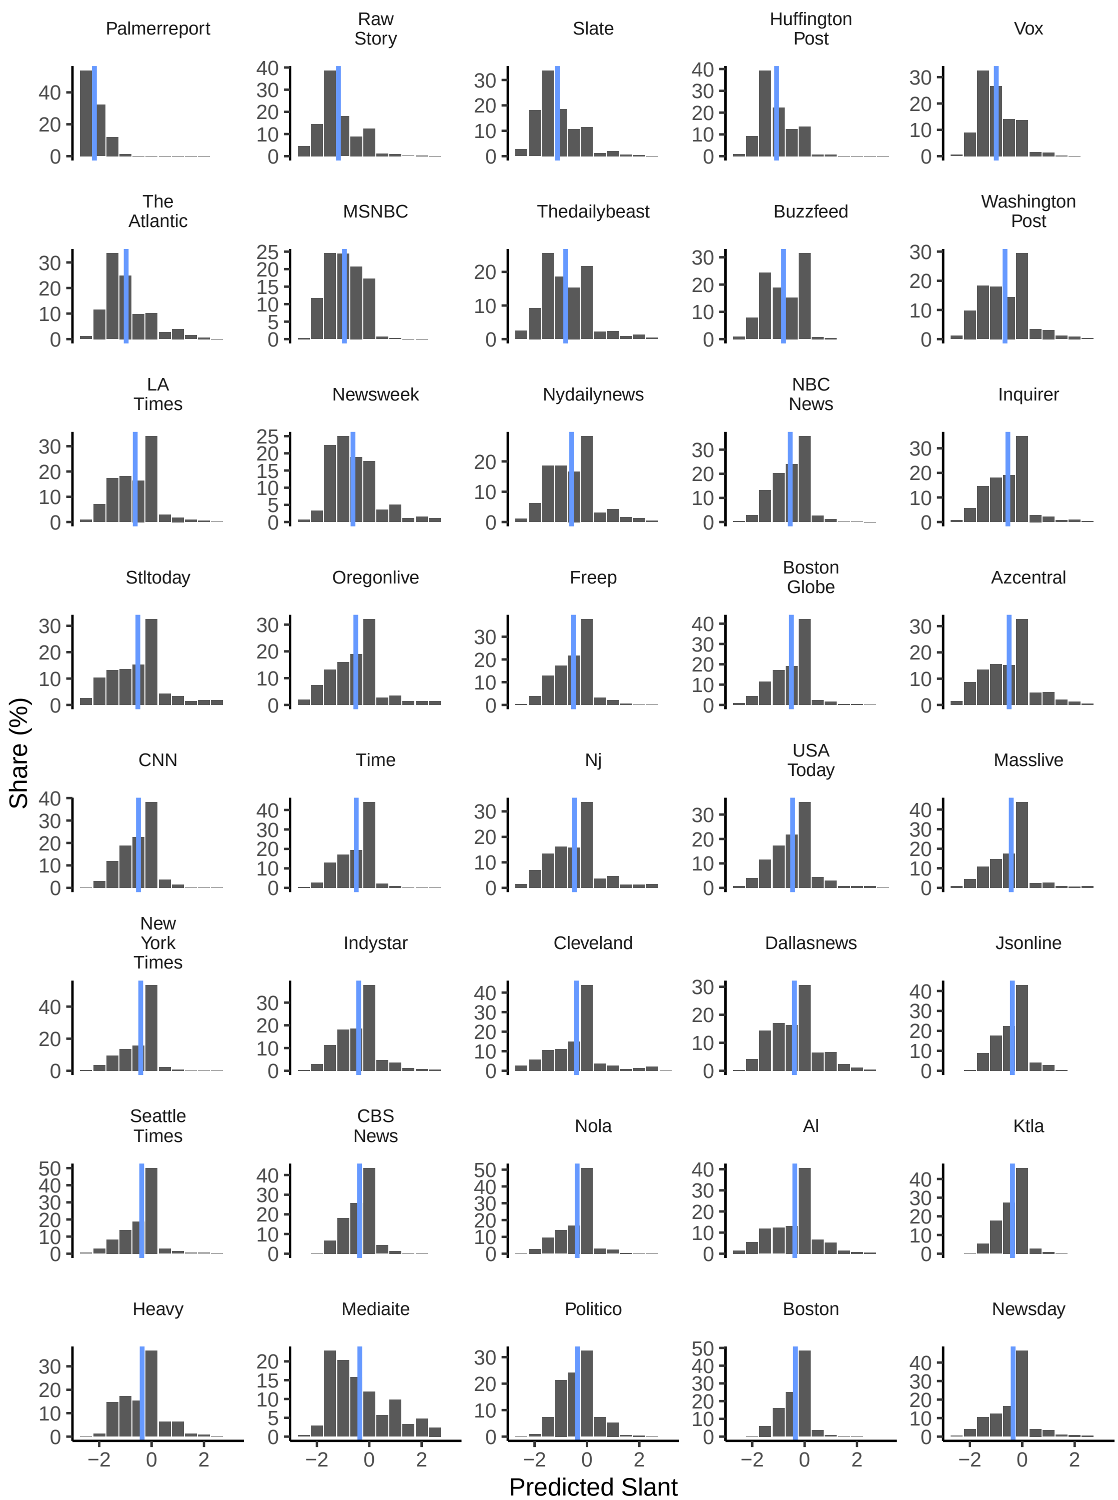

Figure F.1: Slant Distribution within Outlets (Continued)

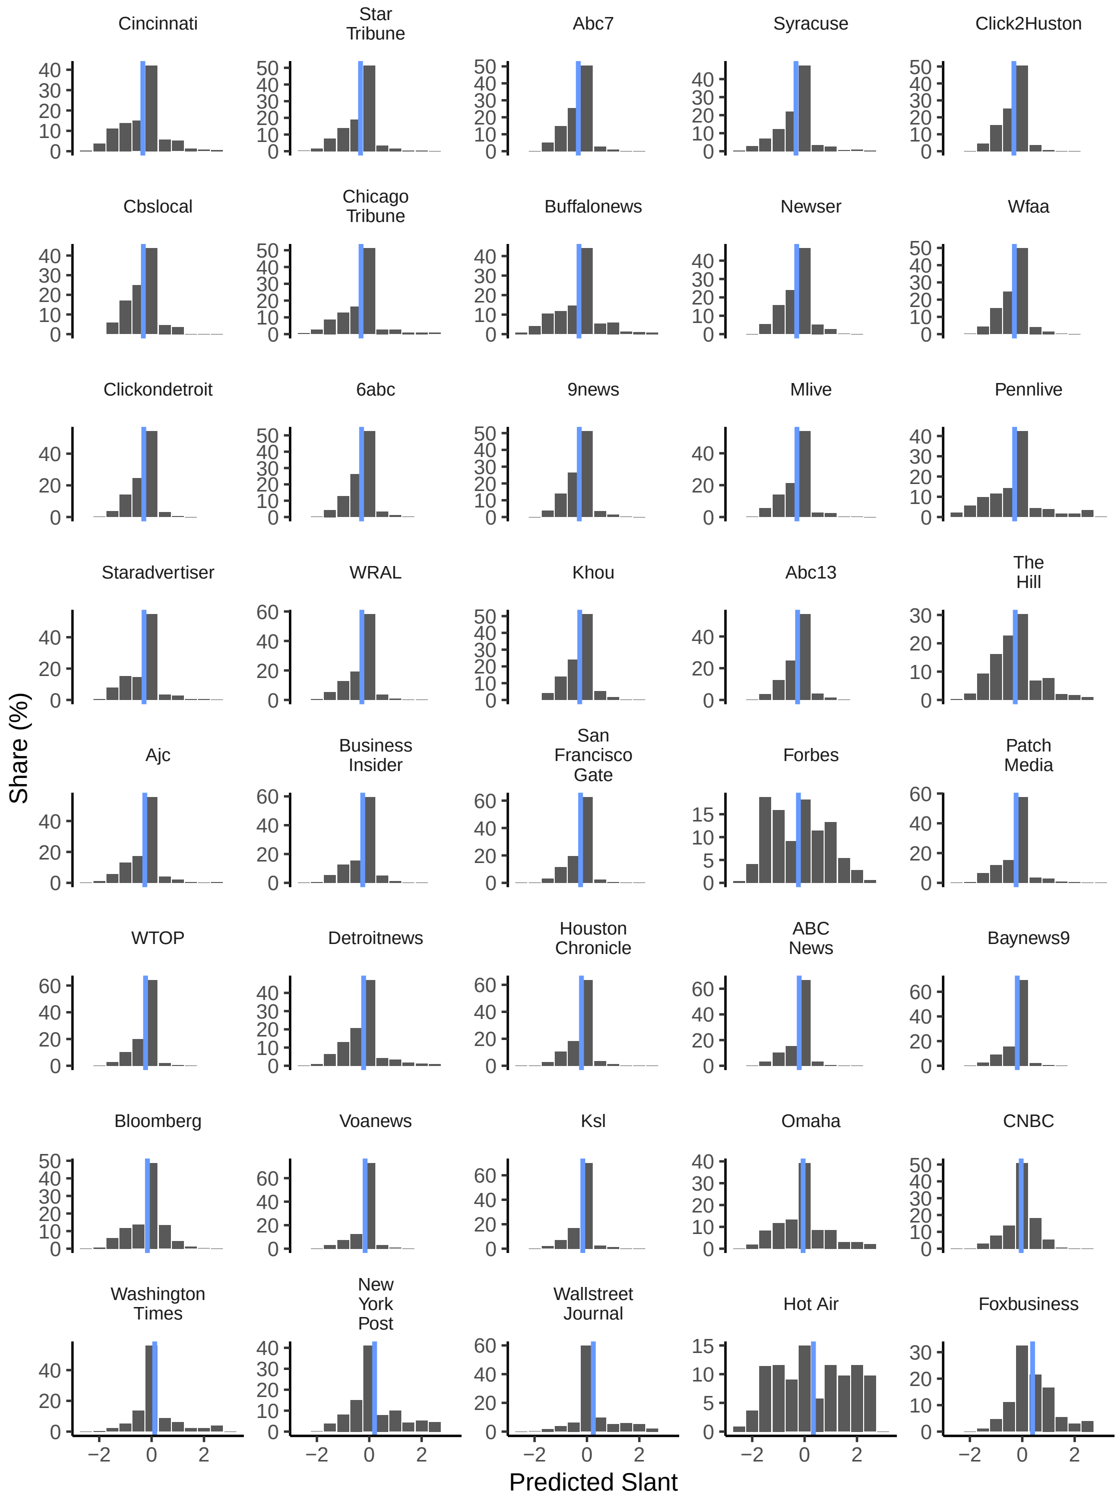

Figure F.1: Slant Distribution within Outlets (Continued)

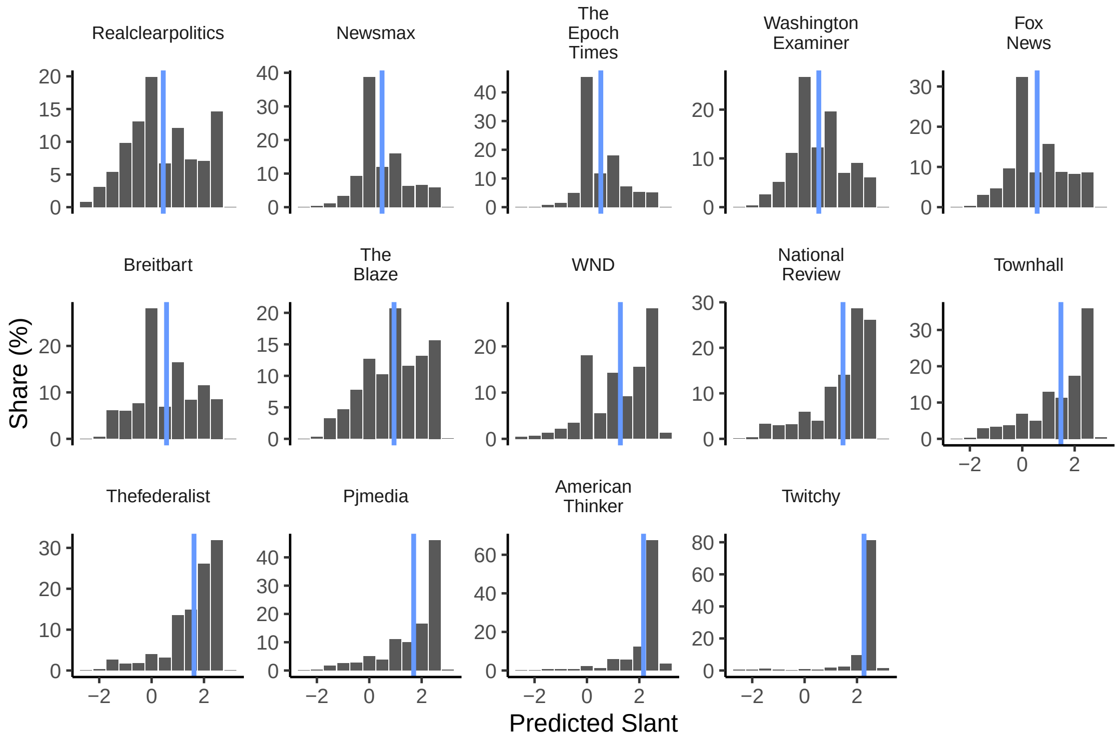

Notes: This figure shows the distribution of slant within the outlets included in our analysis. The vertical blue line indicates the mean slant for all articles within each outlet. Slant is given by the fine-tuned GPT-4o's prediction as described in Section 4.2. The slant scale goes from -3 (extremely favorable to the Democratic party) to 3 (extremely favorable to the Republican party); 0 signifies centrist.