Platforms as Editors: How Algorithmic Curation Shape Online News Slant

Yikang Shen *

Latest Version

Abstract

Digital platforms shape not only what news people see but also what outlets choose to produce. This paper quantifies the equilibrium effects of algorithmic curation on news supply and exposure. Linking print and online headlines from major U.S. newspapers, I document that online headlines are more ideologically slanted and more emotional than their print counterparts. Using Facebook's announcement of an algorithmic update as a quasi-experiment, I show that newsrooms have an immediate response to algorithm changes. A structural model shows that readers value both like-minded and unexpectedly slanted posts, but platforms overweight these forces. Media outlets respond by adjusting content to maximize platform-driven viewership, accounting for about 80% of the extra slant in user exposure. Counterfactuals indicate that roughly 90% of this extra slant can be eliminated through an algorithm redesign, substantially flattening exposure across users. These findings highlight the central role of platform in shaping both the supply and distribution of news, and a trade-off between engagement and exposure diversity in platform design.

Keywords: media economics, recommendation algorithms, polarization, dynamic games

*Ph.D. Candidate, Tepper School of Business, Carnegie Mellon University. Email: yikangs@andrew.cmu.edu. This draft: April 28, 2026. Comments welcome.

1 Introduction

Digital platforms now play a central role in shaping the information that people consume. A growing body of research shows that social media recommendation algorithms create 'echo chambers' by promoting ideologically aligned content over moderate or counter-attitudinal news (Levy, 2021, Nyhan et al., 2023, Guess et al., 2023a, Guess et al., 2023b, González-Bailón et al., 2023). In contrast, evidence on how the supply of news changes, especially in response to algorithms, is limited. Media bias has long been documented (Groseclose and Milyo, 2005), in the sense that outlets can use different linguistic framings to describe the same set of facts and thereby convey sharply divergent impressions of reality. In traditional markets, prior work has shown that outlets adjust their editorial slant in response to consumer demand for ideologically aligned content (Gentzkow and Shapiro, 2010). In this paper, I address several follow-up questions in the platform era. First, do newspapers supply more polarized content in the online market compared to the print market? If so, is this difference a response to how algorithms distribute and promote news on social media platforms? Second, once such behavior is established, how much of the polarization observed in online news can be attributed to algorithmic curation rather than consumer demand?

The answers to these questions are central to understanding how social media and online news markets should be regulated—domains that have become integral to the functioning of modern democracies. It has been well documented empirically that news content can have significant effects on political attitudes and outcomes (Strömberg, 2004, Gentzkow, 2006). In addition, slanted news has been shown to have persuasive effects on ideologies and voting behaviors (DellaVigna and Kaplan, 2007, Martin and Yurukoglu, 2017, Alesina et al., 2020). Particularly on social media platforms, more exposure to content expressing antidemocratic attitudes or partisan animosity has been shown to increase affective polarization and elevate users' negative emotions (Piccardi et al., 2025). When such content is more prevalent on the platform, users are more likely to encounter and consume it, amplifying risks of polarization and psychological distress. These patterns underscore the importance of understanding how

news supply responds to algorithmic design.

The central findings of this paper are twofold. First, online news headlines are on average more liberally slanted and express more negative emotion than their print counterparts across major U.S. newspapers—except for The Wall Street Journal, whose online headlines are comparatively more conservative. A significant part of this divergence between online and print content can be attributed to algorithmic curation on social media platforms. Second, users display a clear preference for ideologically slanted content, and recommendation algorithms amplify these preferences by disproportionately promoting such material. In turn, media outlets strategically increase the slant of their online content to maximize the probability of being recommended by algorithms. Counterfactual simulations show that removing algorithmic amplification would reduce the extra slant observed in online news by approximately 90 percent.

To generate these results, I assemble a new dataset by matching print and online headlines from eleven major U.S. newspapers between 2017 and 2024, and merging them with the corresponding Facebook posts with consumption measures. Then, I measure each headline's ideological slant using a phrase-frequency measure based on the *Congressional Record* (following Gentzkow and Shapiro, 2010), and its emotional level as sentiment using a pretrained machine learning model. I focus on the posts and the headlines, instead of the articles linked to them, as those are the main products consumed by social media users. As a direct evidence, Sundar et al. (2025) finds that around 75% of the Facebook posts with links to news were shared without clicking into the linked web pages. And such phenomenon is more serve for extreme and user-aligned political content.

I begin by providing direct evidence that media outlets adjust their ideological slant in response to algorithmic incentives. Across all major newspapers, online headlines are on average more negative and use more partisan language than their print counterparts. However, this difference alone cannot distinguish whether it arises from user demand or algorithmic influence. To isolate the latter, I exploit Facebook's July 19, 2022 announcement

that it had implemented an algorithmic change to reduce political content in users' feeds as a quasi-experimental shock to platform incentives. Following the announcement, online headlines from U.S. newspapers became significantly closer in slant to their print versions, while sentiment differences remained unchanged. This event study provides direct causal evidence that social-media algorithms shape the ideological supply of news content online.

Then, to quantify how much of the extra slant is driven by algorithms, I develop a structural model that jointly estimates user preferences, algorithmic decision rules, and media outlets' dynamic strategies for slant provision. With the model, I show that readers prefer both like-minded and unexpected news. They are more likely to engage with posts whose slant matches their own ideological orientation, consistent with Gentzkow et al. (2014) and Nyhan et al. (2023). At the same time, they also respond positively when a familiar outlet displays an atypical ideological or emotional tone. For example, when New York Times publishes a more conservative post, or when Fox News publishes a more liberal one. This reaction reflects the salience of surprise, a key driver of attention documented by Bordalo et al. (2013) and Bordalo et al. (2022). Such cross-ideological content can also appear more trustworthy, similar to the mechanism in Chiang and Knight (2011), who show that partisan endorsements are less persuasive than endorsements from neutral or ideologically unexpected sources. Treating the platform algorithm as an abstract decision maker that selects which news to display, I estimate that Facebook's recommendation system effectively tripled the marginal return to politically slanted content relative to what would arise from user preferences alone, accounting for about 20% of slant in user exposure. On the supply side, newspapers choose the ideological slant of their social-media posts to maximize the probability of being recommended to users, with smaller outlets placing greater weight on visibility than larger ones. Counterfactual simulations show that removing the algorithm's amplification of ideological signals would eliminate roughly 90% of the extra polarization between online and print headlines. More importantly, such a policy substantially flattens the distribution of content exposure across users with different ideologies while generating only

modest welfare losses. In contrast, an intervention similar to Facebook's algorithm update produces comparable welfare losses but has limited effects on exposure diversity. Together, these results show that recommendation systems do not merely mediate information—they actively shape the equilibrium supply of news.

Methodologically, this paper extends tools from structural industrial organization to the domain of media economics. Because the market includes many competing outlets, I adopt the concept of an oblivious equilibrium (Weintraub et al., 2008), which allows each outlet to optimize against aggregate market conditions without tracking rivals individually—a tractable and behaviorally realistic framework for large digital platforms. In traditional media markets, like cable news or newspapers in a given geographic area, typically there are two to three major outlets. However, in the online market like Facebook, all outlets–cable, broadcast, newspapers, influencers–are competing with each other, and none of which has a dominant share in the market. For estimation, I extend the share inversion method of Berry et al. (1995) to a setting with algorithmic decision and two demand layers: one for recommendation and one for engagement. Then I combine it with a simulated conditional choice probability approach (Hotz et al., 1994) for the dynamic editorial problem. Together, this framework provides the first quantitative decomposition of how much of online polarization arises from user demand versus algorithmic design.

This paper contributes to four strands of literature.

First, it relates to a large literature on the demand for slanted media, which shows that consumers prefer ideologically congenial information and that media firms respond to these preferences in equilibrium (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006; Gentzkow and Shapiro, 2010; Martin and Yurukoglu, 2017; Chiang and Knight, 2011; Guess, 2021). Recent work using article-level data emphasizes that much of the variation in slant arises within outlets rather than across them, highlighting the importance of studying content at the level of individual articles (Braghieri et al., 2024). Building on this insight, I show that user engagement reflects not only ideological alignment but also surprise—content

that deviates from an outlet's typical positioning attracts disproportionate attention. More importantly, I demonstrate that these demand-side forces do not only shape consumption, but—through algorithmic curation—feed back into the equilibrium supply of news slant.

Second, the paper contributes to the literature on algorithms, social media, and news dissemination. Existing work shows that ranking and recommendation systems shape exposure, engagement, and polarization by determining which content users see (Bakshy et al., 2015; Levy, 2021; Calzada and Gil, 2020; Moehring, 2024). This paper shifts the focus from exposure to production. Exploiting Facebook's 2022 reduction in political content as a quasi-experimental shock, I provide direct evidence that media outlets adjust the slant of their online headlines in response to changes in platform incentives. I then develop a structural model that separately identifies user preferences, algorithmic amplification, and outlet responses. This allows me to quantify how recommendation algorithms affect not just which articles are seen, but which articles are written.

Third, this paper contributes to the literature on news production, editorial strategy, and media competition. Prior work shows that media firms tailor content in response to consumer demand, competitive pressures, and technological changes (Gentzkow et al., 2014; Allcott and Gentzkow, 2017; Cagé et al., 2019; Leung and Strumpf, 2023). Recent work also highlights the influence of social media on mainstream media coverage decisions (Cagé et al., 2024), showing that events receiving more attention on social media are more likely to be covered by traditional media. In contrast, this paper focuses on the intensive margin of editorial choice. Conditional on covering a given event, I show that media outlets systematically adjust the language and ideological slant of content to capture attention from platform algorithms and online users. To study this, I develop a dynamic, article-level model of editorial choice in which outlets optimally choose slant for each piece of content. Editorial decisions trade off algorithmic visibility, user engagement, and the evolution of outlet reputation over time. This framework allows me to quantify the equilibrium effects of platform algorithms on both the supply and distribution of news, and to evaluate counterfactual platform designs. In this

sense, platform algorithms become a key input into newsroom decision-making, analogous to prices, advertising demand, or competition in earlier media environments.

Finally, the paper contributes to the literature on platform design and information provision with strategic responses (Bergemann and Bonatti, 2019; Armona et al., 2024; Saeedi et al., 2024). This literature studies how intermediaries design information environments when agents respond strategically, but typically abstracts from endogenous content creation. I incorporate both the platform and content producers in a unified equilibrium framework, combining a two-layer demand system for recommendation and engagement with a dynamic supply model of media outlets. This structure allows the data to discipline the relative roles of user preferences and algorithmic amplification, and to quantify how platform design shapes equilibrium exposure diversity and ideological slant.

Taken together, the evidence suggests that social-media algorithms play a quantitatively critical role in shaping the supply of news in the digital era. By amplifying the engagement returns to ideologically charged content, platform design shifts the incentives of professional journalism toward polarization. The findings highlight how the architecture of online platforms, rather than only consumer preferences, contributes to the transformation of the modern news ecosystem.

2 Institutional Overview: News on Social Media

This section summarizes how news distribution and consumption on social media evolved from 2017 to 2024 and how major U.S. newspapers adapted organizationally by hiring search engine optimization (SEO) editors and social media editors.

2.1 From Web & Print to Social Distribution

In 2017, social media was already central to news access: two-thirds of U.S. adults reported getting at least some news from social platforms, with 20% doing so often.[1] Platform composition changed over time. Using a 2023 Pew survey reported in 2024, about 30% of U.S. adults said they regularly get news on Facebook, with smaller shares on Instagram (16%), TikTok (14%), and X/Twitter (12%).[2] By late 2024, Pew additionally documents the rise of "news influencers": about 21% of U.S. adults (and 37% of those aged 18–29) report regularly getting news from influencers on social media.[3] These trends underscore a structural shift: audience discovery increasingly occurs in feeds governed by platform algorithms, not on publisher homepages.

2.2 Organizational Responses in Legacy Newsrooms

Major newspapers responded by creating and professionalizing roles dedicated to platform discovery (SEO) and off-platform packaging and distribution (social media editing). While precise first-hire dates are not always publicly archived, the pattern is consistent: SEO and audience roles diffused widely in the mid-to-late 2010s and expanded again with short-form video in the early 2020s. The typcial job descriptions include optimizing headlines and metadata for search intent, identifying trending, and monitoring ranking performance.

The New York Times named its first Social Media Editor in May 2009 (Jennifer Preston), formalizing a newsroom role to use social tools to find sources, shape coverage, and engage audiences.[4] By 2023, the Times advertised an *Associate SEO Editor* to "maintain visibility on the biggest news stories of the day," provide search recommendations, and plan

[1] Pew Research Center (Sept. 7, 2017), "News Use Across Social Media Platforms 2017." https://www.pewresearch.org/journalism/2017/09/07/news-use-across-social-media-platforms-2017/.

[2] Pew Research Center (June 12, 2024), "How Americans Get News on TikTok, X, Facebook and Instagram." https://www.pewresearch.org/journalism/2024/06/12/how-americans-get-news-on-tiktok-x-facebook-and-instagram/.

[3] Pew Research Center (Nov. 18, 2024), "America's News Influencers" (report + methodology). Report overview: https://www.pewresearch.org/journalism/2024/11/18/americas-news-influencers/.

[4] The Guardian (May 26, 2009), "New York Times names first social media editor." https://www.theguardian.com/media/pda/2009/may/26/new-york-times-twitter.

around anticipated events—indicating a mature editorial search function.[5] I also document newsroom changes of other major newspapers in Section B.

Between 2017 and 2024, news discovery shifted toward feeds and influencers, with platform composition diversifying (Facebook remained largest; TikTok rose sharply among younger adults). These facts rationalize newsroom investments in SEO/social roles and help explain the supply-side adjustments modeled in this paper. As platforms changed recommendation policies, newsrooms adapted content packaging to preserve visibility—consistent with the mechanisms I estimate in the structural model.

2.3 Platform Evolution of Facebook

Since 2017, Facebook's algorithm and its role as a news distributor have undergone substantial transformations—both in how content is ranked and in how many Americans rely on the platform for news.

In the years following 2017, Facebook's News Feed algorithm increasingly prioritized posts that generated engagement—such as comments, shares, and reactions—and those that fostered "meaningful social interactions" among users. In January 2018, Facebook formally announced a major algorithm update aimed at shifting attention away from passive consumption of public content toward posts from friends, family, and community groups.[6] This adjustment substantially reduced the visibility of news content, particularly from publishers that relied on organic distribution through the platform.

Beginning in 2021, Facebook initiated a series of experiments designed to "reduce the amount of political content" shown in user feeds.[7] These tests were expanded in 2022, when Meta reported that "placing less emphasis on shares and comments for political content is

[5] Talking Biz News (Aug. 10, 2023), "The New York Times seeks an associate SEO editor." https://talkingbiznews.com/biz-news-help-wanted/the-new-york-times-seeks-an-associate-seo-editor/.

[6] *TechCrunch* (Jan. 28, 2018), "How Publishers Will Survive Facebook's News Feed Change." https://techcrunch.com/2018/01/28/how-publishers-will-survive-facebooks-newsfeed-change/.

[7] Meta (Feb. 10, 2021), "Political Content in Feeds." https://about.fb.com/news/2021/02/reducing-political-content-in-news-feed/.

an effective way to reduce the amount of political content shown to users." As a result, engagement-driven amplification of political news fell sharply, and many outlets observed measurable declines in reach for politically oriented posts. For news organizations, these changes implied that visibility was increasingly determined by platform algorithms rather than editorial priorities—effectively tying exposure to engagement metrics. This evolution illustrates Facebook's ongoing experimentation with balancing engagement optimization, user satisfaction, and public scrutiny over algorithmic bias.

According to the Pew Research Center, about 30% of U.S. adults regularly received news from Facebook in 2023.[8] This share has declined from earlier years: in 2017, roughly 45% of Americans reported getting news from Facebook.[9] At the same time, as of 2024, approximately 54% of U.S. adults say they sometimes get news from social media platforms more broadly.[10] These statistics indicate that while Facebook remains a major gateway to news, its dominance has stabilized or modestly declined as users diversify toward other platforms such as YouTube, X (formerly Twitter), and TikTok.

Together, these trends highlight how algorithmic design and media incentives evolve jointly. Facebook's periodic recalibration—from prioritizing engagement (2017–2021), to demoting political content (2022–2024)—created shifting incentives for news outlets. Each adjustment affected how publishers frame and distribute content on social media, motivating this paper's central analysis of how algorithmic objectives translate into observable changes in news supply and polarization.

[8]Pew Research Center (June 12, 2024), "How Americans Get News on TikTok, X, Facebook and Instagram." https://www.pewresearch.org/journalism/2024/06/12/how-americans-get-news-on-tiktok-x-facebook-and-instagram/.

[9]Pew Research Center (Sept. 7, 2017), "News Use Across Social Media Platforms 2017." https://www.pewresearch.org/journalism/2017/09/07/news-use-across-social-media-platforms-2017/.

[10]Pew Research Center (June 12, 2024), "Social Media and News Fact Sheet." https://www.pewresearch.org/journalism/fact-sheet/social-media-and-news-fact-sheet/.

3 Data

This section describes the construction of the dataset that links print and online newspaper headlines to social media posts and user engagement measures. The final dataset combines multiple large-scale sources—ProQuest, the New York Times Developers API, Meta's Content Library, and the U.S. Congressional Record—to analyze how algorithms influence news production and online presentation. The sample spans April 2017 to December 2024 and includes major U.S. news outlets and prominent news influencers.

3.1 Newspaper Headlines: ProQuest and NYT Developers

The core of the dataset consists of matched print and online headlines for major U.S. newspapers (including all 7 newspapers with more than 300,000 combining circulation and online subscriptions in 2023). For the New York Times, through the *NYT Developers API*, precise metadata linking online publication time, URL, and the corresponding full print version information is provided. For other newspapers, both online and print headlines are obtained from the *ProQuest TDM Database*, which provides text and metadata for newspaper articles, including publication date, section, and full article text. Together, I acquire 2.7 million online headlines under the name of each outlet's official website and 1.8 million print headlines separately under each newspaper itself.

To construct a one-to-one match between print and online versions of the same article, I match records by the first 300 characters of article body text after removing boilerplate and punctuation. In the case of multiple headlines for the same news article, I use the latest one as it reflects the editor's final decision. The resulting dataset allows direct comparison between the print and the online headlines of each article. This alignment makes it possible to measure how editorial tone and ideological framing differ between print and online dissemination.

3.2 Facebook Posts: Meta Content Library

The social media component of the dataset comes from Meta's *Content Library*, which provides comprehensive post-level data on all publicly available Facebook pages. From this source, I collect approximately 3 million posts from 23 major media outlets and 11 top news influencers with more than three million followers. For each post, the dataset includes the posting date, textual content, URL link to the associated article, and engagement metrics—number of views, shares, comments and reactions.

The outlets represented in the sample span all major media segments:

- **National Newspapers:** New York Times, Washington Post, USA Today, Wall Street Journal, The Guardian;
- **Regional Newspapers:** Boston Globe, Chicago Tribune, Los Angeles Times, Orlando Sentinel, Tampa Bay Times, Baltimore Sun;
- **Digital Newspapers:** Voice of America, POLITICO, Bloomberg, NPR, Reuters;
- **Broadcast:** ABC News, CBS News, NBC News, BBC News;
- **Cable:** Fox News, CNN, MSNBC;
- **Influencers:** Heather Cox Richardson, Tomi Lahren, Ben Shapiro, Glenn Beck, Robert Reich, Charlie Kirk, Dan Bongino, Sean Hannity, Graham Allen, Terrence K. Williams, Jay Sekulow.

Each post is matched to the corresponding article in the newspaper dataset using the URL link when possible.

3.3 Congressional Record and Political Slant

To measure ideological slant, I extend the dictionary-based measure from Gentzkow and Shapiro (2010). The original index quantifies how likely each phrase is to be spoken by

Democrats versus Republicans in the U.S. Congress. I reconstruct the measure using *Congressional Record transcripts*[11] from 2017 to 2024 following the method of phrase extraction in Gentzkow et al. (2019) and update the phrase-level log-odds ratios for each year to reflect contemporary political language.

Each article or post is assigned a slant score computed as the mean of phrase-level weights. Positive values indicate a right-leaning language, negative values indicate a left-leaning language, and values near zero correspond to neutral phrasing. Because both print and online headlines are processed through the same metric, the difference between them for the same article directly captures the shift in ideological framing between platforms.

This measure serves as the key input for both the reduced-form difference-in-differences regressions and the structural estimation of user and algorithmic preferences.

3.4 Topics and Sentiment: Pretrained NLP Model

To complement the slant measure, I estimate the sentiment and topical content of each headline or post using modern transformer-based language models. Sentiment scores are produced using the `Twitter-RoBERTa-base` model[12], a large-scale transformer pretrained on 124 million tweets and finetuned for sentiment classification. Although it was pretrained on tweets, it has been shown to perform well on short text like news headlines. This model outputs a sentiment score in the range $(-1, 1)$, where -1 indicates negative tone, 1 indicates positive tone, and values near zero correspond to neutral statements. Because it is trained on short, informal text, this model is well-suited to classify the tone of headlines and social-media posts.

To identify whether posts are news-related, I employ the `tweet-topic-21-multi` classifier[13]—another transformer model trained on the same corpus but finetuned for multi-label topic classification. This filter removes entertainment, sports, and promotional posts, en-

[11] Files obtained from https://www.congress.gov/congressional-record.
[12] See https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest.
[13] See https://huggingface.co/cardiffnlp/tweet-topic-21-multi.

suring that the remaining observations represent genuine news coverage. See Table 8 for a summary of how Facebook engagement metrics on different topics changed.

The sentiment measure provides a fine-grained view of how outlets frame their news across platforms and how these framing choices evolve after algorithmic interventions. It also serve as a comparison and robust check to the different editorial behaviors for slant.

4 Combined Dataset

The merged dataset thus integrates the print and online newspaper headlines for leading U.S. outlets. Then given Facebook posts and engagement metrics (Meta Content Library), the corresponding online and print headlines are connected to them through the url links.

As motivation, here I show you how the monthly average of sentiment and slant change over time for articles with both online and print headlines. As shown in Figure 1, there are systematic differences between the two measures. On average, online headlines have more negative sentiment than print headlines, which matches the conventional wisdom that negative events tend to be more newsworthy, as tested in Armona et al. (2024). On the other hand, online headlines are more liberal slanted than the print counterparts. Given that most major newspapers in the US (and thus in the data) are liberal leaning, it might suggest that newspapers provide more polarized content online.

Next, I provide an example of the Facebook post, online headline and print headline of the same article from the New York Times. As in the example and in general, the posts and online headlines are similar to each other. But they can be quite different from the print headlines, both in terms of slant and sentiment. Therefore, in the following analysis, I will focus on the difference between the posts and the print headlines.

Format	Headlines of the Same Article	Sentiment	Slant
Post	Republicans' 4-Step Plan to Repeal the Affordable Care Act	0	-7.17
Online	Republicans' 4-Step Plan to Repeal the Affordable Care Act	0	-7.17
Print	Republicans' Four-Step Plan for Dismantling the Care Act	0	0

Table 1: Post, Online and Print Headlines from NYT

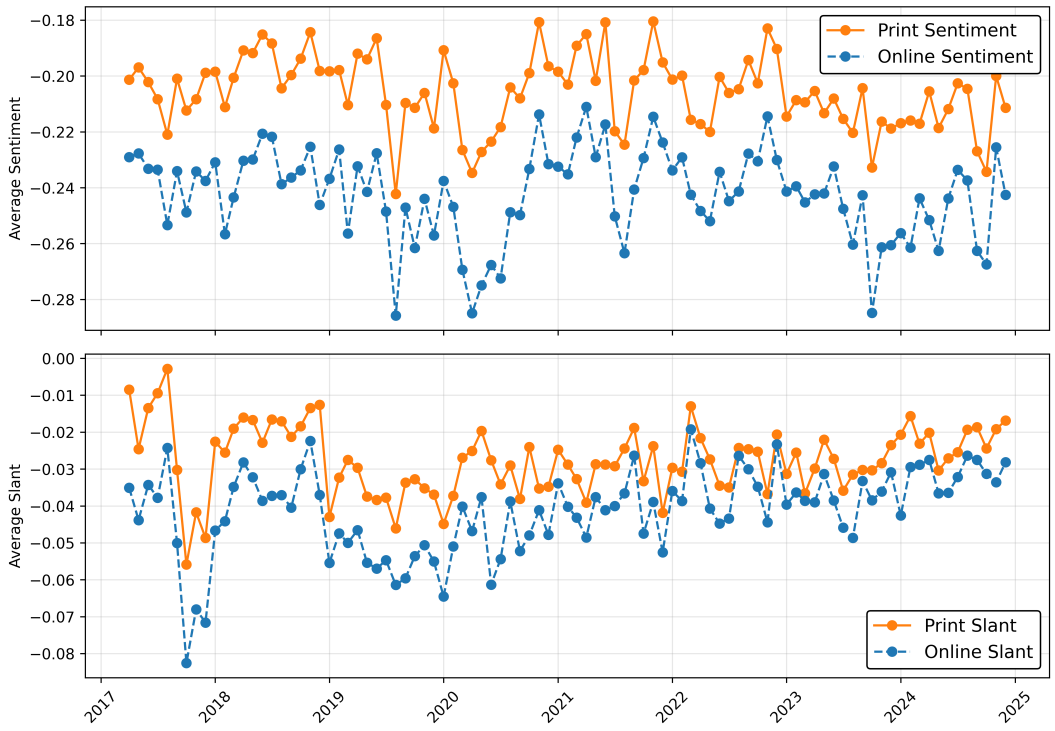


Figure 1: Monthly Average Sentiment and Slant

Monthly average sentiment and slant for articles with both online and print headlines from the 11 major newspapers.

5 Regression Analysis

This section presents reduced-form evidence on whether and how newspapers adjusted their online headlines relative to print in response to Facebook's July 19, 2022 announcement to reduce political content in user feeds. The goal of this analysis is to document whether newspapers altered their editorial behaviors about ideological slant of their online posts relative to their corresponding print headlines in response to this policy change. If so, it provides a quasi-experiment to examine whether media outlets strategically adjusted their social-media presentation after the announcement.

5.1 Empirical Setting and Measures

Each observation is a post n by outlet j on date t that is linked to a print headline. Let $s^n_{j,t}$ denote the slant (or sentiment) of the online/Facebook post and $p^n_{j,t}$ the slant (or sentiment) of its print counterpart. Our primary outcome is the absolute online–print divergence, $|s^n_{j,t} - p^n_{j,t}|$, which captures editorial adjustment. Although I have filter the posts with topic as news, not all news are political. Since the announcement specifically targeted political content, the intensity of the treatment should be different for different categories of news. Therefore, treatment intensity is proxied by $|p^n_{j,t}|$, the absolute print slant, interpreting larger values as more political content.

5.2 Binary Approximations

I first provide nonparametric/binned evidence by splitting stories along ex ante political intensity and comparing pre/post means of $|s^n_{j,t} - p^n_{j,t}|$. Figure 2 shows slant divergence falls after the announcement for posts that is more likely to be political; sentiment shows no comparable shift as shown in Figure 3. Similar results are shown for comparing pre/post means of $s^n_{j,t}$ using $p^n_{j,t}$ as control as well in Figure 11 and Figure 12.

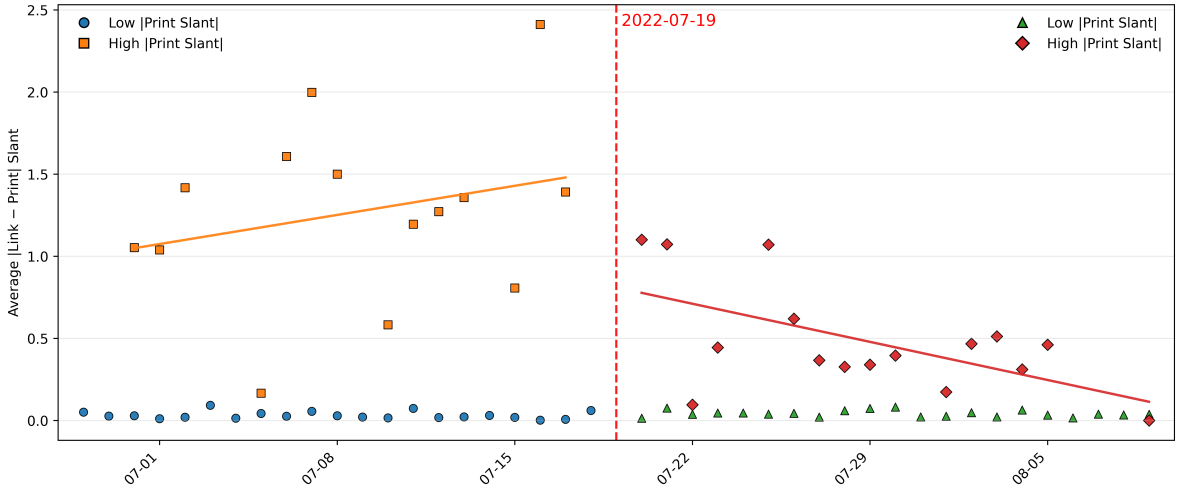


Figure 2: Two-bin split for Slant divergence $|s - p|$ (cut at $|p| = 0.75$).

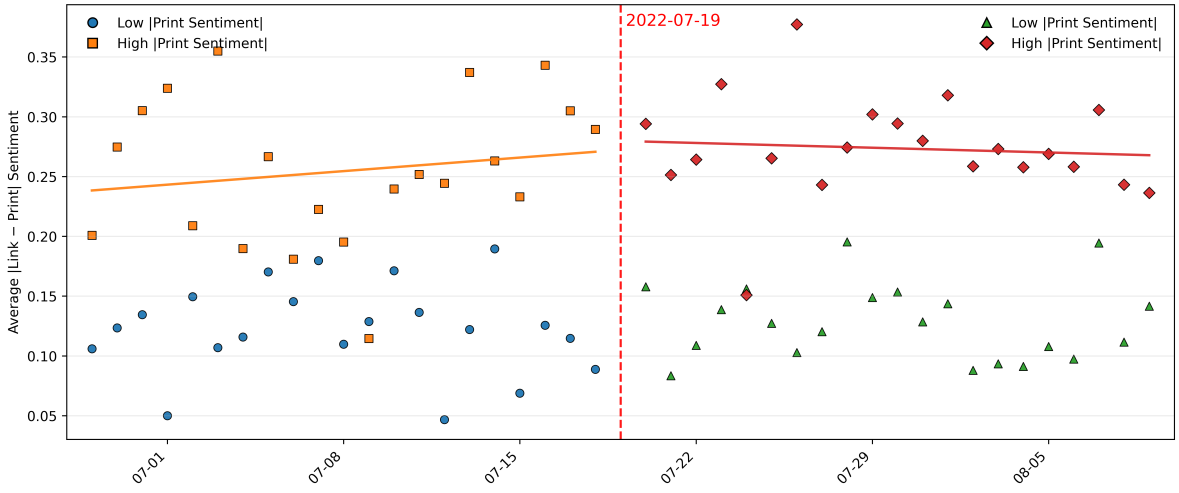


Figure 3: Two-bin split for Sentiment divergence $|s - p|$ (cut at $|p| = 0.35$).

5.3 Difference-in-Differences with Continuous Treatment

To quantify the effect of the announcement, I estimate a difference-in-differences regression with a continuous treatment intensity. The treatment intensity is proxied by the absolute slant of the print headline, $|p^n_{j,t}|$, which captures the ex ante likelihood that a post is perceived as political. The estimating equation is:

$$|s^n_{j,t} - p^n_{j,t}| = \gamma_j + \delta_t + b_0|p^n_{j,t}| + b_1\big(Post_t \times |p^n_{j,t}|\big) + \varepsilon^n_{j,t}, \tag{1}$$

where $Post_t$ is an indicator equal to one for days after July 19, 2022, and γ_j and δ_t are media and date fixed effects, respectively.

The interaction coefficient b_1 captures whether more political news ($|p^n_{j,t}|$ large) experienced larger changes in online-print divergence after the algorithm update. A negative b_1 indicates that politically extreme posts became more similar to their print counterparts after the update—that is, outlets reduced editorial work on social media in response to the platform's announcement.

Table 2: DID with Continuous Treatment around Facebook's Announcement

Duration (Weeks)	2	3	4	3
Measure	Slant	Slant	Slant	Sentiment
$\lvert p^n_{j,t}\rvert$	0.688 (0.054)	0.714 (0.052)	0.669 (0.050)	0.170 (0.021)
$Post_t \times \lvert p^n_{j,t}\rvert$	-0.284 (0.094)	-0.304 (0.088)	-0.206 (0.092)	0.020 (0.029)
n	2,055	3,202	4,224	3,202
R^2	0.356	0.372	0.348	0.090

Notes: OLS with outlet and date FE; standard errors clustered by outlet. The interaction b_1 captures the post-announcement change in online–print divergence for more political stories. Slant effects are negative and significant (largest at three weeks), while sentiment shows no detectable change.

Equation (1) is estimated by ordinary least squares separately for different time windows around the announcement date (two, three and four weeks). All specifications include media and date fixed effects to absorb unobserved heterogeneity in outlet reputation and day-specific news shocks. Standard errors are clustered at the media-outlet level.

As shown in Table 2, for both slant and sentiment, b_0 and $b_0 + b_1$ are always significantly positive, suggesting that news that has greater print slant and is thus more likely to be political, have greater gap between its online and print slant. Across all time windows, the interaction coefficient b_1 is negative and statistically significant for the slant regressions, particularly in the three-week window ($b_1 = -0.304$). This indicates that, immediately following the announcement, politically slanted articles reduced their deviation between online and print versions. In contrast, the effect on sentiment ($b_1 \approx 0$) is small and statistically insignificant, suggesting that newspapers specifically adjusted ideological framing rather than emotional tone. Since the announcement is specifically on political content, such response of slant change but sentiment seems to be reasonable.

To show such results do not rely on the specific variable section, I run similar regressions

$$Y_{j,t}^n = \gamma_j + \delta_t + b_0 X_{j,t}^n + b_1 \big(Post_t \times X_{j,t}^n\big) + \varepsilon_{j,t}^n, \tag{2}$$

for various combination of $Y_{j,t}^n$ and $X_{j,t}^n$ on both slant and sentiment. A richer set of outcomes confirms the pattern as in Table 3.

Table 3: DID with Alternative Outcomes and Intensities

$Y_{j,t}^n$	$\lvert s-p\rvert$		$\lvert s\rvert$		s		s	
$X_{j,t}^n$	$\lvert p\rvert$		$\lvert p\rvert$		$\lvert p\rvert$		p	
Measure	Slant	Sent.	Slant	Sent.	Slant	Sent.	Slant	Sent.
b_0	0.714	0.170	0.339	0.548	-0.312	-0.504	0.339	0.599
	(0.052)	(0.021)	(0.067)	(0.022)	(0.067)	(0.028)	(0.066)	(0.022)
b_1	-0.304	0.020	0.280	-0.026	-0.233	0.045	0.290	-0.030
	(0.088)	(0.029)	(0.102)	(0.031)	(0.104)	(0.039)	(0.101)	(0.031)
R^2 (n=3,202)	0.372	0.090	0.279	0.311	0.220	0.223	0.282	0.373

Notes: Model specification as (2) for three weeks before and after the announcement date, varying outcomes Y and intensity X. Significant impact on slant related variables post announcement; no systematic effect for sentiment.

5.4 Validity Check

I first test whether user demand shifted mechanically at the announcement by examining average views and engagement conditional on views. As long as a post is recommended and shown on the users screen, it is considered as a view count. Therefore, the number of views captures only the algorithms behavior. Once a post is recommended and viewed, a user can choose to engage or not through commenting or sharing. So the ratio of engagements over views purely reflect users' behavior. Later, I use this difference of view and engagement to disentangle algorithm and users' preference in Section 7. As shown in table 4, for four weeks before and after the announcement date, the average views fall significantly, indicating an algorithmic shift in distribution. Engagement conditional on views is statistically unchanged, implying the response is not driven by an immediate change in user demand but by editorial adjustments to anticipated algorithmic demotion.

Table 4: Views, Engagement per View, and Posting Activity Around the Announcement

Measure	Views	Engagements / Views	Posts (Daily)
Before	1.04×10^5 (6.22×10^3)	3.60×10^{-3} (1.10×10^{-4})	145.3 (7.27)
After − Before	-2.53×10^4 (7.01×10^3)	-2.75×10^{-4} (1.41×10^{-4})	−0.42 (10.37)

Notes: Entries are estimated from regressions of the form $y = b_0 + b_1 \cdot \text{Post}$, where Post is an indicator for the period after the announcement date using a four-week window on each side. The coefficient b_0 represents the average level before the announcement, and b_1 represents the change after the announcement relative to before. Views and engagement per view are computed at the post level, while the number of posts is aggregated at the daily level. Standard errors are reported in parentheses and are clustered at the post level (by date for post-level outcomes). For three weeks before and after the announcement date, views fall significantly, indicating an algorithmic shift in distribution. Engagement conditional on views is statistically unchanged, consistent with stable user preferences at the intensive margin. Number of post is also statistically unchanged.

I then check whether outlets adjusted posting frequency around the announcement. The last column of table 4 and figure 4 shows no contemporaneous change in the count of daily posts, suggesting the response operated on intensive editorial margins (slant) rather than quantity. Notice there is a weekly seasonal pattern of the number of posts. As a large proportion of news are related to government behaviors, which do not work during weekends, there are significantly more posts on weekdays than on weekends.

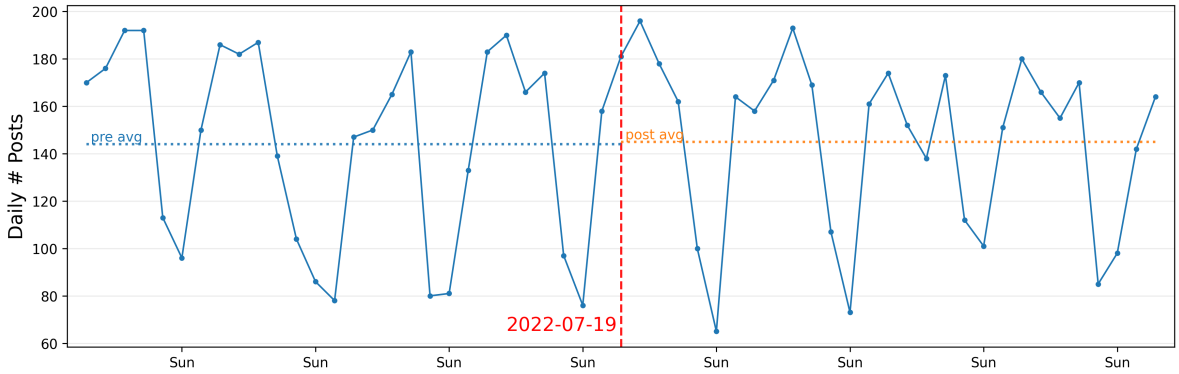


Figure 4: Daily number of posts for three weeks before and after the announcement date: no discernible change on the extensive margin.

I also conduct placebo test by randomly picking various dates that no announcement from Facebook was made, to run DID regression as equation 1. None of them has significant estimate on b_1.

These results are consistent with strategic supply-side adjustments by media outlets in anticipation of algorithmic demotion of political content. Because Facebook's update was expected to lower exposure for political news, outlets have less incentive to adjust the given slant from the print headlines. As they need to put in the same level of effort but will expect much less returns.

However, the above results do not show how exactly outlets choose the extra slant. It is also not able to quantify how the effect of users' preference and algorithms design affect the news supply of slant and sentiment change over time. To accomplish these, a more specific structural model is needed.

6 Model

I develop a dynamic structural model of online news markets in which three types of agents interact: (i) Unit measure of *users* who consume news recommended to them by the platform, (ii) One *algorithm* that mediates content delivery, and (iii) J *media outlets* that strategically choose the slant of their news posts. The model formalizes how recommendation algorithms

and evolving user preferences jointly shape the ideological distribution of news supply.

Time is discrete and indexed by $t \in \{0, 1, 2, \ldots\}$. Each outlet is endowed with multiple news to post per period, and decides how to add slants before publishing them. With the posts in the period published, the algorithms choose one to recommend to a user given the user's ideology. Once a user views a post, they need to decide whether to engage with the post or not. In the following subsections, I introduce each agent's decision making problem backwardly.

6.1 Users

In each period t, there is a continuum of users indexed by $i \in [0, 1]$. Each user i is characterized by an ideology $\mu_{i,t}$ that determines their taste for politically slanted content, where $\mu_{i,t}$ follows a distribution F_t.

If a post $\{n_{j,t}\}$ from media outlets $j \in \{1, \ldots, J\}$ is recommended to user i, denote $R^n_{ij,t} = 1$. Conditional on being recommended, the user decides whether to engage based on their utility. Denote the engagement decision as $I^n_{ij,t} \in \{0, 1\}$.

The utility of user i from engaging with post $n_{j,t}$ is:

$$U^n_{ij,t}(I^n_{ij,t}1 \mid R^n_{ij,t} = 1) = W^n_{ij,t} + \varepsilon^n_{ij,t} \tag{3}$$

where $\varepsilon^n_{ij,t} \sim \text{Logistic}(0, 1)$ and the mean utility for engagement is

$$W^n_{ij,t} = u^n_{j,t} + \xi_{j,t} + \alpha_t(s^n_{j,t} - \lambda_{j,t-1})^2 - \beta_t(s^n_{j,t} - \mu_{i,t})^2 \tag{4}$$

where $u^n_{j,t}$ and $s^n_{j,t}$ are the content value and slant of post $n_{j,t}$, $xi_{j,t}$ and $\lambda_{j,t-1}$ are the media fixed effect on current period capturing outlet quality or reach and expected slant from the media outlet from last period.

The term $-\beta_t(s^n_{j,t} - \mu_{i,t})^2$ captures the distaste for reading a news post whose slant $s^n_{j,t}$ deviates from the preferred slant $\mu_{i,t}$, as proposed in Mullainathan and Shleifer (2005)

and applied in Gentzkow and Shapiro (2010) and Martin and Yurukoglu (2017). The term $\alpha_t(s^n_{j,t} - \lambda_{j,t-1})^2$ captures the effect of salience from surprise or the effect of informativeness from trustworthiness.

Given the logistic error, the probability that a user engages after a recommendation is

$$\Pr(I^n_{ij,t} = 1 \mid R^n_{ij,t} = 1) = \frac{\exp(W^n_{ij,t})}{1 + \exp(W^n_{ij,t})} \tag{5}$$

It also gives the expected utility conditional on viewing as $\mathbb{E}[U^n_{ij,t} \mid R^n_{ij,t} = 1] = \log(1 + \exp(W^n_{ij,t}))$, which will be included in the algorithm's decision problem.

If $\alpha_t > \beta_t > 0$, users with modest ideology prefer slant that matches their ideology but may still engage more with more extreme posts because of the salience effect. In such case, there should be more polarized content recommended if the algorithm just maximize users' utility. On the other hand, if $\beta_t > \alpha_t > 0$, the utility is maximized when $s^n_{j,t} = (\beta_t \mu_{i,t} - \alpha_t \lambda_{j,t-1})/(\beta_t - \alpha_t)$, which is greater than $\mu_{i,t}$ when $\mu_{i,t} > \lambda_{j,t-1} > 0$ and is smaller than $\mu_{i,t}$ when $\mu_{i,t} < \lambda_{j,t-1} < 0$. In other words, the ideal slant is more extreme than the user's ideology, when the ideology is more extreme than the media outlet's expected slant.

Notice that there is no price in the users' utility function. Although several newspapers in the data have their articles behind the paywall, the main product to be consumed here is the post or the headline itself, instead of the article linked to it. In fact, Sundar et al. (2025) finds that around 75% of the Facebook posts with links to news were shared without clicking into the linked web pages. Such phenomenon is more serve for extreme and user-aligned political content. Therefore, I focus on the posts as the product, which does not require a payment to be consumed.

6.2 Algorithm

I abstractly model the algorithm as a decision maker that maximizes its utility. One can consider it as the algorithm designers have such objective and develop a programming to

implement it automatically. The platform's recommendation algorithm observes all posts and predicts expected engagement utility for each user. Its predicted utility for recommending post $n_{j,t}$ is:

$$\widehat{U}^n_{ij,t}(R^n_{ij,t}=1) = \widehat{W}^n_{ij,t} + \widehat{\varepsilon}^n_{ij,t} \tag{6}$$

where $\widehat{\varepsilon}^n_{ij,t} \sim$ i.i.d. EV1$(0,1)$ and mean utility for recommendation is

$$\widehat{W}^n_{ij,t} = \widehat{u}^n_{j,t} + \widehat{\xi}_{j,t} + \widehat{\alpha}_t(s^n_{j,t} - \lambda_{j,t-1})^2 - \widehat{\beta}_t(s^n_{j,t} - \mu_{i,t})^2 + \zeta_t \log(1+\exp(W^n_{ij,t})) \tag{7}$$

where $\widehat{u}^n_{j,t}$ and $\widehat{\xi}_{j,t}$ are the extra content value and the extra media fixed effect considered by the algorithm.

The key parameter ζ_t measures how much the algorithm internalizes user utility through engagements. The two terms $\widehat{\alpha}_t(s^n_{j,t}-\lambda_{j,t-1})^2$ and $\widehat{\beta}_t(s^n_{j,t}-\mu_{i,t})^2$ captures the amplification on slanted related terms from the algorithm. If one believes that the algorithms are designed to maximize users' engagement, all the hat parameters can be considered as the measurement errors during the training of the algorithms.

Considering non-news posts as outside options with utility equals to 0, the probability of recommendation is

$$\Pr(R^n_{ij,t}=1) = \frac{\exp(\widehat{W}^n_{ij,t})}{1+\sum_{m,h}\exp(\widehat{W}^m_{ih,t})} \tag{8}$$

Then the unconditional probability that a post is both recommended and engaged is

$$\Pr(I^n_{ij,t}=1, R^n_{ij,t}=1) = \frac{\exp(W^n_{ij,t})}{1+\exp(W^n_{ij,t})}\frac{\exp(\widehat{W}^n_{ij,t})}{1+\sum_{m,h}\exp(\widehat{W}^m_{ih,t})} \tag{9}$$

Aggregating over user types yields the population-level viewing and engagement shares:

$$\Pr(R^n_{j,t}=1) = \int_i \frac{\exp(\widehat{W}^n_{ij,t})}{1+\sum_{m,h}\exp(\widehat{W}^m_{ih,t})} dF(\mu_{i,t}) \tag{10}$$

$$\Pr(I^n_{j,t}=1, R^n_{j,t}=1) = \int_i \frac{\exp(W^n_{ij,t})}{1+\exp(W^n_{ij,t})}\frac{\exp(\widehat{W}^n_{ij,t})}{1+\sum_{m,h}\exp(\widehat{W}^m_{ih,t})} dF(\mu_{i,t}) \tag{11}$$

There are two channels of amplification on effects from slant. The more direct one is through $\widehat{\alpha}_t$ and $\widehat{\beta}_t$. The other one is through ζ_t. If users care about the slant when considering engagement decisions and the algorithm cares about the users' engagement, a greater ζ_t will make political content, in particular, slanted political content, more likely to be recommended than the others.

Here, I model users and the algorithm as short-lived. Since there are many updates and experiments on the algorithm, and the users group keeps changing throughout time, this is a reasonable design. It also allows me to focus on the media outlets' decision on the supply of slant, which is the main focus of the paper.

6.3 Media Outlets

Media outlets make decisions dynamically. In each period t, the media outlet is endowed with N news with each denoted by $n_{j,t} \in \{1, 2, ..., N_{j,t}\}$. Each news $n_{j,t}$ is characterized by the print slant $p^n_{j,t}$, content value $\mu^n_{j,t}$, and extra content value for the algorithm $\widehat{\mu}^n_{j,t}$. Next, I am going to describe the dynamic decision problem for a given media outlet j in period t, and drop the j, t subscripts for cleaner notations. Denote $\omega^n = \{n, \mu^n, \widehat{\mu}^n, p^n\}$ as all the properties of news n. The editors of the outlet decide whether to adjust the slant from the print version of post n by choosing $a^n \in \{-1, 0, 1\}$, which gives the slant as

$$s^n(a^n) = p^n + a^n \epsilon^n \tag{12}$$

where $\epsilon^n > 0$ is random with known distribution but realizes after the decision is made.

Once all news are edited and posted. The expected slant of the outlet $\lambda_{j,t}$ evolves as a moving average of slants from its past posts,

$$\lambda_t = (1 - \rho)\lambda_{t-1} + \rho \frac{1}{N} \sum_n s^n. \tag{13}$$

such that if an outlet posts more right slanted content this period, its expected slant will

shrift to the right in the next period.

Denote θ as the exogenous states for the current period and $\Omega = \{\omega^n\}_{n=1}^N$ as the properties of all news from the outlet in the period. The timing of the games is as in figure 5. The media outlet is endowed with the news articles and the expected slant λ_- at the beginning of the period t. Then it makes decisions given the properties of the news Ω and based on the exogenous state variables θ. The slants are realized after the decisions are made. At the same time, all other outlets are making decisions and publishing their posts as well. Once all posts are published, the algorithm picks one post to recommend to each user. Given the recommended post, the user decides whether to engage with it or not. With the realized slant of all the posts from the outlet, its expected slant is updated to λ and the game goes to the next period, where θ' is realized given the known transit probabilities $G(\theta'|\theta)$.

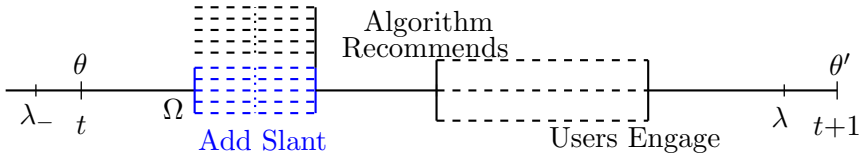


Figure 5: Timing of the Game

Denote $A = \{a^n\}_{n=1}^N$ as the decisions on all news from the outlet in the period. The value function of the outlet j can be written as

$$V_j(\Omega; \lambda_-, \theta) = \max_{A \in \{-1,0,1\}^N} \tau_j^V U^V(A) + \tau_j^E U^E(A) - c_j^A + \upsilon^A + \beta_j \mathbb{E}[V_j(\Omega'; \lambda, \theta')] \tag{14}$$

where $\upsilon^A \sim i.i.d.EV1\,(0,1)$.

For the flow utilities $U^V(A) = \sum_{n=1}^N \mathbb{E}[\Pr\,(R^n = 1|\, s^n(a^n), \omega^n, \lambda_-, \theta)]$ is the expected total probability of any post being recommended or being viewed. Similarly, $U^E(A) = \sum_{n=1}^N \mathbb{E}[\Pr\,(I^n = 1, R^n = 1|\, s^n(a^n), \omega^n, \lambda_-, \theta)]$ is the expected total probability of any post being engaged. Since only one post will be recommended, each event can be considered as independent and the total probability will just be the sum of them. c_A is the action specific cost, in the sense that it might be easier for some editors to come up with more right slanted

phrases than more left slanted phrases.

The level of τ_j^V, τ_j^E tells whether editors are adding extra slants to maximizing views or engagements, or just randomly doing so. The comparison between the costs of adding negative slant or positive slant can suggest the political leaning of the editors as well.

6.4 Oblivious Equilibrium

The above environment describes a dynamic game wherein media outlets add slants dynamically and strategically. In other words, an outlet's past choices affect that its current choice through their effect on the outlet's expected slant $\lambda_{j,t-1}$. In addition, other outlets' current and past choices affect a given outlet's choice, as the probability of each post being viewed or engaged will change accordingly. The literature pioneered by Pakes and McGuire (2001) often uses Markov Perfect Equilibrium (MPE) to analyze players' behavior in such dynamic games. However, in my setting with more than thirty media outlets and hundreds of news posts in each period, the state variable for the game becomes the slant of each post as well as the expected slant of each media outlet which is a vector of dimension easily more than a thousand. This is the curse of dimensionality, which is particularly a problem in dynamic settings with large state space, as discussed in Pakes and McGuire (2001).

In order to curb this problem, I use the oblivious equilibrium as developed by Weintraub et al. (2008). Additionally, I use the adapted version that accounts for time-varying state space as in Saeedi (2019). Under this equilibrium concept, it is assumed that the total utility of all posts are the state variables; that is, all histories of the game that lead to the same level of total utility of all posts must lead to the same strategies. This is in contrast with MPE, in which variations in the past individual states affect equilibrium strategies beyond their effect on average industry variables. This equilibrium concept is particularly applicable to markets that are comprised of many news suppliers. Note that with many smaller outlets, an approximate law of large numbers implies that the distributions of individual states are somewhat invariant, which would then mean that industry average becomes the relevant

state variable in the game. In my sample, there are 32 media outlets, each with dozens of posts in a period. Among the 3 millions Facebook posts, less than 10 posts have more than 5% of the views of that period, and most posts have less than 1% of view share; therefore, oblivious equilibrium can be a good approximation.

The application of the oblivious equilibrium has two key implications: first, the only relevant endogenous state variable is the expected slant of the given media outlet; second, each outlet is small and, as a result, an outlet's choice does not affect the state of the industry. In other words, outlets are oblivious to their knowledge of other outlets' state.

Given this discussion, I can proceed with the formal definition of the equilibrium. Let $\theta = \{\xi, \alpha, \beta; \zeta, \widehat{\xi}, \widehat{\alpha}, \widehat{\beta}; \{M_\mu\}_\mu, F\}$ be the exogenous states for the current state, where M_μ is the market denominator for the probability of being recommend to users with belief μ by $\Pr(R^n_\mu = 1) = \exp(\hat{W}^n_\mu)/M_\mu$. An oblivious equilibrium is defined as a set of policy functions, $A^*(\Omega; \lambda_-, \theta)$, properties of the news Ω, exogenous states θ, such that

- given the action and states, U^V and U^E are the outcome of users and algorithm's utility optimization;
- $A^*(\Omega; \lambda_-, \theta)$ is maximizing outlets' value function given all the state variables;
- total market denominator for each belief M_μ is consistent with individual media outlet's choices—that is, the sum of $\exp(\hat{W}^n_\mu)$ across outlets is equal to M_μ.

7 Identification and Estimation

This section describes the identification and estimation of the model parameters introduced above. Estimation proceeds in two stages. First, I estimate the demand-side parameters that govern user behavior $(\{\xi_{j,t}\}_j, \alpha_t, \beta_t)$ and algorithmic recommendation $(\{\widehat{\xi}_{j,t}\}_j, \widehat{\alpha}_t, \widehat{\beta}_t, \zeta_t)$ using post-level view and engagement data. Estimation is performed separately for each period t to allow for time variation in the platform and user environment. Second, I estimate the

supply-side parameters that govern media outlet behavior $(\tau_j^V, \tau_j^E, c_{a,j}, \beta_j)$ using observed editorial slant choices and the recovered conditional choice probabilities (CCPs).

7.1 Users and Algorithms (Demand Side)

For each post $n_{j,t}$ published by outlet j in period t, I observe the number of views $V_{j,t}^n$ and the number of engagements (comments plus shares) $E_{j,t}^n$. I also observe the post slant $s_{j,t}^n$ and the number of posts from outlet j in period t, $N_{j,t}$.

Before estimating the parameters, two state variables are estimated separately. First is the expected slant of outlet j is estimated by the equation

$$\lambda_{j,t} = (1-\rho)\lambda_{j,t-1} + \rho \frac{1}{N_{j,t}} \sum_n s_{j,t}^n, \tag{15}$$

where the initial $\lambda_{j,0}$ is the average slant of posts in several periods preceding the estimation window, and the decay parameter ρ is calibrated. Second is the distribution of the user's ideology. I assume three representative types: $\mu_{i,t} \in \{\mu_{l,t}, 0, \mu_{r,t}\}$ with probabilities $(p_l, 1 - p_l - p_r, p_r)$ where $\mu_{l,t} < 0$ (left), $\mu_{r,t} > 0$ (right), and ideology may slowly drift over time. A positive shift in $(\mu_{r,t} - \mu_{l,t})$ indicates increasing ideological polarization of the user base.

User ideological types are estimated directly from engagement-weighted averages of post slants:

$$\mu_{l,t} = \frac{\sum_{s_n<0} s_n E_n}{\sum_{s_n<0} E_n}, \quad p_{l,t} = \frac{\sum_{s_n<0} E_n}{\sum_n E_n}, \tag{16}$$

$$\mu_{r,t} = \frac{\sum_{s_n>0} s_n E_n}{\sum_{s_n>0} E_n}, \quad p_{r,t} = \frac{\sum_{s_n>0} E_n}{\sum_n E_n}. \tag{17}$$

These moments summarize both the average ideological positions and the relative mass of left- and right-leaning users as revealed through engagement behavior. Consistent with Guess (2021), most Americans consume a largely moderate and overlapping mix of online news, while a relatively small set of highly partisan users generates a disproportionate share of

interactions with ideologically extreme content. The three-type representation is designed to capture this empirical pattern in a parsimonious way—allowing a small but influential group of polarized users to coexist with a large centrist majority. I have experimented with richer specifications for the distribution of user ideology, and the estimated model outcomes are highly robust to these alternative choices.

With the above two states estimated, for each period t, the demand-side parameter vector is $\theta_t^D = \{\{\xi_{j,t}\}_j, \alpha_t, \beta_t; \{\widehat{\xi}_{j,t}\}_j, \widehat{\alpha}_t, \widehat{\beta}_t, \zeta_t\}$. The key behavioral parameters are $\beta_t, \widehat{\beta}_t$ (ideological sensitivity), $\alpha_t, \widehat{\alpha}_t$ (salient or informativeness effect), and ζ_t (algorithmic alignment with user utility).

I can rewrite the mean utility for engagement and recommendation as

$$W_{ij,t}^n = \delta_{j,t}^n - \beta_t (s_{j,t}^n - \mu_{i,t})^2 \tag{18}$$

$$\widehat{W}_{ij,t}^n = \widehat{\delta}_{j,t}^n - \widehat{\beta}_t (s_{j,t}^n - \mu_{i,t})^2 + \zeta_t \log(1 + \exp(W_{ij,t}^n)) \tag{19}$$

where

$$\delta_{j,t}^n = u_{j,t}^n + \xi_{j,t} + \alpha_t (s_{j,t}^n - \lambda_{j,t-1})^2 \tag{20}$$

$$\widehat{\delta}_{j,t}^n = \widehat{u}_{j,t}^n + \widehat{\xi}_{j,t} + \widehat{\alpha}_t (s_{j,t}^n - \lambda_{j,t-1})^2 \tag{21}$$

Identification relies on the variation in engagement and view shares across posts with different slants. The model predicts that both viewing probabilities $\Pr(R_{j,t}^n = 1)$ and engagement probabilities $\Pr(I_{j,t}^n = 1, R_{j,t}^n = 1)$ are increasing in $\delta_{j,t}^n$ and $\widehat{\delta}_{j,t}^n$ by Equation 10. The estimation procedure extends the two-step share inversion method of Berry et al. (1995) to a setting with algorithmic decisions and two coupled markets: one for views (recommendation) and one for engagement. In the inner loop, I apply a contraction mapping to match observed view and engagement shares, recovering the mean utilities $\delta_{j,t}^n$ and $\widehat{\delta}_{j,t}^n$ for each post given candidate $(\beta_t, \widehat{\beta}_t, \zeta_t)$. The outer loop then estimates the structural parameters by regressing the recovered utilities on each independent variables. Unlike standard BLP applications,

there are no prices or equilibrium constraints, and the recommendation utility $\widehat{W}$ explicitly depends on the expected engagement utility W through the term $\zeta_t \log(1 + \exp(W))$. Since the sum of the observed share is far less than 1, an inversion can be used to estimate $\delta^n_{j,t}$ and $\widehat{\delta}^n_{j,t}$ of each post, as in Berry et al. (1995). To my knowledge, this approach—linking share inversion to a two layers demand system—has not been implemented in previous work.

With the number of views $V^n_{j,t}$ and the number of engagements (like comments or shares) $E^n_{j,t}$ for each post, the demand-side parameters are estimated in two loops:

Inner loop (inversion of shares): Given candidate values $(\beta_t, \widehat{\beta}_t, \zeta_t)$, I invert the predicted shares to recover $\delta^n_{j,t}$ and $\widehat{\delta}^n_{j,t}$.

(i) Initialize $\delta^{n,(0)}_{j,t}$ and $\widehat{\delta}^{n,(0)}_{j,t}$.

(ii) Compute predicted shares under the model based on equation 10:

$$\widehat{S}^{n,E}_{j,t} = \Pr(I^n_{j,t} = 1, R^n_{j,t} = 1 \mid \beta_t, \widehat{\beta}_t, \zeta_t, \delta^{n,(0)}_{j,t}, \widehat{\delta}^{n,(0)}_{j,t})$$
$$\widehat{S}^{n,V}_{j,t} = \Pr(R^n_{j,t} = 1 \mid \beta_t, \widehat{\beta}_t, \zeta_t, \delta^{n,(0)}_{j,t}, \widehat{\delta}^{n,(0)}_{j,t})$$

(iii) Update using share inversion:

$$\delta^{n,(d+1)}_{j,t} = \delta^{n,(d)}_{j,t} + \log S^{n,E}_{j,t} - \log \widehat{S}^{n,E}_{j,t}$$
$$\widehat{\delta}^{n,(d+1)}_{j,t} = \widehat{\delta}^{n,(d)}_{j,t} + \log S^{n,V}_{j,t} - \log \widehat{S}^{n,V}_{j,t}$$

where $S^{n,E}_{j,t} = E^n_{j,t} / \sum_k \sum_m V^m_{k,t}$ and $S^{n,V}_{j,t} = V^n_{j,t} / \sum_k \sum_m V^m_{k,t}$ are the observed share of engagement and view. And $\sum_k \sum_m V^m_{k,t}$ represents the number of views from all posts in period t, including non-news posts. See a proof that the above operation is a contraction mapping in Appendix C.

(iv) Repeat (ii) and (iii) until convergence.

Outer loop (structural regression): Given the recovered $\delta^n_{j,t}$ and $\widehat{\delta}^n_{j,t}$, estimate:

$$\delta^n_{j,t} = \xi_{j,t} + \alpha_t(s^n_{j,t} - \lambda_{j,t-1})^2 + u^n_{j,t}, \tag{22}$$

$$\widehat{\delta}^n_{j,t} = \widehat{\xi}_{j,t} + \widehat{\alpha}_t(s^n_{j,t} - \lambda_{j,t-1})^2 + \widehat{u}^n_{j,t}. \tag{23}$$

Then choose $(\beta_t, \widehat{\beta}_t, \zeta_t)$ to minimize the total sum of squared residuals:

$$\min_{\beta_t, \widehat{\beta}_t, \zeta_t} \sum_j \sum_n ((u^n_{j,t})^2 + (\hat{u}^n_{j,t})^2).$$

This procedure identifies how user ideology and algorithmic incentives jointly determine which posts are shown and engaged with. I allow users with different ideologies to be recommended with posts with different probabilities, i.e., personalized recommendation, which is the core concept in algorithm design. Since prices are absent and slant is plausibly exogenous at the post level, I use a minimum distance estimator rather than IV-GMM. I also include other terms, such as the length of the post, slant $s^n_{j,t}$ as a linear term, etc and none of them affects the estimation on the above parameters.

7.2 Media Outlets (Supply Side)

On the supply side, each media outlet chooses slant adjustments $a^n_{j,t} \in \{-1, 0, 1\}$ given state variables $\theta_{j,t}$ post-level characteristics $\omega^n_{j,t}$, which will then result in different $s^n_{j,t}$ following Equation 12. The outlet-level parameters to estimate are $\tau^V_j, \tau^E_j, c_{-1,j}, c_{+1,j}, \beta_j$ where τ^V_j measures the value of view traffic, τ^V_j the value of engagement traffic, and β_j the discount factor. The cost of no slant change is normalized to 0 such that $c_{-1,j}$ and $c_{+1,j}$ are the cost of adding slant to the left and to the right relative to no change on slant.

Notice that the action space grows exponentially with the number of posts in the period. To handle this issue, I assume that the decision on each post is made sequentially and that the editors know how many news they need to publish at the beginning of the period, which

is set to be daily in the estimation. Then the value function of media outlet j becomes

$$V_j\Big(\omega^n, \hat{\lambda}, \lambda_-, N, \theta\Big) = \max_{a^n \in \{-1,0,1\}} \tau_j^V U^V(a^n) + \tau_j^E U^E(a^n) - c_{a^n} + \upsilon_a^n + \mathbb{E}\Big[V_j\Big(\omega^{n+1}, \hat{\lambda} + \frac{s^n(a^n)}{N}, \lambda_-, N, \theta\Big)\Big]$$

where $\upsilon_a^n \sim i.i.d.EV1\,(0,1)$ and $U^V(a^n), U^E(a^n)$ are the expected probabilities of view and engagement, which depends on $\omega^n, \lambda_-, \theta$ and the distribution of s^n conditional on a^n. $\hat{\lambda}$ is used for updating the expected slant at the end of the period. In the special case of the last post in the period, the value function becomes

$$V_j\Big(\omega^N, \hat{\lambda}, \lambda_-, N, \theta\Big) = \max_{a^N \in \{-1,0,1\}} \tau_j^V U^V(a^N) + \tau_j^E U^E(a^N) - c_{a^N} + \upsilon_a^N + \beta_j \mathbb{E}\Big[V_j\Big(\omega^{1'}, 0, \lambda, N', \theta'\Big)\Big]$$

where $\lambda = (1-\rho)\lambda_- + \rho(\hat{\lambda} + \frac{s^n(a^n)}{N})$ is how the expected slant being updated and $\omega^{1'}$ is the property of the first news in next period. Once the game enters the next period, a new N' is drawn from the empirical distribution and the next period state θ' is drawn from the known transit probabilities $G(\theta'|\theta)$.

With the new specification, identification exploits observed variation in slant choices conditional on state variables and recovered CCPs. I use a two-step conditional choice probability (CCP) estimator following Hotz and Miller (1993).

Let Θ collect the (observed) state $\omega^n, \hat{\lambda}, \lambda_-, N, \theta$. The action-specific utility function can be written as

$$Q_j(a \mid \Theta) \;=\; \tau_j^V\, U^V(a \mid \Theta) \;+\; \tau_j^E\, U^E(a \mid \Theta) \;-\; c_{-1,j}\, \mathbf{1}\{a=-1\} \;-\; c_{+1,j}\, \mathbf{1}\{a=+1\},$$

where $U^V(\cdot)$ and $U^E(\cdot)$ are the model-implied probabilities (from the estimated demand/algorithm side) that the post is viewed and engaged, respectively. Private action shocks are i.i.d. EV1,

so the conditional choice probabilities (CCPs) obey

$$P_j(a \mid \Theta) \;=\; \frac{\exp\big(Q_j(a \mid \Theta) + \beta_j(\Theta)\,\mathbb{E}[V_j' \mid a, \Theta]\big)}{\sum_{a' \in \{-1,0,1\}} \exp\big(Q_j(a' \mid \Theta) + \beta_j(\Theta)\,\mathbb{E}[V_j' \mid a', \Theta]\big)}.$$

Let $a_0 = 0$ denote the reference action. The continuation value after choosing a can be written as a discounted sum of future flow components that are linear in the same economic primitives that enter current payoffs:

$$\mathbb{E}[V_j' \mid a, \Theta] \;=\; \tau_j^V\,\phi(a \mid \Theta) \;+\; \tau_j^E\,\psi(a \mid \Theta) \;-\; c_{-1,j}\,\kappa_{-1}(a \mid \Theta) \;-\; c_{+1,j}\,\kappa_{+1}(a \mid \Theta),$$

where, for a given discount factor β_j,

$$\begin{aligned}
\phi(a \mid \Theta) &:= \mathbb{E}\left[\sum_{\ell=1}^{\infty} \beta_j^{\ell-1} \sum_n U_{n,t+\ell}^V \,\middle|\, a, \Theta\right] \\
\psi(a \mid \Theta) &:= \mathbb{E}\left[\sum_{\ell=1}^{\infty} \beta_j^{\ell-1} \sum_n U_{n,t+\ell}^E \,\middle|\, a, \Theta\right], \\
\kappa_i(a \mid \Theta) &:= \mathbb{E}\left[\sum_{\ell=1}^{\infty} \beta_j^{\ell-1} \sum_n \mathbf{1}\{a_{j,t+\ell}^n = i\} \,\middle|\, a, \Theta\right], \quad i \in \{-1, +1\}
\end{aligned}$$

The technique has been mentioned in Arcidiacono et al. (2013), where they use polynomials to construct the sieve is that the approximation is linear in parameters. These objects are known functions once I (i) fix β_j, and (ii) simulate forward using the estimated demand-side primitives and the observed state process, as in Hotz et al. (1994).

By the logit structure, for any realized Θ,

$$\log \frac{P_j(a \mid \Theta)}{P_j(a_0 \mid \Theta)} = \big[Q_j(a \mid \Theta) - Q_j(a_0 \mid \Theta)\big] + \beta_j\Big(\mathbb{E}[V_j' \mid a, \Theta] - \mathbb{E}[V_j' \mid a_0, \Theta]\Big).$$

Define the known regressor (all functions of observables and simulated paths, conditional on β_j):

$$\Phi(a \mid a_0, \Theta) := \big[U^V(a \mid \Theta) - U^V(a_0 \mid \Theta)\big] + \big[\phi(a \mid \Theta) - \phi(a_0 \mid \Theta)\big],$$

$$\Psi(a \mid a_0, \Theta) := \big[U^E(a \mid \Theta) - U^E(a_0 \mid \Theta)\big] + \big[\psi(a \mid \Theta) - \psi(a_0 \mid \Theta)\big],$$

$$K_i(a \mid a_0, \Theta) := \big[\mathbf{1}\{a = i\} + \kappa_i(a \mid \Theta) - \kappa_i(a_0 \mid \Theta)\big], \quad i \in \{-1, +1\}$$

Then the estimating equation becomes the linear regression

$$\begin{aligned}\log \frac{P_j(a \mid \Theta)}{P_j(a_0 \mid \Theta)} =& \tau_j^V \, \Phi(a \mid a_0, \Theta) + \tau_j^E \, \Psi(a \mid a_0, \Theta) \\ & - c_{-1,j} \, K_{-1}(a \mid a_0, \Theta) - c_{+1,j} \, K_{+1}(a \mid a_0, \Theta) + \varsigma_j(a \mid \Theta) \end{aligned} \tag{24}$$

$\widehat{\beta}_j$ is chosen to minimize the residual sum of squares from the regression.

Although this two-step conditional choice probability (CCP) approach has been developed and widely used in the structural econometrics literature since Hotz and Miller (1993), its application to media economics—and particularly to modeling editorial decisions in response to algorithmic incentives—is new. Traditional uses of the CCP method focus on dynamic problems in labor supply, industrial organization, or entry–exit games (e.g., Aguirregabiria and Mira, 2007; Bajari et al., 2007; Arcidiacono and Miller, 2011). In contrast, applying it to the dynamic choice of news slant and tone introduces a novel context where firms (media outlets) interact with both algorithmic platforms and heterogeneous audiences in real time. This paper is, to my knowledge, the first to use CCP-based estimation to recover the dynamic incentives of media outlets in a high-dimensional digital environment, allowing for the joint identification of different engagement motives.

8 Empirical Results

This section presents the empirical results from both the demand (user and algorithm) and supply (media outlet) sides of the model. The estimation combines revealed-preference data on engagement and visibility from 1.5 million Facebook posts with matched print and online headlines from eleven major newspapers. The results quantify: (i) how users respond to

ideological features of news content; (ii) how algorithms amplify those preferences; and (iii) how media outlets dynamically adjust their headline slant to maximize performance under algorithmic feedback.

8.1 Algorithmic Amplification and User Preferences

Figure 6 visualize the estimated coefficients governing the degree to which Facebook's recommendation algorithm internalizes user utility. Before 2022, algorithmic weights on user engagement were positive, with ζ_t exceeding 2 for political content—implying that the platform prioritized posts predicted to yield high engagement, even beyond users' intrinsic preferences. After Facebook began rolling out its policy to "reduce political content" on August 31, 2021, the estimated ζ_t starts to decline, indicating a gradual reduction in the algorithm's reliance on engagement signals for politically charged content. By July 19, 2022, when the policy was fully implemented, ζ_t reaches its lowest level. This pattern is consistent with the platform's official statement that it would "place less emphasis on shares and comments for political content." Notably, the estimates suggest that the algorithm may even assign negative weight to engagement for such content in the later period, implying an active demotion of posts predicted to generate high engagement.

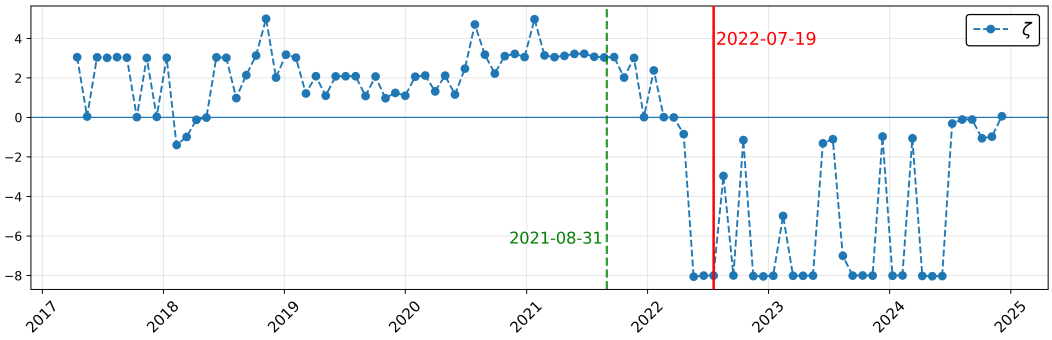


Figure 6: Evolution of algorithmic emphasis ζ_t (30-day intervals). Facebook began rolling out the reduction in political content on August 31, 2021, marking the onset of a decline in algorithmic weighting of politically charged content. The policy was fully implemented on July 19, 2022, at which point ζ_t reaches its lowest level, indicating a substantial reduction in the algorithm's emphasis on such content.

Table 5 reports the estimated demand-side coefficients prior to the algorithmic change. The parameters α_t and $\hat{\alpha}_t$ capture the salience effect from surprise or the trustworthiness from the unexpected slant by the users and the algorithm; the parameters β_t and $\hat{\beta}_t$ capture the how much users prefer like-minded news. ζ_t represent the corresponding algorithmic weights on users engagements. The results show that the algorithm amplifies both dimensions of user preference through two channels: the estimated ζ_t is averaged around 2, implying that recommendation design roughly doubles the effective incentive toward like-minded and salient content through the extra weights on engagements. On the other hand, $\hat{\alpha}_t$ and $\hat{\beta}_t$ are almost at the same level with α_t and $\hat{\beta}_t$, suggesting that the algorithm has extra consideration on the two slant related terms. In total, the algorithm amplification on slant related effect is around three times as much as the users original preferences. For the supply-side estimation, I use the sum of comments and shares as the engagement measure, since these actions are typically more costly and reflect stronger user intent than simple reactions. Moreover, as shown in Aldous et al. (2025), fewer than 1% of combined engagements on social media (reactions, shares, comments) come from the same user, suggesting that these metrics capture aggregate audience responses rather than repeated actions by a small set of users.

Table 5: Demand and Algorithm Parameters Before the Algorithmic Change

Engagements	Comments+Shares	Comments	Reactions
α	0.229 (0.027)	0.229 (0.030)	0.118 (0.031)
β	0.329 (0.034)	0.347 (0.037)	0.179 (0.039)
$\hat{\alpha}$	0.190 (0.023)	0.193 (0.029)	0.162 (0.032)
$\hat{\beta}$	0.244 (0.028)	0.261 (0.034)	0.222 (0.036)
ζ	2.229 (0.177)	2.166 (0.305)	1.255 (0.206)

Notes: Parameters α and β capture user-level preferences for surprise and ideological alignment; $\hat{\alpha}$ and $\hat{\beta}$ capture the corresponding algorithmic preferences; ζ measures how much the algorithm values user utility relative to engagement prediction. Standard errors in parentheses.

With the estimated model, I simulate the composition of slanted posts that are recom-

mended to (viewed by) and engaged with by users. Figure 7 reports the percentage change in the share of liberal-leaning and conservative-leaning posts at the viewed and engaged stages relative to their baseline supply levels.

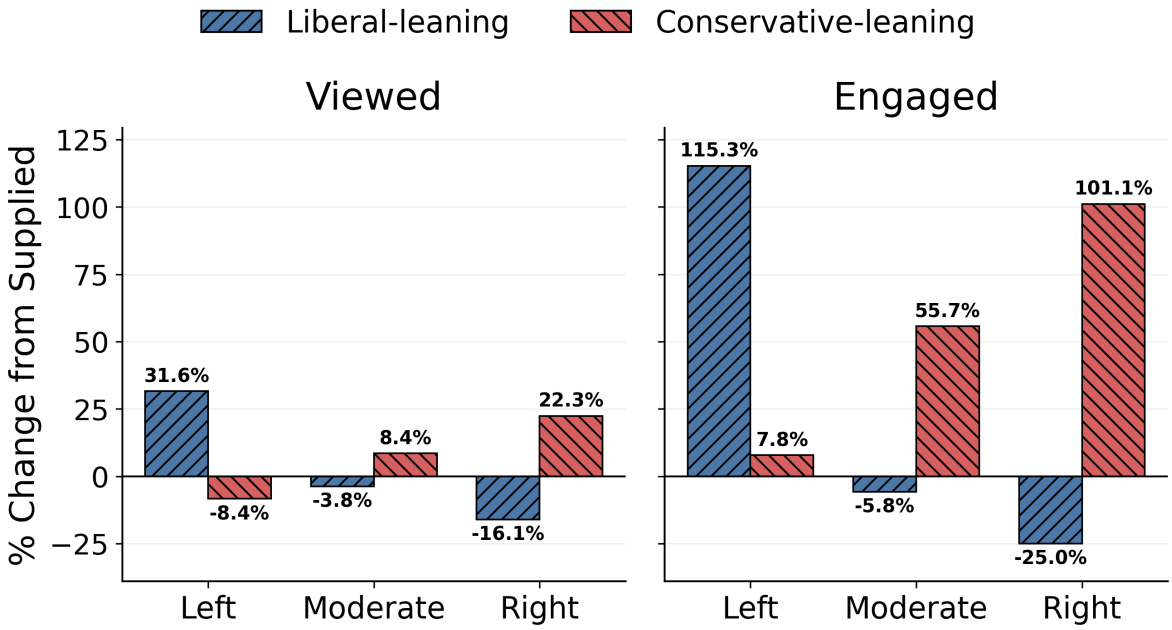


Figure 7: Percentage change in the share of ideologically slanted posts relative to supply. The figure reports the change in the share of liberal-leaning (negative slant) and conservative-leaning (positive slant) posts at the viewed and engaged stages relative to their baseline supply shares. The shares are simulated using the estimated model, with engagement measured by Comments + Shares. Users are grouped by ideology into Left, Moderate, and Right. Posts are classified as liberal-leaning, conservative-leaning, or neutral based on their estimated slant. Positive values indicate that a given type of content is overrepresented relative to supply, while negative values indicate under-representation. The figure highlights how algorithmic curation and user engagement behavior jointly distort the composition of content away from the underlying supply.

Figure 7 shows that ideological asymmetries arise both in exposure and in engagement behavior. Although news influencers, who are mostly conservative-leaning, are included in the dataset, most large media outlets are liberal-leaning and post more frequently. As a result, the supply of posts is tilted toward liberal-leaning slant. Comparing the supplied and viewed panels, the model implies that users are more likely to be exposed to ideologically aligned posts, consistent with the estimated results. In the engaged panel, all user groups are more likely to engage with conservative-leaning posts relative to their viewed (recommended)

share. In contrast, only left-leaning users exhibit a higher engagement rate for liberal-leaning posts, while moderate and right-leaning users do not display a comparable tendency. This pattern may reflect the role of unexpected slant: since most large media outlets are liberal-leaning, their conservative-leaning posts may attract more attention and generate higher engagement.

Overall, these results confirm two key insights. First, readers exhibit strong preferences for both like-minded news and unexpected slant—consistent with selective exposure and curiosity for ideological deviation. Second, Facebook's algorithm prior to mid-2022 substantially amplified these preferences, especially for politically slanted content.

8.2 Media Outlets' Objective and Slant Choice

Turning to the supply side, the estimated parameters $(\tau_j^V, \tau_j^E, c_{-1,j}, c_{+1,j}, \beta_j)$ reveal how outlets value views and engagements relative to the cost of headline adjustment. Figure 8 plots the estimated τ_j^V and τ_j^E coefficients by outlet, scaled by their follower counts. Since views and engagements enter linearly in my model, it implicitly assumes that outlets care about the the first view or engagement the same as when they have thousands of them. However, one should expect those with fewer views or engagements have much higher marginal value on the metrics. In fact, the raw estimation shows that smaller regional newspapers have much higher τ_j^V and τ_j^E than bigger national newspapers. Therefore, I report the estimation scaled by their follower counts on Facebook, which bring the numbers to the same magnitude. However, this does not change how significant the estimates are.

The estimates indicate that all outlets prioritize maximizing views but not engagements, suggesting that algorithms' visibility mechanisms dominate editorial incentives. Even though the algorithm takes users' preference into consideration when it recommends posts, as the decision on which post is viewed by the user is solely made by the algorithm and media outlets only care about views, the algorithm becomes the most critical component in shaping online news slant.

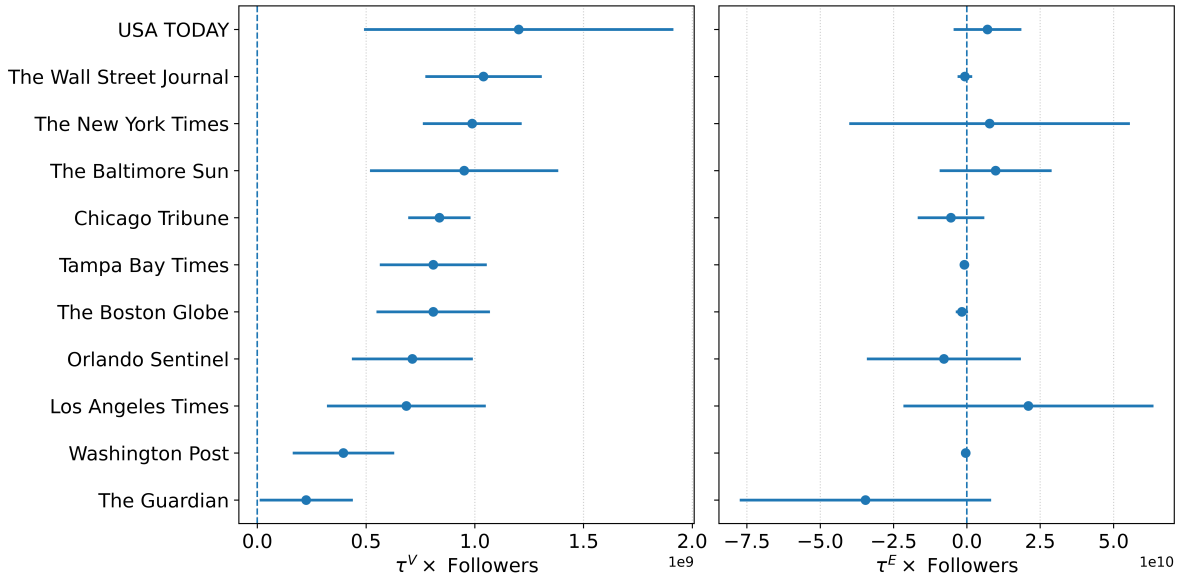


Figure 8: Estimated media-outlet weights on views and engagement (τ_j^V, τ_j^E), scaled by social-media follower counts. All outlets place sigficigantly positive weight on view traffic but not engagement metrics.

Figure 9 reports the estimated costs of left- and right-leaning headline adjustments, $c_{-1,j}$ and $c_{+1,j}$. These costs capture editorial frictions on adding extra slant to the existed print headlines. When $c_{+1,j} > c_{-1,j}$, it is more costly for the outlet to add right-leaning slant, which suggests a more liberal leaning editorial preference. And vice versa. The estimated cost asymmetries are modest on average, except Los Angeles Times exhibits a strong liberal leaning preference as its $c_{+1,j} > c_{-1,j}$ by a big margin.

The combination of demand and supply estimates provides a coherent quantitative picture of how social-media algorithms shape the production of news. Before 2022, Facebook's algorithm effectively tripled the marginal return to slanted content relative to intrinsic user demand. Media outlets responded by shifting the slant of online headlines toward more slanted language, consistent with their estimated τ_j^V values emphasizing view-maximization. Following Facebook's policy change to downweight political content, both algorithmic amplification (ζ_t) and outlet-level slant adjustments declined measurably, confirming that algorithm design can influence not only what users see but also what media outlets produce.

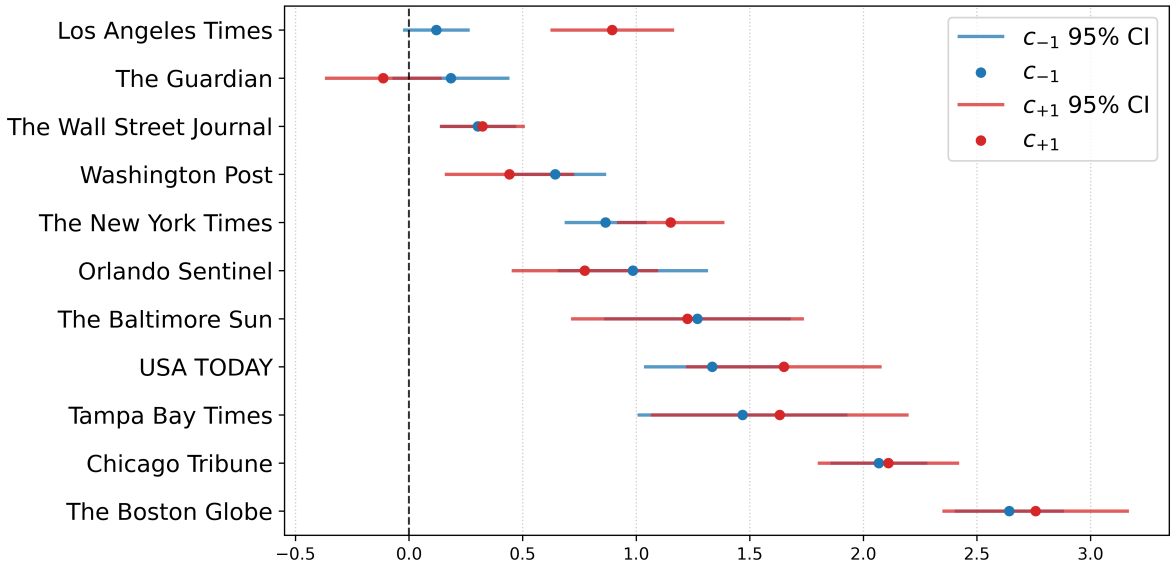


Figure 9: Estimated costs of left- and right-slant adjustments ($c_{-1,j}$ and $c_{+1,j}$). Costs capture reputational and editorial frictions associated with slant changes. Right-leaning deviations are on average more costly, reflecting both outlet branding and reader expectations.

9 Counterfactual Analysis

This section uses the estimated structural model to quantify how much of the observed ideological slant on social media can be attributed to algorithmic amplification. The analysis compares the original equilibrium to several counterfactual scenarios in which the recommendation algorithm's sensitivity to ideological and engagement signals is modified. All counterfactuals are simulated holding the estimated demand-side and supply-side parameters fixed, and recomputing equilibrium outcomes for post-level viewing, engagement, and editorial slant choices.

9.1 Counterfactual Design

The observed case represents the upper bound on algorithmically induced polarization. In this scenario, the algorithm fully amplifies user-level ideological preferences and engagement responses through positive values of $\hat{\alpha}_t$, $\hat{\beta}_t$, and ζ_t. This is the environment estimated from the data, corresponding to the original equilibrium during 2017–2022 when Facebook's feed

emphasized engagement metrics.

To isolate the mechanisms, I consider a sequence of counterfactual experiments that sequentially remove different components of algorithmic amplification:

- Scenario 1: No Extra Effect on Slant. The algorithm no longer rewards ideological alignment beyond users' own preferences. Formally, I set $\hat{\alpha}_t = \hat{\beta}_t = 0$, keeping ζ_t fixed. This removes algorithmic reinforcement of slant but maintains engagement weighting.

- Scenario 2: Ignore Engagement. The algorithm ceases to prioritize predicted engagement in its recommendation function, i.e., $\zeta_t = 0$. It still perceives ideological distances, but places no additional value on posts with higher expected engagement. This experiment identifies the importance of engagement-based amplification separate from ideological bias.

- Scenario 3: No Algorithm Effect. A fully neutral algorithm that neither rewards ideological similarity nor engagement intensity: $\hat{\alpha}_t = \hat{\beta}_t = \zeta_t = 0$. The decision is only based on the costs in this case.

For reference, I also compute a hypothetical lower bound in which all parameters related to ideological distance are shut down—both user and algorithmic components: $\alpha_t = \beta_t = \hat{\alpha}_t = \hat{\beta}_t = 0$ as the Benchmark Scenario. This represents the most extra slant that could ever be removed, since it eliminates ideological differentiation altogether. In practice, total slant cannot fall to zero even under this benchmark because the logit form of outlet's value function (υ_a^n) allow outlets to retain some dispersion in equilibrium slant purely due to randomness and unobserved heterogeneity.

The model simulates the results below by iterating the algorithmic recommendation and outlet choice functions under each parameter restriction. Outlets adjust their post slant choices through the estimated policy functions $P_j(a \mid \Theta)$, and the resulting equilibrium slant distribution is compared to the baseline.

9.2 Effects on Slant: Supply vs Exposure

For each counterfactual scenario, I compute two distinct measures of polarization. The first is the *supplied extra slant*, defined as the average difference between the slant of online posts and their corresponding print headlines. The second is the *viewed extra slant*, which measures the same difference but weighted by the equilibrium probability that each post is recommended and viewed by users.

Formally, the supplied extra slant captures how much media outlets adjust their content, while the viewed extra slant reflects the distribution of information that users are actually exposed to after algorithmic recommendation. Comparing the two allows me to separate changes in editorial behavior from changes in exposure induced by the platform.

Table 6: Counterfactual Reductions in Slant: Supplied vs Viewed

Scenario	Supplied Reduction (%)	Viewed Reduction (%)
Scenario 1	89.52%	90.24%
Scenario 2	16.82%	9.24%
Scenario 3	100.00%	99.86%

Notes: Each scenario recomputes equilibrium slant, visibility, and engagement given the parameter restrictions described above. Supplied reduction is based on the average extra slant across posts, while viewed reduction weights each post by its equilibrium probability of being recommended. All reductions are normalized relative to the benchmark scenario, which provides the maximum achievable reduction in slant.

Table 6 reports the percentage reduction in extra slant under each scenario for both measures, normalized relative to the benchmark case as described above. Several patterns emerge from the comparison between supplied and viewed slant.

First, removing only the algorithm's ideological weighting $(\hat{\alpha}_t, \hat{\beta}_t)$ in Scenario 1 yields a major reduction in extra online slant, indicating that only a small part of the bias originates from users' own preferences. Shutting down engagement weighting $(\zeta_t = 0)$ in Scenario 2 produces a significant but not major decline, showing that the platform's consideration on users' preference is a not major driver of polarized content.

Second, the reduction in viewed slant does not generally coincide with the reduction

in supplied slant. While Scenario 1 delivers nearly identical reductions in both measures, Scenario 2 exhibits a substantial divergence: it achieves a modest reduction in supplied slant but a noticeably smaller reduction in viewed slant. This indicates that even when outlets adjust their content, the algorithm can continue to concentrate attention on relatively more slanted posts.

Finally, the fully neutral algorithm scenario ($\hat{\alpha}_t = \hat{\beta}_t = \zeta_t = 0$) achieves the full reduction. As newspapers only care about views but not engagement, this scenario effectively shuts down all channel of slant entering outlets' decision making problem.

Overall, the counterfactual analysis highlights the contribution of algorithmic amplification to the equilibrium polarization of online news.

9.3 Algorithm Effect

The previous results distinguish between extra slant at the content level and extra slant in user exposure. To quantify the role of the algorithm in bridging these two objects, I decompose viewed extra slant into a component attributable to media supply and a component attributable to algorithmic amplification.

For each scenario, let G^V denote viewed extra slant and G^S denote supplied extra slant. I define the algorithmic share as

$$\text{Algorithm Share} = \frac{G^V - G^S}{G^V}.$$

This measure captures the fraction of viewed extra slant that cannot be explained by editorial choices alone, and is therefore attributed to the recommendation algorithm's allocation of attention across posts.

Table 7 reports the shares across all scenarios, including the original equilibrium and the benchmark. Several insights emerge from this decomposition.

First, the algorithm plays a quantitatively important role in shaping exposure. In the

Table 7: Decomposition of Viewed Extra Slant

Scenario	Algorithm Share (%)
Benchmark	12.65
Scenario 1	13.33
Scenario 2	18.57
Scenario 3	12.66
Original	19.10

Notes: Supplied extra slant measures the average difference between online and print slant across posts, while viewed extra slant weights each post by its equilibrium probability of being recommended. Algorithm share is defined as the fraction of viewed extra slant not accounted for by supplied extra slant.

original equilibrium, approximately 19% of viewed extra slant is attributable to the algorithm, indicating that recommendation behavior significantly amplifies polarization beyond what is generated by media supply alone. At the same time, media outlets' response to algorithm accounts for the majority of viewed extra slant.

Second, a nontrivial share of viewed extra slant persists even in scenarios where slant is irrelevant in the algorithm's objective. In both the benchmark and Scenario 3, the algorithm share remains around 12–13%. One possible explanation is that media outlets differ in their baseline prominence and editorial capacity. More established outlets are more likely to engage in strategic headline editing and are also more likely to receive attention from users. As a result, differences in exposure across outlets can generate residual viewed extra slant even in the absence of explicit algorithmic amplification.

Finally, Scenario 1 delivers an outcome close to this lower bound. Removing the algorithm's direct sensitivity to ideological slant ($\hat{\alpha}_t = \hat{\beta}_t = 0$) reduces the algorithm share to a level similar to the benchmark and Scenario 3. This again indicates that most of the algorithm's contribution to polarization operates through its amplification of slant-related features. Scenario 2 highlights the role of engagement-based amplification. Although shutting down the engagement channel ($\zeta_t = 0$) reduces overall slant, the algorithm share remains relatively high, consistent with previous results.

Overall, these results show that algorithms accounts for a sizable share of slant exposure

and interacts with endogenous supply decisions. Understanding both channels is essential for evaluating how platform design affects the equilibrium distribution of news exposure.

9.4 Welfare and Diversity Trade-offs

The counterfactual analysis reveals a central trade-off between short-run user utility and the diversity of information exposure.

To quantify welfare, I compute the total expected user utility under each scenario using the estimated demand system. The results show that removing algorithmic amplification reduces user utility. Scenarios that shut down the algorithm's sensitivity to slant or engagement lead to substantial utility losses relative to the original equilibrium. This is expected, as the algorithm amplifies components of content that users find engaging, including ideological alignment and salience.

However, these welfare losses are accompanied by important changes in the distribution of viewed content. Figure 10 shows the share of viewed slant categories (negative, neutral, positive) across users with different ideological positions under each scenario. The original equilibrium exhibits clear polarization in exposure: left-leaning users disproportionately consume negatively slanted content, while right-leaning users consume more positively slanted content. This sorting reflects both user preferences and algorithmic amplification.

Scenario 3 delivers the most diverse exposure environment, corresponding to the case in which recommendation is independent of slant. In this scenario, users across the ideological spectrum are exposed to the same mix of content, with less concentration in any ideological categories. Scenario 1 achieves a very similar degree of flattening in exposure.

Importantly, these scenarios differ substantially in their welfare implications. While Scenario 3 generates a large welfare loss of approximately 75%, Scenario 1 results in a much smaller loss of about 20%. This suggests that it is possible for the platform to provide a more balanced and diverse information environment without substantially reducing user utility. In contrast, Scenario 2—closely related to Facebook's policy change—produces a comparable

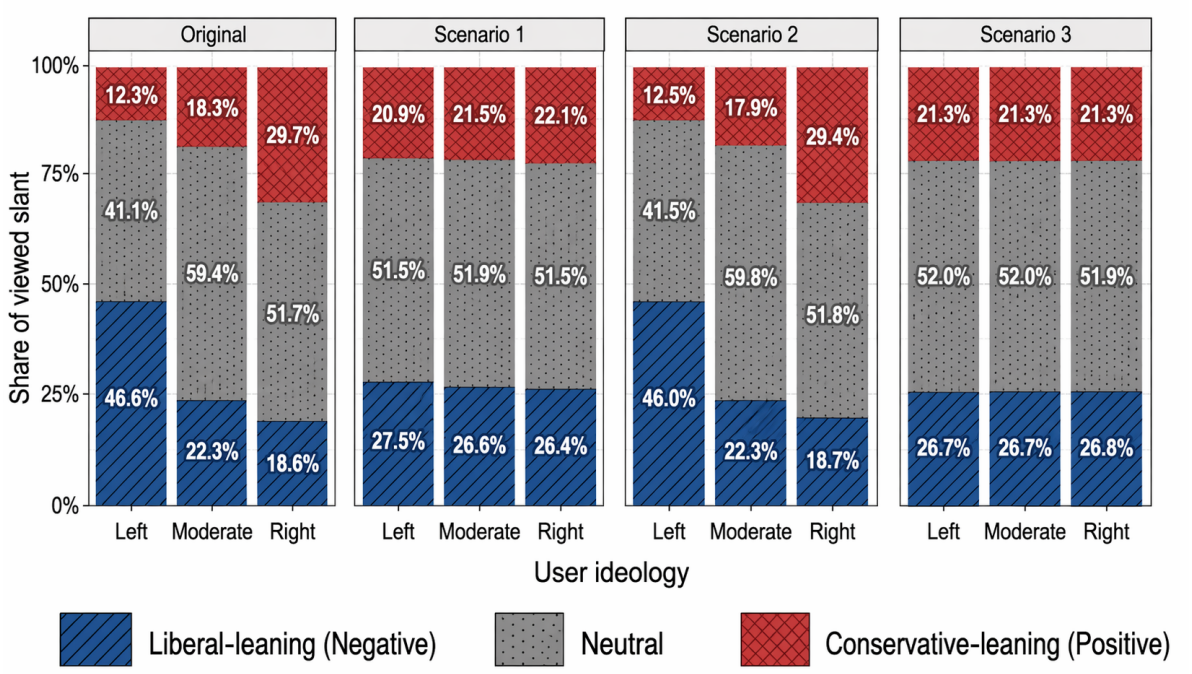


Figure 10: Share of viewed slant by user ideology across counterfactual scenarios. The figure reports the distribution of viewed slant categories (negative, neutral, positive) across user ideology groups under each scenario. For each group, the shares sum to one and reflect the composition of content that users are exposed to in equilibrium. The measure is based on the slant of viewed posts, rather than extra slant relative to print benchmarks. Differences across ideology groups capture polarization in exposure, while differences across scenarios reflect the effects of algorithmic and supply-side changes on the distribution of content.

welfare loss (around 19.5%) but leads to only limited improvements in exposure diversity.

Taken together, these results highlight a fundamental trade-off. Algorithmic amplification increases short-run user utility by directing attention toward content that users are more likely to engage with, but at the cost of greater polarization in exposure. Conversely, reducing amplification lowers engagement-based utility while improving the diversity and balance of information consumption. This trade-off suggests that platform design involves balancing engagement objectives against broader diversity outcomes. While engagement-maximizing algorithms increase immediate user satisfaction, they may also contribute to segregated information environments with potential long-run consequences for belief formation and media bias.

10 Policy Implications

The results of this paper have direct implications for the design and regulation of social-media platforms. A central insight is that algorithmic curation does not merely reallocate attention across existing content; it alters the incentives of media outlets and thereby reshapes the equilibrium supply of news. As a result, policies that focus solely on user exposure or short-run behavioral responses may fail to address the underlying source of distortion. Effective intervention must account for how platform design affects both demand (what users see) and supply (what outlets choose to produce).

10.1 Platform Design as Information Design

The structural estimates show that the platform algorithm overweights slant-related engagement signals relative to user preferences. This amplification creates a wedge between the platform's objective—maximizing engagement—and the socially optimal allocation of information. From an information design perspective, the platform chooses a mapping from content to exposure that shapes both user beliefs and producer incentives.

A natural implication is that platforms can reduce polarization through ex-ante adjustments to their recommendation objectives. In particular, reducing the weight placed on engagement-based signals for political content, regularizing ideological features in the ranking function, or imposing constraints that increase exposure diversity can all attenuate the incentive to produce extreme content. The counterfactual analysis shows that removing algorithmic amplification of ideological signals would reduce extra slant by approximately 90 percent, indicating that relatively small changes in platform objectives can have large equilibrium effects.

Importantly, these interventions operate through the allocation rule rather than through direct restrictions on content. By changing how visibility is assigned, the platform alters the marginal return to slant faced by media outlets, thereby affecting supply in equilibrium.

10.2 Supply-Side Responses and Regulatory Monitoring

A second implication is that regulation should explicitly account for supply responses by content producers. In the model, media outlets primarily optimize for visibility rather than engagement, implying that algorithmically determined exposure is the key driver of editorial decisions. Consequently, policies that target only user-side outcomes may be insufficient if they do not alter producer incentives.

The empirical results provide a framework for separating algorithmic effects from user preferences. Following the platform's policy change in 2022, exposure (measured by views) declined significantly, while engagement conditional on exposure remained stable. This pattern indicates that the observed changes in content supply were driven by shifts in algorithmic allocation rather than changes in user demand.

This decomposition suggests two complementary approaches to regulatory monitoring. First, exposure-based auditing can test whether content with different ideological characteristics receives systematically different recommendation probabilities, conditional on observables. Second, comparing engagement conditional on exposure provides a way to distinguish whether observed disparities arise from user preferences or from algorithmic design. Together, these tools allow regulators to identify distortions in the allocation of attention without requiring direct observation of the algorithm itself.

10.3 Polarization and Welfare

The counterfactual results reveal a central trade-off in platform design. Reducing algorithmic amplification lowers user engagement and short-run utility, but simultaneously leads to a more balanced distribution of information exposure.

On the one hand, the algorithm amplifies components of user utility that are positively valued—namely ideological alignment and the salience of unexpected content. As a result, removing these amplification channels mechanically reduces expected utility in the short run. This effect is reflected in the utility loss estimates reported in Section 9.4.

On the other hand, weakening algorithmic amplification substantially changes the composition of viewed content. As shown in Figure 10, users are exposed to a more diverse and ideologically balanced set of news when amplification is reduced. This suggests that the algorithm not only increases engagement but also concentrates attention on ideologically extreme content.

These findings highlight a tension between engagement maximization and exposure diversity. While engagement-based metrics favor content that is aligned with user preferences or highly salient, such content may reduce exposure to opposing viewpoints and contribute to polarization. A large literature is based on the idea that more diverse media environments can mitigate bias and improve information aggregation (e.g., Gentzkow et al., 2014, Bakshy et al., 2015, Levy, 2021).

Together, the results imply that policies aimed at reducing polarization through algorithmic design are likely to involve short-run costs in terms of engagement and user utility. However, these costs may be offset by long-run gains from a more balanced and informative media environment. This trade-off is central to evaluating the welfare consequences of platform regulation.

10.4 Broader Implications

More broadly, the findings highlight a general principle for digital markets: when intermediaries control visibility, their objective functions shape equilibrium outcomes in upstream markets. In such environments, algorithm design becomes a central policy lever, analogous to pricing or entry regulation in traditional industrial organization settings.

This insight extends beyond news to other contexts in which platforms mediate information or product discovery. By shaping the mapping from content to exposure, platforms influence not only what users consume, but also what producers choose to supply. As a result, understanding and regulating algorithmic design is essential for addressing distortions in modern information markets.

11 Conclusion

This paper documents and quantifies how social-media algorithms shape the supply of news. Using newly linked data that match print and online newspaper headlines to millions of Facebook posts and user engagement metrics from 2017–2024, I show that the design of platform algorithms not only determines what users see, but also feeds back to influence what news outlets produce.

The reduced-form evidence around Facebook's July 2022 announcement to "reduce political content" provides a clean quasi-experiment. When the platform changed its recommendation policy, politically slanted news headlines became significantly less divergent from their print counterparts, while sentiment and posting frequency remained stable. This pattern indicates a deliberate editorial adjustment—outlets strategically moderated their ideological framing on social media when algorithmic incentives shifted, rather than responding to any contemporaneous change in user demand.

Building on this evidence, I develop a dynamic structural model in which users, algorithms, and media outlets interact in a digital news market. Users engage with content based on ideological alignment and salience; algorithms recommend posts that maximize predicted engagement; and media outlets dynamically choose headline slant to attract visibility. Estimating the model with post-level data reveals that prior to Facebook's policy change, the algorithm effectively tripled the marginal return to politically slanted content relative to intrinsic user preferences. Outlets responded by increasing online slant to maximize expected views, consistent with their estimated objective parameters that heavily weight visibility but place little value on engagement metrics.

Counterfactual simulations using the estimated model quantify the magnitude and channels of algorithmic amplification. Eliminating ideological weighting from the recommendation system reduces the extra slant of online news by roughly one order of magnitude and substantially flattens the distribution of exposure across users with different ideologies, while generating only modest welfare losses. In contrast, reducing engagement-based tar-

geting—similar to recent platform interventions—yields comparable welfare losses but leads to only limited improvements in exposure diversity. A fully neutral algorithm that removes both channels would eliminate nearly all of the observed divergence between print and online headlines, but at a large welfare cost. Because newsrooms primarily optimize for visibility, the platform's design shapes not only how content is distributed but also the equilibrium level of slant in online news supply.

Taken together, the findings imply that algorithmic curation creates a general-equilibrium feedback loop between user engagement, platform design, and editorial strategy. When engagement-based algorithms reward polarization, outlets respond by producing more polarized content; when platforms demote such content, supply moderates even without a shift in user tastes. These results provide the first structural quantification of how digital recommendation systems reshape the incentives of media producers, highlighting that algorithm design has consequences not only for the demand for information, but for its very supply. The approach also demonstrates how structural tools from industrial organization and information design can be applied to media economics to study dynamic platform markets where firms, algorithms, and consumers interact in real time.

Acknowledgments. This research did not receive any external funding. I have no financial relationships or conflicts of interest to disclose.

References

Aguirregabiria, V. and Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53.

Aldous, K. K., Şengün, S., Salminen, J., Farooq, A., Jung, S.-G., and Jansen, B. J. (2025). Multiple engagement by an individual on a social media post is rare: Insight from an anal-

ysis of 3.5 million instagram user accounts and 29 user interviews. *Data and Information Management*, page 100113.

Alesina, A., Miano, A., and Stantcheva, S. (2020). The polarization of reality. *AEA Papers and Proceedings*, 110:324–28.

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–36.

Arcidiacono, P., Bayer, P., Bugni, F. A., and James, J. (2013). Approximating high-dimensional dynamic models: Sieve value function iteration. 31:45–95.

Arcidiacono, P. and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867.

Armona, L., Gentzkow, M., Kamenica, E., and Shapiro, J. M. (2024). What is newsworthy? theory and evidence. NBER Working Paper 32512, National Bureau of Economic Research. Available at SSRN: https://ssrn.com/abstract=4851786.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Bakshy, E., Messing, S., and Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132.

Bergemann, D. and Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11:85–107.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2022). Salience. *Annual Review of Economics*, 14(Volume 14, 2022):521–544.

Braghieri, L., Eichmeyer, S., Levy, R., Mobius, M. M., Steinhardt, J., and Zhong, R. (2024). Article-level slant and polarization of news consumption on social media. Working Paper.

Cagé, J., Hervé, N., and Mazoyer, B. (2024). Social media influence on mainstream media: Evidence from two billion tweets. *SSRN Electronic Journal*.

Cagé, J., Hervé, N., and Viaud, M.-L. (2019). The production of information in an online world. *The Review of Economic Studies*, 87(5):2126–2164.

Calzada, J. and Gil, R. (2020). What do news aggregators do? evidence from google news in spain and germany. *Marketing Science*, 39(1):134–167.

Chiang, C.-F. and Knight, B. (2011). Media bias and influence: Evidence from newspaper endorsements. *The Review of Economic Studies*, 78(3):795–820.

DellaVigna, S. and Kaplan, E. (2007). The fox news effect: Media bias and voting*. *The Quarterly Journal of Economics*, 122(3):1187–1234.

Gentzkow, M. (2006). Television and voter turnout*. *The Quarterly Journal of Economics*, 121(3):931–972.

Gentzkow, M. and Shapiro, J. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2):280–316.

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? evidence from u.s. daily newspapers. *Econometrica*, 78(1):35–71.

Gentzkow, M., Shapiro, J. M., and Sinkinson, M. (2014). Competition and ideological diversity: Historical evidence from us newspapers. *The American Economic Review*, 104(10):3073–3114.

Gentzkow, M., Shapiro, J. M., and Taddy, M. (2019). Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4):1307–1340.

González-Bailón, S., Lazer, D., Barberá, P., Zhang, M., Allcott, H., Brown, T., Crespo-Tenorio, A., Freelon, D., Gentzkow, M., Guess, A. M., Iyengar, S., Kim, Y. M., Malhotra, N., Moehler, D., Nyhan, B., Pan, J., Rivera, C. V., Settle, J., Thorson, E., Tromble, R., Wilkins, A., Wojcieszak, M., de Jonge, C. K., Franco, A., Mason, W., Stroud, N. J., and Tucker, J. A. (2023). Asymmetric ideological segregation in exposure to political news on facebook. *Science*, 381(6656):392–398.

Groseclose, T. and Milyo, J. (2005). A measure of media bias. *The Quarterly Journal of Economics*, 120(4):1191–1237.

Guess, A. M. (2021). (almost) everything in moderation: New evidence on americans' online media diets. *American Journal of Political Science*, 65(4):1007–1022.

Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., González-Bailón, S., Kennedy, E., Kim, Y. M., Lazer, D., Moehler, D., Nyhan, B., Rivera, C. V., Settle, J., Thomas, D. R., Thorson, E., Tromble, R., Wilkins, A., Wojcieszak, M., Xiong, B., de Jonge, C. K., Franco, A., Mason, W., Stroud, N. J., and Tucker, J. A. (2023a). How do social media feed algorithms affect attitudes and behavior in an election campaign? *Science*, 381(6656):398–404.

Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., González-Bailón, S., Kennedy, E., Kim, Y. M., Lazer, D., Moehler, D., Nyhan, B., Rivera, C. V., Settle, J., Thomas, D. R., Thorson, E., Tromble, R., Wilkins, A., Wojcieszak, M., Xiong, B., de Jonge, C. K., Franco, A., Mason, W., Stroud, N. J., and Tucker, J. A. (2023b). Reshares on social media amplify political news but do not detectably affect beliefs or opinions. *Science*, 381(6656):404–408.

Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529.

Hotz, V. J., Miller, R. A., Sanders, S., and Smith, J. (1994). A simulation estimator for dynamic models of discrete choice. *The Review of Economic Studies*, 61(2):265–289.

Leung, T. C. and Strumpf, K. S. (2023). All the headlines that are fit to change. Working Paper.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3):831–70.

Martin, G. J. and Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–99.

Moehring, A. (2024). Personalization, engagement, and content quality on social media: An evaluation of reddit's news feed. Working Paper.

Mullainathan, S. and Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4):1031–1053.

Nyhan, B., Settle, J., Thorson, E., Wojcieszak, M., Barberá, P., Chen, A. Y., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., González-Bailón, S., Guess, A. M., Kennedy, E., Kim, Y. M., Lazer, D., Malhotra, N., Moehler, D., Pan, J., Thomas, D. R., Tromble, R., Rivera, C. V., Wilkins, A., Xiong, B., de Jonge, C. K., Franco, A., Mason, W., Stroud, N. J., and Tucker, J. A. (2023). Like-minded sources on facebook are prevalent but not polarizing. *Nature*, 620(7972):137–144.

Pakes, A. and McGuire, P. (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse' of dimensionality. *Econometrica*, 69(5):1261–1281.

Piccardi, T., Saveski, M., Jia, C., Hancock, J., Tsai, J. L., and Bernstein, M. S. (2025).

Reranking partisan animosity in algorithmic social media feeds alters affective polarization. *Science*, 390(6776):eadu5584.

Saeedi, M. (2019). Reputation and adverse selection: theory and evidence from ebay. *The RAND Journal of Economics*, 50(4):822–853.

Saeedi, M., Shen, Y., and Shourideh, A. (2024). Getting the agent to wait. Carnegie Mellon University Working Paper.

Strömberg, D. (2004). Radio's impact on public spending*. *The Quarterly Journal of Economics*, 119(1):189–221.

Sundar, S. S., Snyder, E. C., Liao, M., Yin, J., Wang, J., and Chi, G. (2025). Sharing without clicking on news in social media. *Nature Human Behaviour*, 9(1):156–168.

Weintraub, G. Y., Benkard, C. L., and Van Roy, B. (2008). Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411.

A Additional Tables and Figures

Table 8: Facebook Posts Metrics in Different Period

	17-04-01 to 22-05-01			22-05-01 to 24-12-31		
	All	News	Ratio	All	News	Ratio
Posts (t.)	1,565	997	64%	847	486	57%
Views (b.)	357	218	61%	79	27	34%
Comments (m.)	755	559	74%	163	98	60%
Shares (m)	661	464	70%	46	23	50%

Between April 2017 and May 2022, news-related posts accounted for roughly 64% of all media posts on Facebook, attracting 61% of total views. After Facebook's mid-2022 algorithm change to reduce political content, the share of news in total views dropped sharply to 34%, and the ratio of news-related shares fell from 70% to 50%. This pattern confirms that the platform update substantially lowered visibility for political news.

Table 9: Summary of articles by measure, 2017-04-01 onward. Sentiment is estimated for every article; partisanship is estimated for articles with topic as news & social concern.

Measure	Total	Print	Online	Both	Different	Online Higher	Print Higher
Sentiment	3,427,116	1,760,040	2,659,357	992,281	925,449	209,091	248,052
Partisanship	1,621,428	717,189	1,344,451	440,212	405,712	18,165	23,833

	Sentiment				Partisanship			
Signs	Online Higher	Print Higher	Equal	Total	Online Higher	Print Higher	Equal	Total
$-\|\|+$	2,990	4,635	0	7,625	124	225	0	349
$-\|\|0$	62,514	94,317	0	156,831	4,607	15,923	0	20,530
$0\|\|+$	52,437	50,322	0	102,759	6,639	1,968	0	8,607
$-\|\|-$	65,803	73,114	0	138,917	5,966	3,267	1,183	10,416
$+\|\|+$	25,347	25,664	0	51,011	829	2,450	316	3,595
$0\|\|0$	0	0	468,306	468,306	0	0	362,215	362,215
Total	209,091	248,052	468,306	925,449	18,165	23,833	363,714	405,712

Table 10: Summary of Facebook posts with links from 2017-04-01 onward. Sentiment is estimated for every post; partisanship is estimated for posts with topic as news & social concern.

Measure	Total	With Text	Positive Link	Negative Link	Positive Text	Negative Text
Sentiment	2,929,991	2,825,485	200,933	914,269	341,425	908,671
Partisanship	1,792,168	1,717,723	44,538	96,484	79,238	215,071

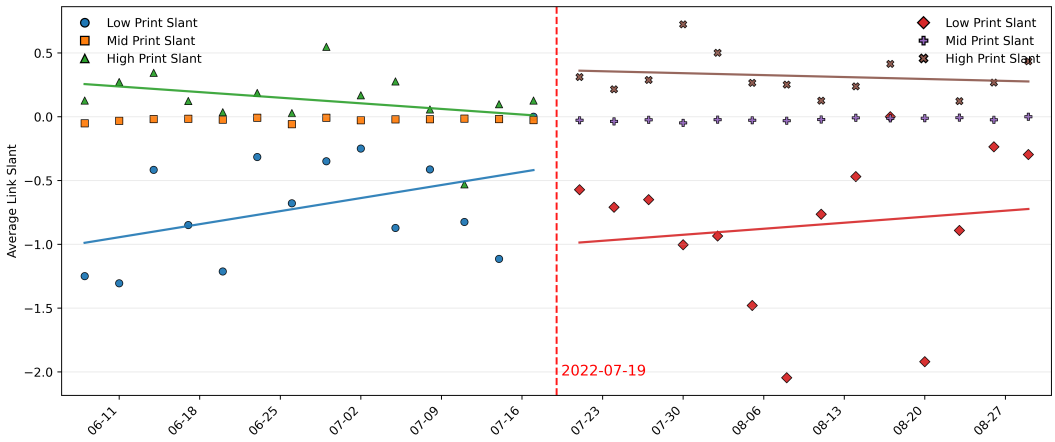


Figure 11: Binary Approximation for Online Slant $s^n_{j,t}$ by Print Slant $s^n_{j,t}$ (three bins) pre vs. post announcement.

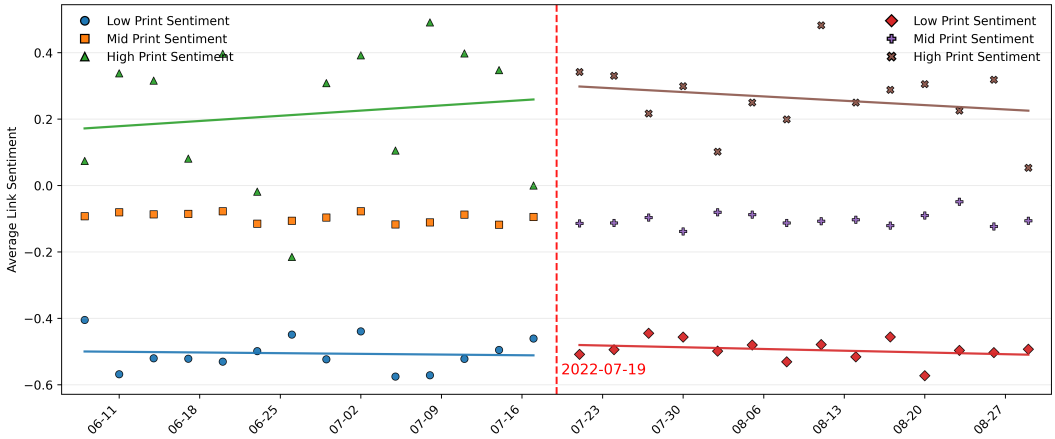


Figure 12: Binary Approximation for Online Sentiment $s^n_{j,t}$ by Print Sentiment $s^n_{j,t}$ (three bins) pre vs. post announcement.

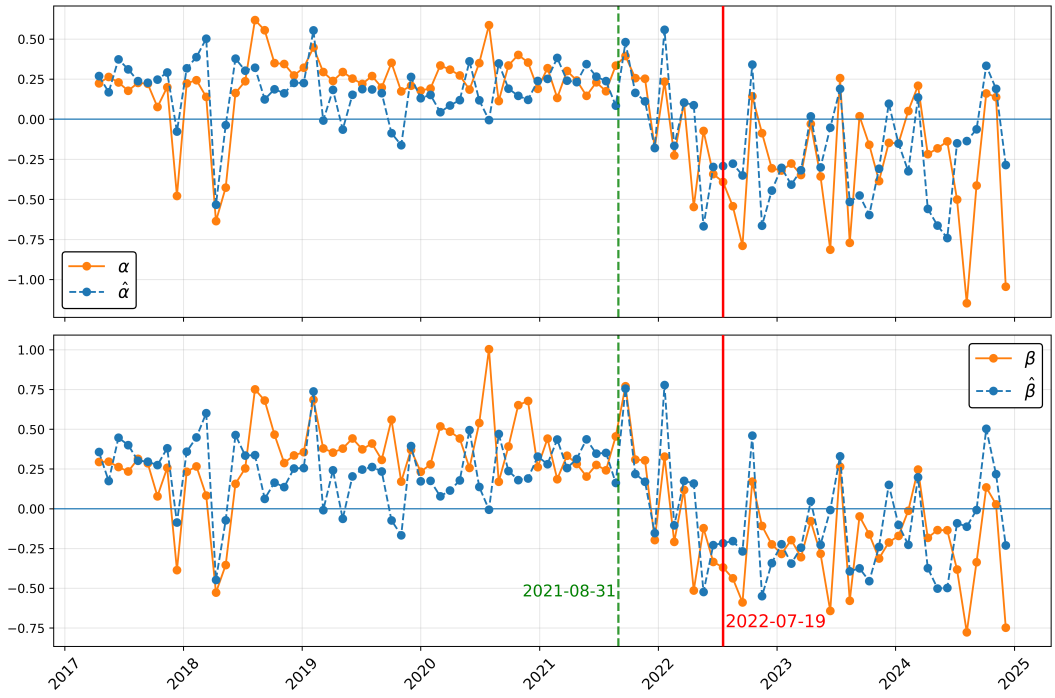


Figure 13: Estimated algorithmic weight $\alpha_t, \beta_t, \widehat{\alpha}_t, \widehat{\beta}_t$ before and after the 2022 policy change. The parameter measures how users and algorithm's utility related to slant.

B Newsroom Changes

Wall Street Journal. In August 2019, the *WSJ* built a *Newsroom Innovation* group with a dedicated SEO team and hired an SEO editor (Edward Hyatt), embedding search strategy alongside newsletters and product experimentation.[14]

Washington Post. The Post institutionalized platform work via an "Audience" team and later social-video initiatives (e.g., its well-known TikTok presence launched in 2019 under Dave Jorgenson). Public postings in the mid-to-late 2010s and early 2020s include social media editors and SEO leadership roles (e.g., "Head of AI Discovery & SEO").[15]

USA Today (Gannett). Gannett/USA Today Network has long maintained centralized audience/SEO and social teams across properties, reflected in frequent postings for audience, trending, and platform roles (mid-2010s onward). These roles focus on search optimization, real-time packaging, and cross-network distribution.[16]

Los Angeles Times. The *LA Times* scaled short-form and platform teams in the early 2020s (e.g., explicit TikTok/Instagram focus noted in 2023 coverage of legacy outlets' youth strategies), complementing multiplatform editor roles responsible for headline testing, metadata, and social copy.[17]

[14] NiemanLab (Aug. 12, 2019), "How The Wall Street Journal is building an incubator into its newsroom with new departments and plenty of hires." https://www.niemanlab.org/2019/08/how-the-wall-street-journal-is-building-an-incubator-into-its-newsroom-with-new-departments-and-plenty

[15] Vanity Fair (Feb. 2023), "Legacy Media Wants a Piece of That Gen Z TikTok Mojo" (notes WaPo's TikTok team). https://www.vanityfair.com/news/2023/02/tiktok-media-new-york-times-washington-post. See also Mediabistro job listing, "Head of AI Discovery & SEO," https://www.mediabistro.com/jobs/1658352006-the-washington-post-is-hiring-head-of-ai-discovery-and-seo-in-washington. For a staff bio reflecting social editor roles, see Nina Zafar, Social Media Editor: https://www.washingtonpost.com/people/nina-zafar/.

[16] Gannett careers portal: https://www.gannett.com/careers/. Representative network postings emphasize analytics-driven packaging and social programming.

[17] Vanity Fair (Feb. 2023), https://www.vanityfair.com/news/2023/02/tiktok-media-new-york-times-washington-post. LAT newsroom announcements and postings reference multiplatform/audience roles; see company press and careers pages.

Chicago Tribune & Boston Globe. By the mid-2010s, both organizations maintained digital desks with multiplatform editors and social roles; public job histories and postings indicate ongoing platform work (web producing, social packaging, SEO best practices).[18]

C Contraction of the Share Inversion Mapping

Here I drop the subscript $\{j,t\}$ for cleaner notations. Fix parameters $(\beta, \widehat{\beta}, \zeta)$ and the distribution $F(\mu_i)$. Define the predicted shares of post n under equations (10) as

$$\widehat{S}^{n,E}(\delta^n, \widehat{\delta}^n) = \int_i \frac{\exp(W_i^n)}{1+\exp(W_i^n)} \frac{\exp(\widehat{W}_i^n)}{1+\sum_m \exp(\widehat{W}_i^m)} dF(\mu_i)$$
$$\widehat{S}^{n,V}(\delta^n, \widehat{\delta}^n) = \int_i \frac{\exp(\widehat{W}_i^n)}{1+\sum_m \exp(\widehat{W}_i^m)} dF(\mu_i)$$

Let the share–inversion operator be

$$\mathcal{T}_E(\delta^n, \widehat{\delta}^n) = \delta^n + \log S^{n,E} - \log \widehat{S}^{n,E}(\delta^n, \widehat{\delta}^n) \tag{25}$$
$$\mathcal{T}_V(\delta^n, \widehat{\delta}^n) = \widehat{\delta}^n + \log S^{n,V} - \log \widehat{S}^{n,V}(\delta^n, \widehat{\delta}^n) \tag{26}$$

and denote $\mathcal{T}(\delta, \widehat{\delta}) = (\mathcal{T}_E(\delta, \widehat{\delta}), \mathcal{T}_V(\delta, \widehat{\delta}))$. Then $\mathcal{T}$ is a contraction mapping on the product space of $(\delta, \widehat{\delta})$ under the logit structure with a large outside option (i.e. $\sum_n S^{n,V} \ll 1$), and therefore possesses a unique fixed point $(\delta^*, \widehat{\delta}^*)$ satisfying

$$S^{n,E} = \widehat{S}^{n,E}(\delta^*, \widehat{\delta}^*), \qquad S^{n,V} = \widehat{S}^{n,V}(\delta^*, \widehat{\delta}^*).$$

To see this, let $\Delta = (\delta_1 - \delta_2, \widehat{\delta}_1 - \widehat{\delta}_2)$ and denote the corresponding predicted shares

[18] Examples: Boston Globe Media careers portal (audience/multiplatform roles): https://www.bostonglobemedia.com/opportunities/. Public job histories illustrating Tribune digital roles (e.g., digital editors since 2015) can be found through staff CVs/profiles.

$(\widehat{S}_1^E, \widehat{S}_1^V)$ and $(\widehat{S}_2^E, \widehat{S}_2^V)$. By the mean–value theorem,

$$\mathcal{T}_E(\delta_1, \widehat{\delta}_1) - \mathcal{T}_E(\delta_2, \widehat{\delta}_2) = -J_E(\tilde{\delta}, \tilde{\widehat{\delta}})\,\Delta, \qquad \mathcal{T}_V(\delta_1, \widehat{\delta}_1) - \mathcal{T}_V(\delta_2, \widehat{\delta}_2) = -J_V(\tilde{\delta}, \tilde{\widehat{\delta}})\,\Delta,$$

where J_E and J_V are the Jacobians of $\log \widehat{S}^E$ and $\log \widehat{S}^V$ with respect to $(\delta, \widehat{\delta})$, evaluated at some intermediate point $(\tilde{\delta}, \tilde{\widehat{\delta}})$.

Denote the individual shares as

$$\begin{aligned}
\widehat{S}_i^{n,E|V}(\delta^n, \widehat{\delta}^n) &= \frac{\exp(W_i^n)}{1 + \exp(W_i^n)} \\
\widehat{S}_i^{n,V}(\delta^n, \widehat{\delta}^n) &= \frac{\exp(\widehat{W}_i^n)}{1 + \sum_m \exp(\widehat{W}_i^m)} \\
\widehat{S}_i^{n,E}(\delta^n, \widehat{\delta}^n) &= \frac{\exp(W_i^n)}{1 + \exp(W_i^n)} \frac{\exp(\widehat{W}_i^n)}{1 + \sum_m \exp(\widehat{W}_i^m)}
\end{aligned}$$

which gives Jacobian elements as

$$\begin{aligned}
\frac{\partial \widehat{S}_i^{n,V}}{\partial \widehat{\delta}^m} &= \widehat{S}_i^{n,V} \begin{cases} 1 - \widehat{S}_i^{m,V}, & n = m, \\ -\widehat{S}_i^{m,V}, & n \neq m, \end{cases} \\
\frac{\partial \widehat{S}_i^{n,E}}{\partial \delta^m} &= \widehat{S}_i^{n,E} \begin{cases} (1 - \widehat{S}_i^{m,E|V}) + \zeta(\widehat{S}_i^{m,E|V} - \widehat{S}_i^{m,E}), & n = m, \\ -\zeta \widehat{S}_i^{m,E}, & n \neq m, \end{cases}
\end{aligned}$$

The argument follows Berry et al. (1995). When the outside option share $s_0 = 1 - \sum_n \widehat{s}^{n,V}$ is positive (as in the data, where total view share $\ll 1$) and ζ is small enough such that the share of non-engagement $1 - \zeta \sum_n \widehat{s}^{n,E}$ is also positive (as in the data, where engagement share $\ll 1$ even after being multiplied by ζ), each Jacobian is strictly diagonally dominant with spectral radius $\rho(J_E), \rho(J_V) < 1$. Therefore

$$\|\mathcal{T}(\delta_1, \widehat{\delta}_1) - \mathcal{T}(\delta_2, \widehat{\delta}_2)\| \leq \rho\,\|\Delta\|, \qquad \rho = \max\{\rho(J_E), \rho(J_V)\} < 1,$$

so $\mathcal{T}$ is a contraction. By the Banach fixed–point theorem, the iterative updates in the text converge to the unique $(\delta^*, \widehat{\delta}^*)$ satisfying the observed shares.