LLMs as Gatekeepers: Source Concentration, Factual Quality, and Political Slant in Information Search

Pengxiang (Shawn) Zhou[1] Davide Proserpio[2] Ali Goli[3]

May 31, 2026

[1]University of Southern California; Pengxiang.Zhou@marshall.usc.edu.
[2]University of Southern California; proserpi@marshall.usc.edu.
[3]University of Rochester; agoli@ur.rochester.edu.

Abstract

As LLM-based search platforms displace traditional search, the information environment they curate increasingly shapes what users see. We study these platforms as a new class of information intermediary. Using 48 months of referral traffic data for 8,082 news domains across five AI platforms (ChatGPT, Perplexity, Gemini, Claude, Grok) and traditional search, linked to publisher-level bot-blocking records, we document three facts. First, AI platforms surface more concentrated content than organic search, but with higher factual quality and, on average, a comparable political slant. No platform dominates: Perplexity is the most factually reliable and most centrist but also the most concentrated, Grok distributes traffic most broadly but is the least factually reliable, and Claude is the most left-leaning. Second, publishers that block AI crawlers via `robots.txt` are disproportionately the largest, most factually reliable, and centrist-to-left outlets. This selective exit shifts ChatGPT's and Perplexity's pools toward higher concentration, lower factual quality, and rightward slant; Gemini's concentration also rises, but its factual quality and slant remain largely flat. Third, focusing on Perplexity, we show that blocking Perplexity's crawler cuts the blocking domain's referral traffic by 46–48%. Moreover, the referral traffic that would have gone to banning domains flows to the largest and highest-factual-quality remaining accessible outlets, with size being the dominant channel, so the referral redistribution reinforces concentration among larger domains rather than diluting it. Together, these findings show that selective publisher opt-outs reduce the factual quality of the pool AI platforms can refer to, and that the resulting referral-traffic redistribution reinforces concentration while failing to undo the compositional shift toward less factual and more right-leaning publishers.

Keywords: Large Language Models, AI Search, Information Intermediaries, Publisher Opt-outs, Source Concentration, Factual Quality, Political Slant, Staggered Difference-in-Differences

1 Introduction

Large language models (LLMs) are emerging as a new class of information intermediary. Platforms such as ChatGPT, Perplexity, Gemini, Claude, and Grok now attract hundreds of millions of users who use them to query for news, analysis, and factual information, bypassing traditional search engines. Recent clickstream evidence shows that LLM adoption reduces traditional search activity and downstream browsing, particularly to smaller sites, and that AI-generated summaries suppress click-through to external links (Padilla et al., 2025). As AI platforms displace traditional search, the information environment they curate increasingly shapes what users see. This raises a first question: do AI platforms distribute attention across the information ecosystem in the same way as traditional search, or do they systematically favor certain sources in terms of source concentration, factual quality, and political orientation? And how do AI platforms differ from one another? We measure this environment through publisher referral traffic, i.e., the share of user attention each AI platform directs to each source.

The content AI platforms can surface, however, is itself shaped by publishers. Since mid-2023, publishers have responded to AI intermediation on two fronts: a surge of `robots.txt` opt-outs against AI crawlers, and a wave of lawsuits led by *The New York Times*' December 2023 suit against OpenAI.[1] By December 2025, Gemini, ChatGPT, and Perplexity had accumulated the largest ban counts, with bans concentrated among the most factually reliable outlets. This raises two further questions. First, how does the disproportionate exit of publishers shift the content *supply*, i.e., the source concentration, factual quality, and political slant of the pool still available to AI platforms? Second, how do platforms redistribute the referral traffic among remaining accessible domains, and how does that reshape the *exposure* users ultimately receive?

In this paper, we study these three questions using a panel of 8,082 news and media

[1] Complaint, *The New York Times Company v. Microsoft Corporation et al.*, No. 1:23-cv-11195 (S.D.N.Y. filed Dec. 27, 2023); Complaint, *Dow Jones & Co., Inc. v. Perplexity AI, Inc.*, No. 1:24-cv-07984 (S.D.N.Y. filed Oct. 21, 2024).

domains observed monthly from January 2022 through December 2025. The panel combines three data sources. First, we obtain domain-level quality ratings from Media Bias/Fact Check (MBFC), an independent organization that rates publishers on political bias and factual reporting. MBFC publishes both categorical labels and numeric ratings on two dimensions: political bias on a continuous scale from -10 (Extreme Left) to $+10$ (Extreme Right), with Least Biased outlets near zero; and factual reporting on a scale from 0 (Very High) to 10 (Very Low), so lower scores indicate higher factual reliability. Second, we obtain monthly referral traffic estimates from SEMrush across 15 channels, including five AI platforms (ChatGPT, Perplexity, Gemini, Claude, Grok) and traditional channels such as Organic Search and Direct Visit. Third, we reconstruct a monthly `robots.txt` bot-blocking panel from Wayback Machine and Common Crawl archives. This dataset allows us to observe, for each domain and month, both the traffic it receives from each AI platform and whether it has blocked that platform's crawler.

To characterize each platform's content pool, we construct three traffic-weighted indices. The *Factual Index* is the traffic-weighted average of domain factual scores, and lower values indicate a higher-quality pool. The *Slant Index* is the traffic-weighted average of domain political bias scores, and negative values indicate a left-leaning pool, while positive values a right-leaning pool. The *HHI* (Herfindahl-Hirschman Index) is the sum of squared traffic shares, multiplied by 10,000. HHI measures source concentration: a platform that spreads traffic evenly across many sources has a low HHI, while a platform dominated by a few sources has a high HHI. We use these indices in two distinct ways. To describe what each platform *surfaces*, we weight domains by their realized AI referral traffic, so the weights reflect where users are actually sent. To trace how publisher exit reshapes the *available* pool, we instead compute each index over the not-yet-banned domains, weighted by their fixed pre-ban size (2022 Organic Search traffic), so that index movements reflect changes in availability rather than model ranking or demand.

We document three facts. First, in the 2025 cross-section, AI platforms produce a structurally different pattern of source exposure than organic search. Newspapers account for 23.6% of organic search referral traffic but only 6–13% on AI platforms, while academic journals and organizations/foundations are over-represented. AI platforms are 3–8× more concentrated than Organic Search, with HHI ranging from 166 (Grok) to 475 (Perplexity) compared to 57 for Organic Search; but they score 13–33% better on the Factual Index (1.68–2.18 vs. 2.50), indicating higher average factual quality. On slant, where negative values indicate a more left-leaning source mix, AI platforms vary widely: Perplexity (−0.41) and Gemini (−0.42) are slightly more centrist than Organic Search (−0.43), while Claude is markedly more left-leaning (−0.72). No platform dominates across all three dimensions: Perplexity combines the highest factual quality and most centrist slant with the highest concentration, Grok distributes traffic most broadly but has the lowest factual quality, and ChatGPT, Claude, and Gemini occupy different points along this concentration–quality–slant frontier.

Second, the publishers that block AI crawlers via `robots.txt` are not a random draw, and the selective exit moves aggregate pool quality. Newspapers, centrist-to-left outlets, and Very High/High factual quality domains account for the vast majority of bans, while low-quality and politically peripheral publishers are nearly absent. Tracking these changes over time, with each index computed over the domains that have not yet banned a platform's crawler, we find rising HHI, declining factual quality, and rightward slant drift for ChatGPT and Perplexity. Gemini's HHI also rises but its Factual and Slant indices remain largely flat. Claude and Grok stay comparatively stable because fewer publishers have banned them so far.

Third, we estimate how traffic redistributes when publishers exit the accessible pool. This analysis requires observing referral traffic both before and after publishers block a platform's crawler. In our SEMrush data, Perplexity provides the strongest setting for this test since it has the longest AI-platform coverage window, beginning in December 2022, and substantial variation in publisher ban timing. We therefore focus on Perplexity and exploit staggered `robots.txt` bans in a difference-in-differences design. We first show that bans are effective: blocking Perplexity's crawler cuts the blocking domain's own referral traffic by 46–48%. We then ask how the traffic displaced by a ban is substituted among non-blocking domains. The substitution is primarily in kind, with publisher size: traffic displaced by a blocking outlet is reallocated most strongly to the remaining accessible outlets in the same size tier. Because large publishers ban most often, this concentrates displaced traffic among the remaining large outlets rather than passing it to smaller ones. Factual quality is a weaker but still important channel: traffic displaced by high-quality outlets tends to move to other high-quality outlets that remain accessible, while bans by low-quality outlets reduce traffic to the remaining low-quality outlets rather than benefiting them. Political slant plays a much smaller role: substitution is concentrated around least-biased and left-center outlets, while far-left and far-right outlets lose traffic when their ideological peers block Perplexity. A placebo test that randomly reassigns publisher size, factual quality, and slant ratings corroborates these patterns.

Together, these findings show that AI intermediation produces an information environment that is structurally distinct from traditional organic search. Moreover, when publishers decide to ban these AI platforms from referring them, referral traffic flows to the largest, highest-factual-quality remaining accessible outlets. The traffic redistribution, therefore, accelerates the concentration trend of referral traffic we observed rather than reducing it.

2 Related Literature

How information intermediaries shape the source concentration, factual quality, and political slant of the content users encounter has been a longstanding concern in the political economy of media. Our paper relates to three strands of literature on this question.

First, we contribute to the literature on how intermediaries shape the information environment users encounter as measured by its source concentration, factual quality, and political composition. The factual reporting and political slant of news content are themselves endogenous to market forces. Consumer demand for like-minded news drives slant more than owner identity does (Gentzkow and Shapiro, 2010), and most major U.S. outlets sit to the left of the average member of Congress (Groseclose and Milyo, 2005). Competition among media shapes which facts reach readers (Zhu and Dukes, 2015). Intermediaries then reshape which of this content users actually encounter, as ranking mechanisms steer exposure (Robertson et al., 2018). Online news consumption remains only modestly segregated (Gentzkow and Shapiro, 2011), and search engines' rich-get-richer concentration of traffic is tempered by users' heterogeneous topical interests (Fortunato et al., 2006). Steinacker-Olsztyn et al. (2026) document that trusted news sites are more likely to block AI crawlers than misinformation sites. We add to this literature by categorizing and analyzing AI platforms and their referral traffic across three dimensions (source concentration, factual quality, and political slant) and by documenting how each dimension shifts as high-quality publishers selectively exit.

Second, we contribute to the literature on platform intermediation and the redistribution of attention. News aggregators redirect traffic and lower discovery costs (Athey et al., 2021; Calzada and Gil, 2020; Chiou and Tucker, 2017). LLM adoption further substitutes for traditional search and reduces downstream click-through (Padilla et al., 2025), and Google's AI Overview cuts daily traffic to exposed Wikipedia articles (Khosravi and Yoganarasimhan, 2026). Publisher opt-outs intended as a defensive response often reduce rather than restore traffic (Zhao and Berman, 2026). We instead focus on how news/media referral traffic redistributes when publishers exit the pool accessible to AI platforms, and how that redistribution shapes the concentration, quality, and slant of the information environment users encounter.

Third, we contribute to the emerging literature auditing LLM search for concentration, factual quality, and slant. Closest to our work, Aral et al. (2026) probe Google AI Overviews with real user queries and document that AI search surfaces fewer long-tail sources, lower response variety, and more low-credibility, right- and center-leaning sources than traditional search. Other recent audits find inter-model agreement on credibility with a liberal slant (Yang and Menczer, 2025), fewer unique outlets than Google News when asked to retrieve the latest news (Minici et al., 2025), 37% of cited domains absent from traditional results (Zhang et al., 2025), and pronounced concentration with a left-leaning tilt (Yang, 2025). These studies all measure the cited-source set that AI platforms return in response to probe queries, typically at one or two snapshots. We instead measure the realized referral traffic that AI platforms deliver to publisher domains (SEMrush estimates) across a 48-month panel for five AI platforms (ChatGPT, Perplexity, Gemini, Claude, Grok), supporting both dynamic analysis of compositional shifts and causal identification of publisher opt-out effects.

3 Data

Our analysis combines three data sources: (1) domain-level media quality ratings from Media Bias/Fact Check (MBFC); (2) monthly website traffic from SEMrush; and (3) `robots.txt` bot-blocking records reconstructed from the Wayback Machine and Common Crawl. We describe each in turn and then detail the panel construction.

3.1 Media Bias/Fact Check (MBFC)

We obtain domain-level quality ratings from Media Bias/Fact Check (MBFC), an independent organization that evaluates news publishers on two dimensions: *political slant* (which MBFC labels *political bias*) and *factual reporting*.[2] MBFC has been widely used in prior work on media quality and platform auditing (e.g., Yang and Menczer, 2025; Minici et al., 2025; Zhang et al., 2025). For each domain, MBFC assigns:

- **Political slant**: MBFC publishes a numeric score from −10 (Extreme Left) to +10 (Extreme Right) alongside the ordinal label. Least Biased outlets score near zero.

[2] See `https://mediabiasfactcheck.com/methodology/` for the full rating methodology.

- **Factual reporting**: a numeric score from 0 (Very High) to 10 (Very Low). Lower scores indicate higher factual reliability.
- **Media type**: a categorical classification (e.g., Newspaper, Website, TV Station, Organization/Foundation, Journal).

Grouping scheme. The raw MBFC scales are fine-grained (nine slant categories, six factual categories). For the analysis, we aggregate into coarser groups that balance granularity with cell size. Table 1 summarizes the mapping.

Table 1: Variable grouping scheme for the main analysis.

Dimension	Group	Score range	Sample share
Political slant	Far Left	$[-10, -5]$	8.1%
	Left-Center	$(-5, -2]$	23.1%
	Least Biased	$(-2, +2)$	29.2%
	Right-Center	$[+2, +5)$	21.9%
	Far Right	$[+5, +10]$	17.7%
Factual reporting	Very High / High	$[0, 2)$	54.5%
	Mostly Factual	$[2, 4.5)$	9.8%
	Mixed	$[4.5, 6.5)$	25.6%
	Low / Very Low	$[6.5, 10]$	10.1%

Notes: Slant groups collapse the nine raw MBFC categories by merging the tails (e.g., Extreme Left, Far Left, and Left into "Far Left"). Factual groups collapse six raw categories into four (e.g., Very High and High into "Very High / High"). Sample shares computed over the 8,082 domains with valid MBFC ratings. Full mapping is reported in Appendix A.

Sample composition. The 8,082 rated domains span a range of media types: websites (3,368; 41.7%), newspapers (2,065; 25.6%), organizations/foundations (1,154; 14.3%), TV stations (683; 8.5%), and smaller shares of radio stations, magazines, journals, and news agencies. The sample skews slightly right: 29.2% are Least Biased, 23.1% Left-Center, 21.9% Right-Center, 17.7% Far Right, and 8.1% Far Left. Over half (54.5%) are rated Very

High/High factual quality. Appendix B reports the full distributions.

3.2 Web Traffic Data (SEMrush)

Monthly website traffic estimates come from SEMrush, a commercial digital intelligence platform that provides domain-level analytics across multiple channels.[3] For each domain $\times$ month, SEMrush reports estimated visits by channel:

- **AI platforms**: ChatGPT, Perplexity, Gemini, Claude, and Grok.
- **Channels**: Organic Search, Direct, Referral, Paid Search, Organic Social, Paid Social, Email, Display Ads, Google AI Mode, and an aggregate AI channel for AI-referred visits not attributed to a specific platform.

The panel spans January 2022 through December 2025 (48 months). Each AI platform enters the sample at its commercially relevant launch: Perplexity from December 2022, Claude from July 2023, Gemini from February 2024, ChatGPT from May 2024 (earlier data unavailable due to a SEMrush tracking gap), and Grok from January 2025.[4] Organic Search and other traditional channels are available for all 48 months.

Site size. We measure domain size as a domain's mean monthly Organic Search traffic during 2022. We partition domains into size quintiles, with Q1 denoting the smallest 20% and Q5 the largest.

3.3 Robots.txt Bot-Blocking Panel

To measure which publishers restrict AI crawlers, we construct a monthly panel of `robots.txt` directives for each domain. `Robots.txt` is a plain-text file at the root of a do-

[3]See `https://www.semrush.com/kb/1562-the-ai-traffic-dashboard` for SEMrush's AI Traffic dashboard.

[4]Perplexity launched with web search in December 2022. Claude.ai launched in July 2023. SEMrush tracks gemini.google.com only; the Bard period (March 2023–January 2024) is missing. SEMrush data for ChatGPT have a tracking gap from October 2023 through April 2024. Grok.com launched globally in January 2025.

main that instructs crawlers which content they may access. Major AI companies—including OpenAI, Anthropic, Google, and Perplexity—have publicly committed to honoring these directives.[5]

We recover historical snapshots from two archival sources—the Wayback Machine and Common Crawl—which provide timestamped `robots.txt` snapshots for a large fraction of the web. Combining the two gives the most complete coverage.

Ban definition. We define a platform-specific ban as an *access ban*: a domain is coded as banned in month t if it blocks the platform's search-indexing bot or user-facing retrieval bot. We exclude model-training bots (e.g., `gptbot`), since training restrictions capture consent for model training rather than contemporaneous content access. Full bot-to-platform mapping in Appendix C (Table 9).

3.4 Panel Construction

We merge the three sources into a domain $\times$ platform $\times$ month panel: 8,082 domains with (i) a valid MBFC rating, (ii) at least one month of nonzero SEMrush traffic, and (iii) coverage in the `robots.txt` archive. The panel has one row per (domain, platform, month) triple (8,082 domains, up to 48 months). Table 2 summarizes the key panel dimensions.

Table 2: Panel summary statistics.

Dimension	Value
Domains (MBFC-rated)	8,082
Traffic channels	15 (5 AI platforms + 10 channels)
Panel months	48 (Jan 2022 – Dec 2025)
AI platforms tracked	ChatGPT, Perplexity, Claude, Gemini, Grok
Domains ever banned $\geq$ 1 platform (Jan 2022 – Dec 2025)	3,224

[5]OpenAI: `https://developers.openai.com/api/docs/bots`; Anthropic: `https://support.claude.com/en/articles/8896518-does-anthropic-crawl-data-from-the-web-and-how-can-site-owners-block-the-crawler`; Google: `https://developers.google.com/crawling/docs/crawlers-fetchers/google-common-crawlers`; Perplexity: `https://docs.perplexity.ai/docs/resources/perplexity-crawlers`.

4 LLMs' Referral Traffic Distribution

4.1 Media Type Composition

We first illustrate how AI platforms differ from Organic Search in the types of publishers receiving referral traffic. Figure 1a plots the traffic-weighted composition of referral traffic across Organic Search and five AI platforms using 2025 data.[6]

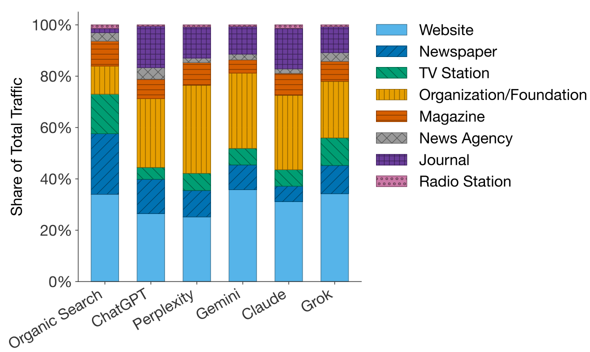

(a) Media type composition (2025).

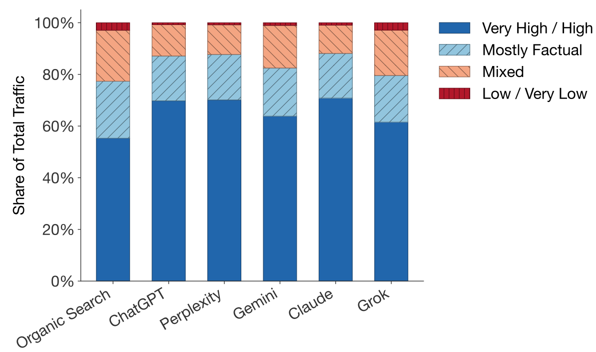

(b) Factual quality composition (2025).

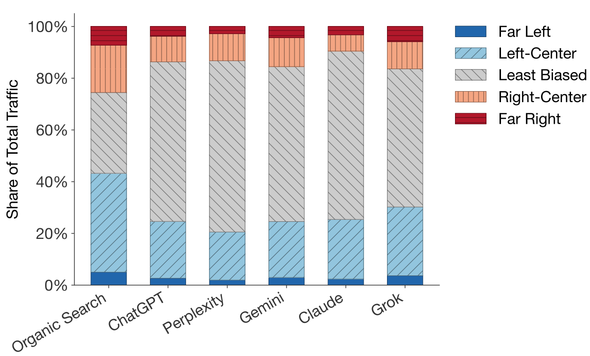

(c) Political slant composition (2025).

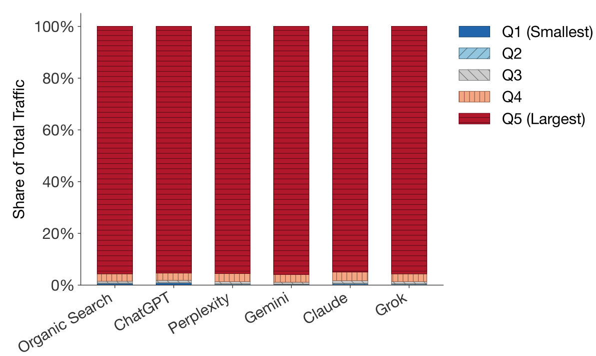

(d) Site size composition (2025).

Figure 1: Traffic-weighted composition of referral traffic by platform (2025).

Three patterns stand out. First, *newspapers* account for 23.6% of Organic Search referral traffic but only 6–13% on AI platforms. This is consistent with newspapers being among the most active `robots.txt` adopters. Second, *organizations and foundations* are over-represented on AI platforms (22–34% vs. 10.9% for Organic Search), suggesting that AI systems draw disproportionately on institutional sources. Third, *academic journals* are

[6] We use 2025 data so that all five AI platforms can be observed; Grok, the most recent entrant, launched in January 2025.

over-represented on AI platforms (10–16% vs. 1.7%). This reflects a tendency for LLMs to cite peer-reviewed literature that traditional search engines rarely surface. Generic *websites* remain broadly stable across platforms (25–36%).

Figures 1b and 1c confirm that these structural differences translate into factual quality and slant differentials. All AI platforms skew toward higher quality: Very High/High factual domains account for 61.5% (Grok) to 70.9% (Claude) of AI traffic, compared with 55.3% for Organic Search. AI platforms are more centrist: Center-rated content accounts for 53–66% of AI traffic versus 31.2% on Organic Search.

Figure 1d shows the size distribution. Across all platforms, Q5 (largest) sites account for roughly 95% of traffic, with the bottom four quintiles contributing negligibly.

4.2 How LLMs Differ from Traditional Search

We next compare the overall mix of referral traffic sent by AI platforms with that sent by Organic Search. The goal is to summarize, in a common set of indices, whether AI-referred traffic is more concentrated, more or less factual, or politically different from traditional search traffic. This motivates the construction of the indices below, which use each domain's share of a platform's referral traffic as its weight.

4.2.1 Index Construction

To summarize and compare these differences across platforms, we construct three aggregate indices per platform p. We focus on the 2025 calendar year, when all five AI platforms are observed.

Let $w_{d,p} = \text{traffic}_{d,p} / \sum_{d'} \text{traffic}_{d',p}$ denote domain d's 2025 referral traffic share from platform p. Two of the three indices weight domain-level MBFC scores by these shares: factual reporting $fr_d \in [0, 10]$ (lower = more factual) and political slant $\text{slant}_d \in [-10, +10]$ (negative = left-leaning). The third index, the Herfindahl–Hirschman Index, is computed directly from the shares $w_{d,p}$; it summarizes how unequally referrals are distributed across source domains:

- **Factual Index**: $I_p^{\mathrm{FR}} = \sum_d w_{d,p} \cdot fr_d$. Lower = higher-quality pool.
- **Slant Index**: $I_p^{\mathrm{Slant}} = \sum_d w_{d,p} \cdot \mathrm{slant}_d$. Negative = left-leaning.
- **HHI**: $I_p^{\mathrm{HHI}} = 10{,}000 \times \sum_d w_{d,p}^2$. Higher = more concentrated.

4.2.2 Cross-Platform Comparison

Figure 2 plots each platform's position across factual–slant and factual–HHI planes using the indices described above. The top-right corner represents the most desirable outcome, i.e, more factual and more centrist for 2a and more factual and less concentrated for Figure 2b.

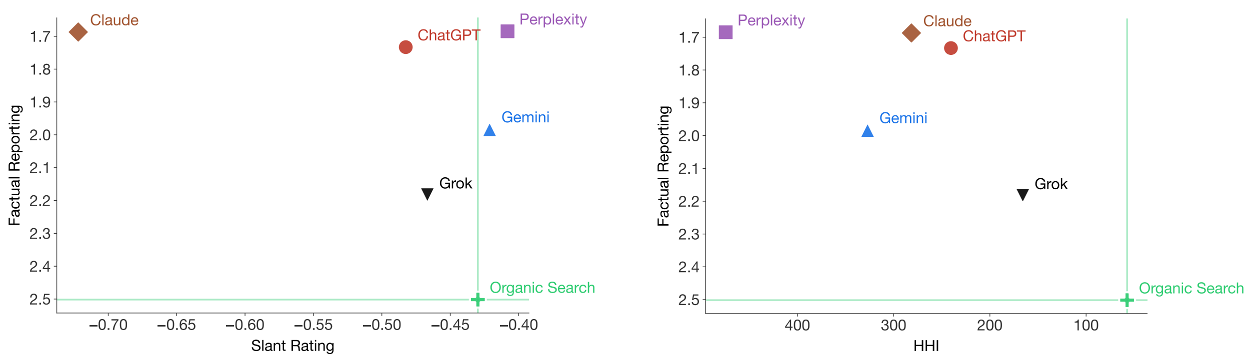

(a) Factual reporting vs. political slant.

(b) Factual reporting vs. concentration .

Figure 2: Platform positioning across content quality dimensions (2025 aggregate).

The AI platforms differ from Organic Search on every dimension. AI referral traffic is more factual (Factual Index 1.68–2.18 vs. 2.50 for Organic Search, a 13–33% improvement) and far more concentrated (HHI 166–475 vs. 57, roughly 3–8×). On political slant, most AI platforms are close to Organic Search ($I^{\mathrm{Slant}} = -0.43$): Perplexity ($-0.41$) and Gemini ($-0.42$) are slightly more centrist, ChatGPT (-0.48) and Grok (-0.47) marginally more left-leaning. Claude is the outlier (-0.72), with a stronger left-leaning tilt.

The scatter plots in Figure 2 make these tradeoffs visible. Figure 2a shows the quality–slant frontier, where Perplexity combines the strongest factual quality with near-neutral political slant and Gemini preserves a similarly centrist orientation at lower factual quality.

In Figure 2b, the tradeoff shifts to quality versus diversity: Perplexity and Grok anchor the two endpoints, while ChatGPT and Claude occupy intermediate positions. Gemini falls behind ChatGPT and Claude on this two-dimensional plane, but its centrist slant prevents it from being dominated in the overall comparison. Table 3 summarizes these tradeoffs in scorecard form.

Table 3: Relative platform positions across concentration, factual quality, and slant. No platform dominates across all three dimensions.

Platform	**Concentration** ($I^{\text{HHI}} \downarrow$)	**Factual** ($I^{\text{FR}} \downarrow$)	**Slant** ($I^{\text{Slant}} \uparrow$)
Perplexity	✗	✓✓	✓✓
ChatGPT	✓	✓✓	✓
Gemini	✓	✓	✓✓
Claude	✓	✓✓	✗
Grok	✓✓	✗	✓

Notes: ✓✓ = within top 10% of cross-platform range; ✗ = within bottom 10%; ✓ = middle.

Table 3 and Figure 2 indicate that no platform dominates the others across all metrics. Perplexity sits at one corner of the frontier, with the best factual quality and most centrist slant, but the highest concentration. Grok sits at the opposite corner, distributing traffic most broadly but with the poorest factual quality. ChatGPT and Claude occupy the middle of the quality–diversity tradeoff: both are high-quality relative to the other AI platforms, but ChatGPT is less concentrated, while Claude has a stronger left-leaning tilt. Gemini is weaker on the concentration–factual plane, but it remains on the full three-dimensional frontier because its slant is among the most centrist.

5 LLM Content Quality Over Time

Section 4 established that AI platforms differ from traditional search in the static composition of their content pools. In the next set of analyses, we examine the dynamic consequences of publisher opt-outs. We proceed in two steps. We first characterize which publishers have blocked LLM crawlers (Section 5.1). We then show how those exits shift aggregate quality indices over time (Section 5.2).

Figure 3 documents the scale and timing of publisher exit. ChatGPT and Perplexity are the most-banned platforms once Gemini is set aside (2,613 and 2,331 by December 2025); Claude and Grok remain far below (484 and 271). Gemini's nominal total (2,871) is the largest of all, but is driven almost entirely by `google-extended`: without it, Gemini's count falls to 412.[7] Publishers began blocking ChatGPT and Gemini in mid-2023, and Perplexity in mid-2024. Claude and Grok have seen few bans throughout.

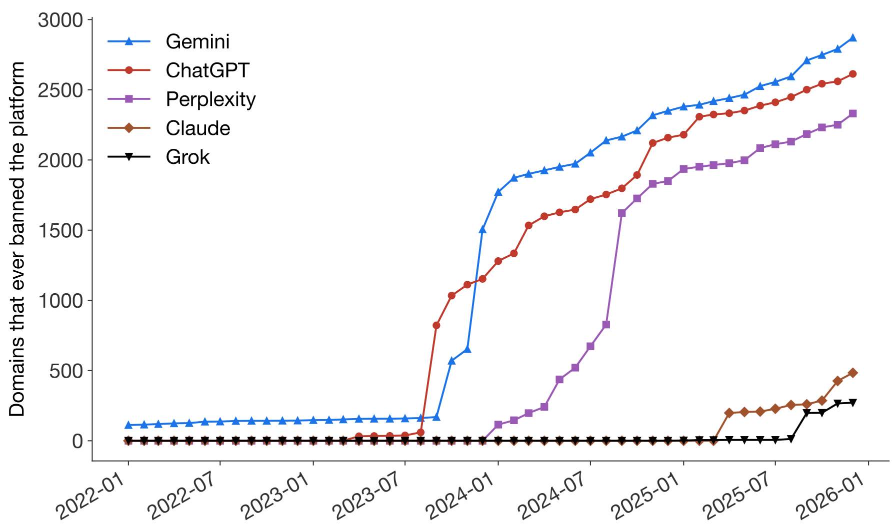

Figure 3: Number of domains that have ever banned each platform's crawlers, by month.

[7]Two Gemini-specific inclusions in our access-ban set: `googlebot` (Gemini's grounded retrieval uses the Google Search index) and `google-extended` (Google documents it as governing both training and grounding).

5.1 Who Has Left the LLM-Indexable Pool?

The composition of exiting domains determines whether the residual pool becomes systematically different in terms of media type, political slant, factual reporting, and size. We classify each of the 8,082 domains as banned or unbanned for each platform as of December 2025 and characterize the ban landscape along four dimensions, summarized in Figure 4.

Banning is highly selective along all four dimensions. By media type, newspapers account for 47–96% of banned domains across platforms, far exceeding their 26% share of the sample. By political slant, Least Biased and Left-Center outlets comprise 75–92% of bans. By factual quality, Very High/High-rated outlets account for 85–97% of bans. By site size, mid-to-large publishers (Q3–Q5) account for 75–95% of bans, while Q1 (smallest) sites account for just 0–8%.

In short, the publishers exiting are disproportionately newspapers, centrist-to-left outlets, and the most factually reliable tier. What remains accessible is skewed toward small websites, politically peripheral sources, and lower-factual quality publishers.

5.2 How Do Bans Change the Content Pool?

We now ask whether these selective exits are large enough to move aggregate content quality. We extend the three indices defined in Section 4.2.1 to a monthly panel. Let $\mathcal{U}_{p,t}$ denote the set of domains that have not blocked platform p's crawler in month t,[8] and let a_d denote domain d's 2022 Organic Search traffic. We weight each unbanned domain by its fixed 2022

[8]The platform-specific ban indicator is cleaned in three steps. First, missing values are forward-filled within a domain, carrying the most recent observed ban status forward. Second, any value still missing before the platform's first observed ban month is set to zero, since no domain blocked the platform before then (March 2023 for ChatGPT, January 2024 for Perplexity, January 2025 for Grok, and April 2025 for Claude; for Gemini we use the September 2023 launch of `google-extended`). Third, any domain still missing a ban value in or after that month is dropped, because its status during the ban period cannot be reliably determined. This leaves per-platform pools of 7,537 (ChatGPT), 7,736 (Perplexity), 7,671 (Gemini), 7,939 (Claude), and 7,903 (Grok) of the 8,082 domains.

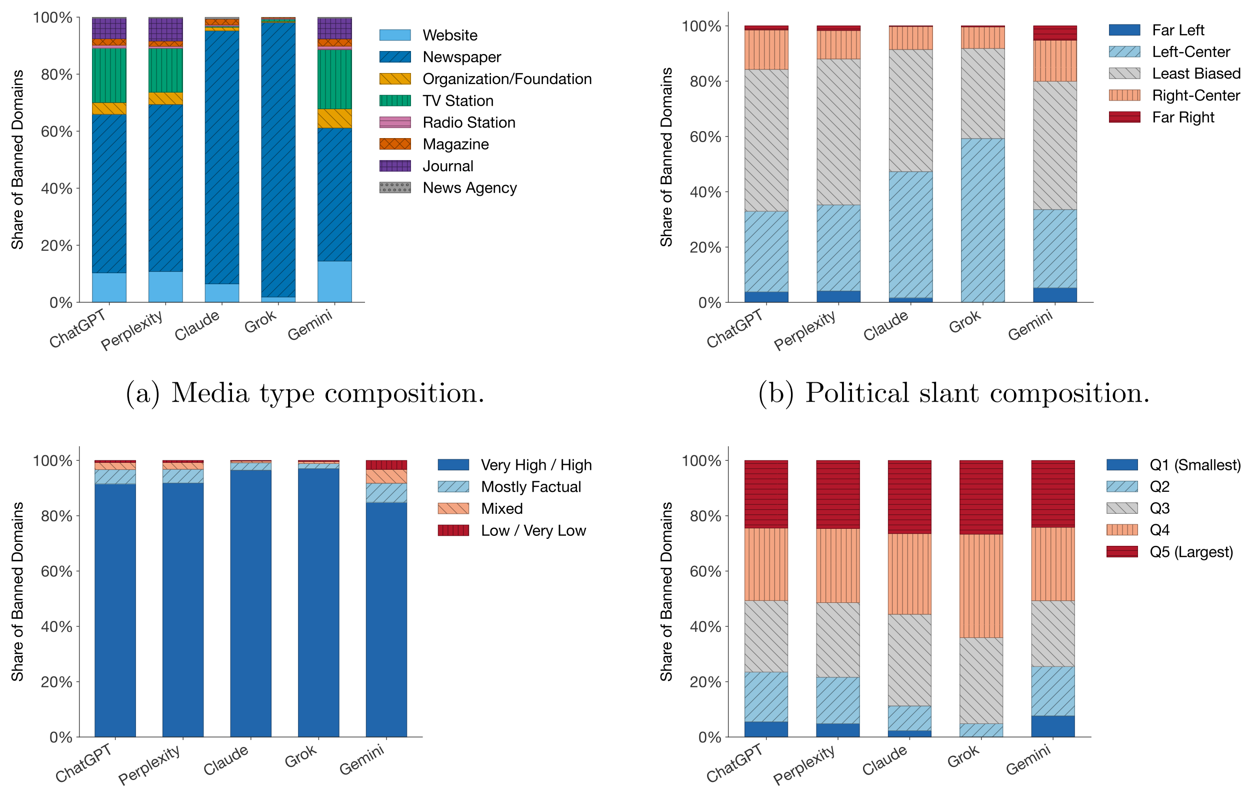

(a) Media type composition.

(b) Political slant composition.

(c) Factual reporting composition.

(d) Site size composition.

Figure 4: Composition of platform-banned domains (December 2025). Each bar shows, within one platform's banned domain set, the share that falls into each category along the indicated dimension.

Organic Search traffic share within platform p's available pool:

$$\widehat{\omega}_{d,p,t} = \frac{a_d}{\sum_{j \in \mathcal{U}_{p,t}} a_j}, \qquad d \in \mathcal{U}_{p,t}. \tag{1}$$

The three indices for platform p in month t are then $I^{\mathrm{FR}}_{p,t} = \sum_{d \in \mathcal{U}_{p,t}} \widehat{\omega}_{d,p,t}\, fr_d$ and $I^{\mathrm{Slant}}_{p,t} = \sum_{d \in \mathcal{U}_{p,t}} \widehat{\omega}_{d,p,t}\, \mathrm{slant}_d$, with HHI defined analogously as $I^{\mathrm{HHI}}_{p,t} = 10{,}000 \times \sum_{d \in \mathcal{U}_{p,t}} (\widehat{\omega}_{d,p,t})^2$. These indices capture changes in the available pool of domains for each platform, rather than contemporaneous changes in platform ranking or user demand. Figures 5–7 plot the three indices from each platform's entry date through December 2025.

Three findings emerge. First, the Factual Index trends upward for ChatGPT and Perplexity through December 2025, indicating a shift toward lower factual quality as high-quality

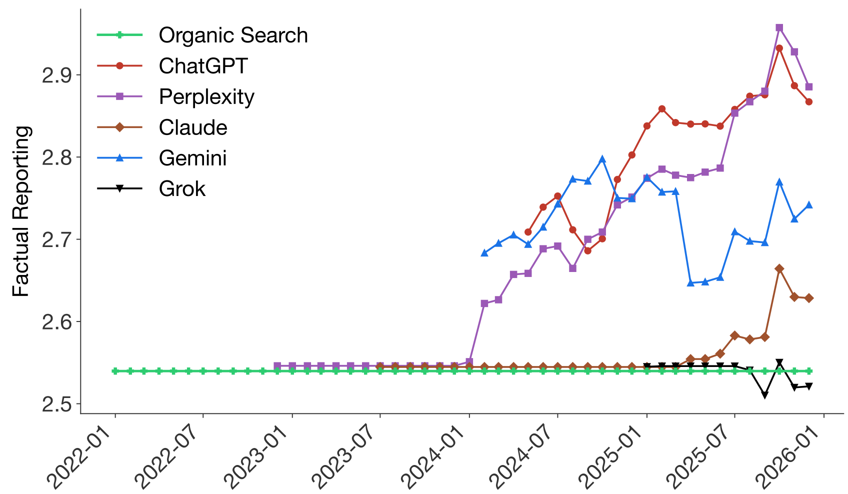

Figure 5: Factual Index over time (higher score means less factual), computed over each platform's unbanned domain pool.

publishers exit. Gemini's index has been largely flat. Claude shows a smaller, later rise (starting in April 2025). Grok remains flat. Second, the Slant Index for ChatGPT and Perplexity drifts rightward, consistent with the disproportionate exit of centrist and left-leaning publishers; the shift is most pronounced for ChatGPT. Gemini's Slant Index remains roughly flat. Claude retains these publishers and remains more left-leaning. Third, pool HHI rises for ChatGPT and Gemini as large institutional domains exit. Perplexity's pool HHI rises less, yet Perplexity has the highest traffic HHI in the cross-section (Section 4.2.2). This suggests that Perplexity's concentration is driven by algorithm- or demand-side selection, not by the supply of its crawlable pool.

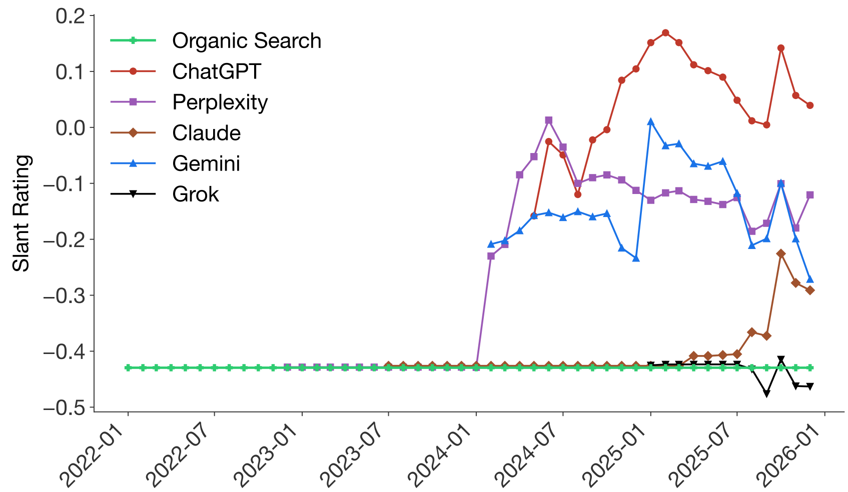

Figure 6: Slant Index over time, computed over each platform's unbanned domain pool.

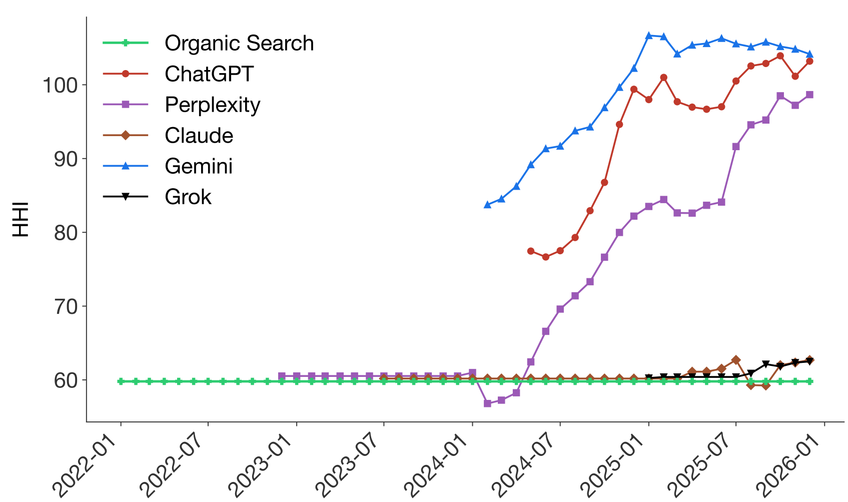

Figure 7: HHI Concentration Index over time, computed over each platform's unbanned domain pool.

6 Quality Composition and Traffic Redistribution

Sections 4 and 5 characterized the pool of domains that each platform can still reach and showed that selective exit degrades this pool. That analysis describes the available supply, not how platforms route traffic across the remaining set of publishers, and the two need not move to the same extent: Perplexity's available pool grows only modestly more concentrated (Section 5.2), yet its realized traffic in 2025 is the most concentrated of any platform (Section 4.2.2).

Figure 8 traces how the quality composition of Perplexity's realized referral traffic changes month by month. Recomputing the three indices of Section 4.2.1 month by month with realized referral shares as weights, $w_{d,p,t} = \text{traffic}_{d,p,t} / \sum_j \text{traffic}_{j,p,t}$,[9] rather than the fixed 2022 organic search weights used in Section 5.2, three trends stand out. Perplexity's realized traffic grows steadily *more concentrated*—HHI climbs from ≈260 at entry to ≈590 by the end of 2025, an order of magnitude above Organic Search and far beyond what its pool alone implies (Figure 7); *less factual*—the Factual Index drifts upward only modestly despite the pool's deterioration (Section 5.2), staying well below (i.e., more factual than) Organic Search throughout; and *more right-leaning*—Slant drifts more noticeably toward Organic Search, tracking the pool's loss of centrist and left-leaning publishers more closely. Together, these patterns hint at how Perplexity redistributes traffic across the remaining accessible publishers: concentration rises far beyond the modest increase in the accessible pool's HHI (Figure 7), factual quality is partially buffered, and slant tracks the pool's loss of centrist publishers more closely.

To understand these redistribution patterns causally, we ask: when publishers exit the accessible pool, how does traffic redistribute among remaining accessible publishers? We exploit the staggered timing of publisher bans on Perplexity, the AI platform for which our data offer the richest combination of ban events and domain-level referral traffic. The

[9] To avoid constructing the indices from too few sources, we compute them for months with more than 200 *active* domains, i.e., domains receiving at least one referral visit from Perplexity that month.

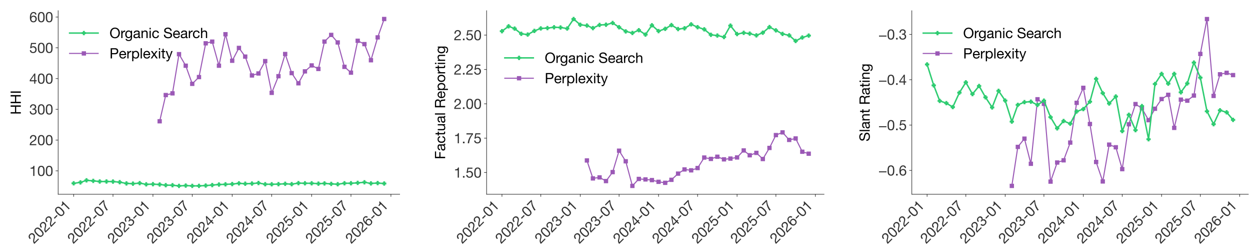

(a) HHI (b) Factual Index (c) Slant Index

Figure 8: Realized traffic-weighted indices over time for Perplexity and Organic Search (for the Factual index, higher score means less factual; for the Slant index, negative score means left-leaning).

analysis proceeds in two steps: Section 6.1 establishes that a ban actually cuts the banning publisher's own referral traffic, and Section 6.2 asks where the referral traffic goes, i.e., whether it flows to remaining accessible publishers that resemble the blocking publishers in size, factual reliability, or political slant.

6.1 Validating Bans as an Effective Treatment

We estimate an event study on the focal domain's own first ban and report average treatment effects on the treated (ATTs) under three estimators: (1) Two-Way Fixed Effect (TWFE), (2) Callaway and Sant'Anna (2021), and (3) Sun and Abraham (2021), the latter two robust to treatment-effect heterogeneity across cohorts and event time, with Callaway–Sant'Anna using not-yet-treated and never-treated domains as the control group and Sun–Abraham using never-treated domains as the control group.

The standard TWFE model takes the following form:

$$\log(1 + \text{traffic}_{d,t}) \;=\; \sum_{k \neq -1} \delta_k \, \mathbf{1}\{\text{event\_time}_{d,t} = k\} + \alpha_d + \lambda_t + \varepsilon_{d,t}, \tag{2}$$

where t_d^* is the first month in which domain d blocks Perplexity's crawler, $\text{event\_time}_{d,t} = t - t_d^*$, $k = -1$ (i.e., the month prior to each domain ban) is the omitted reference period, and α_d and λ_t are domain and year-month fixed effects, respectively. For domains that later remove the ban, we drop observations after the first unban month so that, in the estimation sample, treatment switches on once and remains on thereafter. Standard errors are clustered at the domain level.[10] Figure 9 plots the event study estimates.

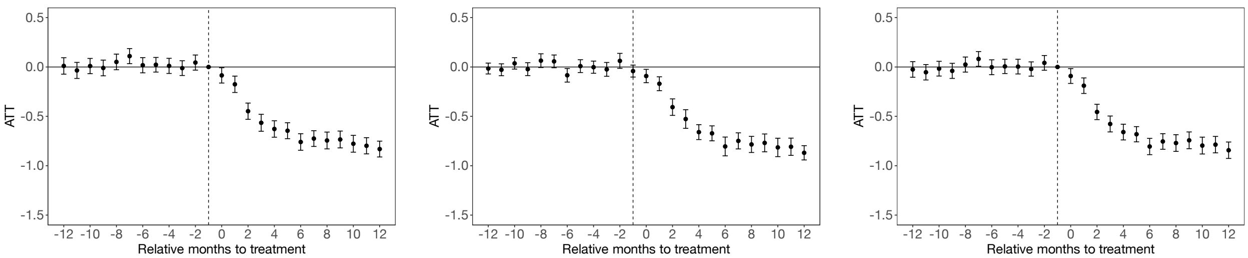

(a) TWFE (b) Callaway-Sant'Anna (c) Sun-Abraham

Figure 9: Own-ban event study of Perplexity referral traffic. Each panel plots period-by-period ATT with 95% CIs. Each model includes domain and year-month fixed effects. Standard errors are clustered at the domain level (by bootstrap ($B = 200$) in (b)).

The event study coefficients are close to zero in the pre-ban months and show no systematic trend, providing evidence of no pre-trends. The overall ATTs range between -0.61 and -0.65 ($p < 0.001$), implying a 46–48% reduction in the banning domain's referral traffic from Perplexity. Table 4 reports overall ATT estimates and confidence intervals.

Table 4: Own-ban effect on Perplexity referral traffic.

Estimator	Estimate	95% CI
TWFE (overall ATT)[11]	-0.61^{***}	$[-0.68, -0.54]$
Callaway-Sant'Anna (overall ATT)	-0.65^{***}	$[-0.71, -0.58]$
Sun-Abraham (overall ATT)	-0.64^{***}	$[-0.70, -0.57]$

Notes: $N = 284{,}732$ domain–months; 2,309 ever-banned, 5,427 never-banned. Standard errors clustered at the domain level are reported in parentheses.
Significance levels: $^{*}\, p < 0.10$; $^{**}\, p < 0.05$; $^{***}\, p < 0.01$.

[10] Perplexity ban indicator is cleaned in three steps. First, missing values are forward-filled within a domain, carrying the most recent observed ban status forward. Second, any value still missing before January 2024 is set to zero, since no domain was banned by Perplexity before then. Third, any domain still missing a ban value in or after January 2024 is dropped entirely, because its status during the ban period cannot be reliably determined. The result is a panel of 7,736 domains observed over a 37-month window (284,732 domain–months; 2,309 ever-banned, 5,427 never-banned).

[11] Equal-weight average of $\hat{\delta}_0, \ldots, \hat{\delta}_{12}$; Standard errors computed via the full cluster-robust covariance matrix.

6.2 Substitution Among Non-Blocking Domains

When a publisher bans Perplexity, the referral it would have received should be reassigned to one of the remaining accessible publishers. We ask whether these referrals flow to publishers of the same type as the blocking publisher. We sort publishers into bins along each of three dimensions we study, a bin being a group of similar publishers, such as a traffic-size quintile, a factual-reporting category, or a political-slant category.[12] For each focal domain, we then measure its *peer-ban exposure*: the share of its same-bin peers—the other publishers in its own bin—that have already banned Perplexity. We then ask how the focal domain's own referral traffic moves with this exposure. A positive coefficient means that as more of a domain's same-type peers exit, the domain gains traffic, i.e., Perplexity refers to remaining accessible publishers of the same type. We compute peer-ban exposure leave-one-out, i.e., excluding the focal domain: for each focal d, month t, and dimension m, it is the 2022-traffic-weighted share of d's same-bin peers that have banned Perplexity by t,

$$E^{(m)}_{-d,t} \;=\; \frac{\sum_{j\neq d:\, g^{(m)}_j = g^{(m)}_d} a_j\, B_{j,t}}{\sum_{j\neq d:\, g^{(m)}_j = g^{(m)}_d} a_j}, \qquad E^{(m)}_{-d,t} \in [0,1], \tag{3}$$

where a_j is domain j's 2022 Organic Search traffic, $B_{j,t} = \mathbf{1}\{j \text{ has banned Perplexity by } t\}$ is absorbing, and $g^{(m)}_j$ is domain j's bin under dimension m. The grouping dimensions are $m \in \{\text{Size, FR, Slant}\}$: 2022 traffic-size quintiles for Size, factual-reporting bins for Factual Reporting, and slant bins for Slant.

The estimating equation pools all three peer-exposure terms so that each coefficient is identified after partialling out exposure on the other two dimensions:

$$\begin{aligned} \log(1+\text{traffic}_{d,t}) = \gamma^{\text{Size}}\, E^{(\text{Size})}_{-d,t} + \gamma^{\text{FR}}\, E^{(\text{FR})}_{-d,t} + \gamma^{\text{Slant}}\, E^{(\text{Slant})}_{-d,t} \\ + \sum_{k\neq -1} \delta_k\, \mathbf{1}\{\text{event\_time}_{d,t} = k\} + \alpha_d + \lambda_t + \varepsilon_{d,t}, \end{aligned} \tag{4}$$

[12] For factual reporting and political slant, we use the categorization provided by MBFC and reported in Table 1.

with α_d domain fixed effects, λ_t year-month fixed effects, and standard errors clustered by domain. We estimate equation (4) using TWFE and Sun and Abraham (2021).[13]

Two features of this design address the main identification threats. First, a focal domain's own ban directly reduces its referral traffic and is correlated with peers' ban status. To separate this direct own-ban effect from peer substitution, equation (4) includes the full set of own-ban event-time dummies δ_k. The peer-exposure coefficients $\gamma^{(m)}$ therefore capture traffic changes associated with peers' bans, net of the focal domain's own ban effect. Second, the exposure must not depend on the focal domain's own ban status, which is why equation (3) is computed leave-one-out, excluding $j = d$ from both sums.

The own-ban event-time dummies δ_k in equation (4) have pre-ban leads ($k < 0$) that cluster near zero under both TWFE and Sun–Abraham, consistent with parallel trends (Appendix D). The implied overall own-ban ATTs from these dummies are reported in Appendix E (Table 10). The two estimators are nearly identical, so we focus our discussion on the Sun–Abraham column.

Table 5: Effect of peer-ban exposure on log Perplexity referral traffic.

	TWFE		Sun–Abraham	
Exposure dimension	$\hat{\gamma}$	(SE)	$\hat{\gamma}$	(SE)
Size ($E^{(\text{Size})}$)	$+4.24^{***}$	(0.11)	$+4.16^{***}$	(0.12)
Factual ($E^{(\text{FR})}$)	$+1.42^{***}$	(0.15)	$+1.58^{***}$	(0.15)
Slant ($E^{(\text{Slant})}$)	$+0.46^{***}$	(0.13)	$+0.39^{***}$	(0.14)

Notes: N = 284,732 domain–months. Standard errors clustered at the domain level are reported in parentheses.
Significance levels: $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

Size dimension. Size is the dominant substitution channel: when a focal domain's same-size peers exit Perplexity, referral traffic flows back to it at high rates. A 10-percentage-point rise in same-size-quintile peer-ban exposure raises the focal domain's referral traffic

[13] Callaway and Sant'Anna (2021) does not allow for time-varying controls.

by $\approx 52\%$ ($\hat{\gamma}^{\text{Size}} = 4.16$, conditional on same-factual and same-slant exposure). On average, referral traffic returns to the blocking publisher's own size tier rather than crossing the size hierarchy—a within-tier pull that, as we show below, comes almost entirely from the largest publishers. A direct implication is that AI referrals become more concentrated within each tier: since banning rates rise with publisher size, the referral traffic accumulates among the few large publishers that have not opted out, raising HHI overall and even within each size tier.

Factual dimension. Factual reliability is the next channel, roughly one-third the strength of size: when peers of similar factual quality exit, referral traffic flows to remaining accessible domains of comparable quality. A 10-percentage-point rise in same-factual-bin peer-ban exposure lifts the focal domain's referral traffic by $\approx 17\%$ ($\hat{\gamma}^{\text{FR}} = 1.58$, conditional on same-size and same-slant exposure). This same-quality matching is, however, much weaker than for size and reverses among the low-factual outlets, as we show below.

Slant dimension. Political slant is the weakest substitution channel: when same-slant peers exit Perplexity, remaining accessible domains in that slant bin gain little traffic. A 10-percentage-point rise in same-slant-bin peer-ban exposure lifts the focal domain's referral traffic by only $\approx 4\%$ ($\hat{\gamma}^{\text{Slant}} = 0.39$)—roughly an order of magnitude weaker than the size channel. Ideological matching thus plays only a marginal role in Perplexity's reallocation of referrals when peers withdraw.

Hierarchy of substitution channels. The estimates show a clear ordering across channels, with size being the strongest, factual quality roughly one-third as strong, and political slant weakest. Referral gains, therefore, accrue mainly to the remaining accessible publishers in the same size or factual-quality tier as the blocking publisher, rather than dispersing toward smaller or lower-quality publishers. Because larger publishers ban more often, the set of accessible large publishers shrinks over time; within-tier substitution then concentrates referral gains among the large publishers that remain accessible, rather than passing traffic down to smaller publishers. The factual channel works the same way, concentrating referral gains among high-quality, remaining-accessible publishers. Overall, Perplexity's reallocation reinforces existing advantages of large, high-quality publishers rather than spreading referrals more evenly.

6.3 Heterogeneity Within Each Tier

The estimates above impose a single substitution parameter for each dimension. For example, the size coefficient averages over focal domains in all five size quintiles, so it does not tell us whether same-size substitution is stronger among large publishers than among small publishers. We now relax this restriction by allowing the peer-exposure parameter to vary with the focal domain's own bin. Specifically, we modify equation (4) by interacting each peer-exposure term with indicators for the focal domain's group within that dimension:

$$\begin{aligned}\log(1+\text{traffic}_{d,t}) = \gamma^{\text{Size}}\, E^{(\text{Size})}_{-d,t} &+ \sum_{g\neq g_0^{\text{Size}}} \gamma_g^{\text{Size}}\, \mathbf{1}\{g_d^{\text{Size}} = g\}\, E^{(\text{Size})}_{-d,t} \\ &+ \gamma^{\text{FR}}\, E^{(\text{FR})}_{-d,t} + \sum_{g\neq g_0^{\text{FR}}} \gamma_g^{\text{FR}}\, \mathbf{1}\{g_d^{\text{FR}} = g\}\, E^{(\text{FR})}_{-d,t} \\ &+ \gamma^{\text{Slant}}\, E^{(\text{Slant})}_{-d,t} + \sum_{g\neq g_0^{\text{Slant}}} \gamma_g^{\text{Slant}}\, \mathbf{1}\{g_d^{\text{Slant}} = g\}\, E^{(\text{Slant})}_{-d,t} \\ &+ \sum_{k\neq -1} \delta_k\, \mathbf{1}\{\text{event\_time}_{d,t} = k\} + \alpha_d + \lambda_t + \varepsilon_{d,t},\end{aligned} \tag{5}$$

where $\gamma^{(m)}$ is the coefficient for the reference bin $g_0^{(m)}$, and $\gamma_g^{(m)}$ the differential slope of bin g relative to the reference bin—so the bin-specific coefficent is $\gamma^{(m)}$ for the reference bin and $\gamma^{(m)} + \gamma_g^{(m)}$ otherwise. We use the top category in each dimension as the reference bin—Q5, the largest size quintile for Size; Very High / High for factual quality; and Least Biased for political slant. The interaction coefficients for the remaining bins are interpreted relative to these reference categories. Table 6 displays the estimates from Equation 5.

Table 6: Per-bin peer-exposure effects from equation (5).

		TWFE		Sun–Abraham	
Dimension	Parameter	Estimate	(SE)	Estimate	(SE)
Size	γ^{Size} (ref.)	$+7.28^{***}$	(0.18)	$+6.99^{***}$	(0.18)
	$\gamma^{\text{Size}}_{\text{Q4}}$	-3.78^{***}	(0.17)	-3.75^{***}	(0.17)
	$\gamma^{\text{Size}}_{\text{Q3}}$	-5.71^{***}	(0.15)	-5.73^{***}	(0.16)
	$\gamma^{\text{Size}}_{\text{Q2}}$	-6.69^{***}	(0.16)	-6.89^{***}	(0.16)
	$\gamma^{\text{Size}}_{\text{Q1}}$	-6.45^{***}	(0.76)	-8.49^{***}	(0.76)
Factual	γ^{FR} (ref.)	$+0.83^{***}$	(0.14)	$+0.77^{***}$	(0.14)
	$\gamma^{\text{FR}}_{\text{MF}}$	$+0.78^{***}$	(0.16)	$+0.73^{***}$	(0.16)
	$\gamma^{\text{FR}}_{\text{Mix}}$	$+0.17$	(0.24)	$+0.08$	(0.24)
	$\gamma^{\text{FR}}_{\text{Low}}$	-1.18^{***}	(0.18)	-1.26^{***}	(0.18)
Slant	γ^{Slant} (ref.)	$+0.27$	(0.18)	$+0.36^{*}$	(0.20)
	$\gamma^{\text{Slant}}_{\text{LC}}$	$+0.04$	(0.14)	-0.08	(0.16)
	$\gamma^{\text{Slant}}_{\text{RC}}$	-0.28	(0.19)	-0.50^{**}	(0.20)
	$\gamma^{\text{Slant}}_{\text{FL}}$	-0.31	(0.22)	-0.63^{***}	(0.24)
	$\gamma^{\text{Slant}}_{\text{FR}}$	-0.74^{***}	(0.21)	-0.91^{***}	(0.22)

Notes: Each row reports the effect of same-bin peer-ban exposure for focal domains in that bin. The reference parameter $\gamma^{(m)}$ is the bin-specific peer-exposure effect for the reference bin; the subscripted $\gamma^{(m)}_g$ are effects relative to the reference (the bin-specific effect equals $\gamma^{(m)} + \gamma^{(m)}_g$). Reference bins are Q5 (Size), Very High / High (Factual), and Least Biased (Slant). Subscripts index the focal bin: Q1–Q4 (Size); MF = Mostly Factual, Mix = Mixed, Low = Low / Very Low (Factual); LC = Left-Center, RC = Right-Center, FL = Far Left, FR = Far Right (Slant). N = 284,732 domain–months. Standard errors clustered at the domain level are reported in parentheses
Significance levels: $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

On the size dimension, the bin-specific coefficients rise with publisher size. They are the strongest for the largest publishers (Q5), and decline monotonically across tiers, reaching near zero by Q2 and turning negative for the smallest (Q1) under Sun–Abraham, leaving little to no within-tier substitution for small domains. Because large publishers most often ban Perplexity, this top-heavy substitution concentrates referral gains among an ever-smaller set of large remaining accessible publishers, raising concentration within the top tier.

On the factual dimension, substitution is positive at every tier except the low-factual domains, where it turns sharply negative. When a low-factual publisher is banned, referral traffic that would have flowed through it does not return to other remaining accessible low-factual publishers; instead, it flows to higher-quality domains. Therefore, banning a low-factual-quality source improves the factual quality of referrals rather than recirculating its traffic within the low-quality tier.

On the slant dimension, substitution is positive only at the center—weakly for Least-Biased and Left-Center domains—then near zero for Right-Center and negative for both extremes, most strongly for Far-Right. Therefore, ideological matching steers reallocation only weakly: when a partisan outlet bans, remaining outlets in the same slant tier do not systematically gain more traffic. Therefore, unlike the size dimension, slant substitution is not self-reinforcing.

Together, the per-bin coefficients show that referral gains are not evenly distributed across bins but instead concentrate at the top of the size and quality dimensions (the largest publishers and the higher-factual-quality outlets), while political slant steers them only weakly.

6.4 Placebo Test: Randomized Publisher Quality Ratings

We conclude with a placebo test that checks whether the peer-exposure effects we discussed and presented in Table 5 are indeed driven by publisher bans and the actual assignment of publishers to size, factual-quality, and slant bins, rather than by spurious correlations between ban timing and traffic outcomes. In each of 100 placebo draws, we randomly permute the domains and assign each domain the size, factual quality, and slant attributes of another domain in the permutation. The domain's own traffic outcomes and ban timing are left unchanged; only the attributes used to construct the peer-exposure measures are reassigned. We then recompute the peer-exposure measures using these placebo attributes and re-estimate equation (4). This preserves the marginal and joint distribution of publisher attributes while breaking the link between a domain's actual attributes and the peer-exposure measures. If the effects in Table 5 reflect substitution among similar publishers, the placebo estimates should be centered near zero. If instead they are driven by spurious correlations in the ban schedule or traffic outcomes, the placebo attributes should produce effects similar to the real estimates.

Figure 10 plots the resulting placebo distribution. First, the distribution of placebos is centered around zero. Second, the real Sun–Abraham estimate presented in Table 5 (red dashed line) lies in the tail of the placebo distribution for all three attributes: Size and Factual ($p < 0.01$) and Slant ($p = 0.07$). These results confirm that our findings are indeed due to the peer-exposure measure we defined and the types of domains that are banned.

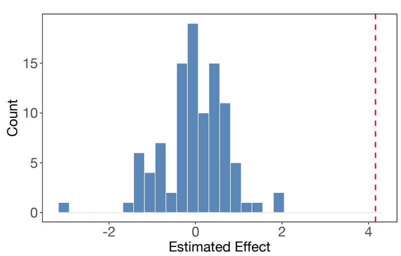

(a) Size

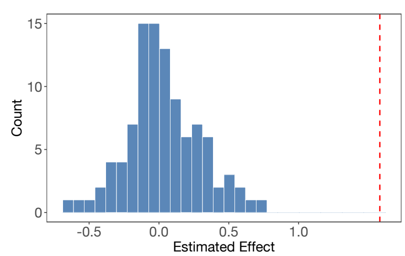

(b) Factual

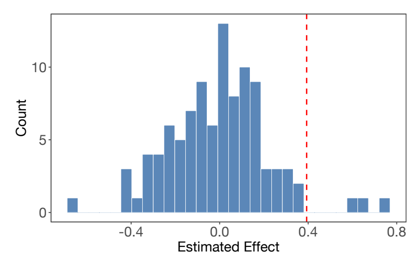

(c) Slant

Figure 10: Placebo distribution of each peer-exposure slope across 100 draws that reassign all three attributes under a single shared donor permutation. Red dashed line marks the correct Sun–Abraham estimate from Table 5.

7 Conclusion

We study LLM-based search platforms as information intermediaries focusing on the news and media industry. We document how these platforms differ from organic search in source concentration, factual quality, and political slant, and how publisher opt-outs reshape these properties over time. Using a 48-month panel of 8,082 news domains across five AI platforms, we find that AI referral traffic is more concentrated than organic search but of higher factual quality and comparable political slant on average; that publishers blocking AI crawlers via `robots.txt` are disproportionately the largest, most factually reliable, and centrist-to-left outlets; and that AI platforms reallocate referral traffic back to the largest and most factually reliable remaining accessible outlets.

We find that this reallocation is dominated by domain size, followed by factual quality, with political slant showing negligible reallocation. Because size dominates redistribution, AI platforms' referral traffic flows primarily back to the largest remaining accessible publishers, leaving smaller high-quality publishers with little to gain in terms of referral traffic when their peers exit. The redistribution itself, therefore, accelerates the concentration trend of referral traffic we observed rather than reducing it. Redistribution does not, however, undo the compositional shifts that exit produces. Although referrals reallocate toward the remaining most reliable publishers, the pool continues to lose its most reliable members, so its average factual quality still declines. And because the slant channel is roughly an order of magnitude weaker than the size channel, redistribution does little to counter the rightward tilt generated by the exit of centrist and left-leaning outlets.

AI platforms are quickly becoming a primary interface between users and the open web, and what they surface will shape both which publishers are likely to benefit from these platforms and which sources users are actually referred to. The patterns we document point to a narrowing of the accessible source pool that is unlikely to reverse on its own. Whether the resulting information environment is welfare-improving depends on whose voices users miss and whose remain. Our estimates quantify how that mix is being reshaped, and provide a baseline against which licensing arrangements, revenue-sharing frameworks, or changes to retrieval design can ultimately be evaluated.

References

Aral, S., Li, H., and Zuo, R. (2026). The rise of AI search: implications for information markets and human judgement at scale. *arXiv preprint arXiv:2602.13415*.

Athey, S., Mobius, M., and Pal, J. (2021). The impact of aggregators on internet news consumption. Technical report, National Bureau of Economic Research.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

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.

Chiou, L. and Tucker, C. (2017). Content aggregation by platforms: the case of the news media. *Journal of Economics & Management Strategy*, 26(4):782–805.

Fortunato, S., Flammini, A., Menczer, F., and Vespignani, A. (2006). Topical interests and the mitigation of search engine bias. *Proceedings of the National Academy of Sciences*, 103(34):12684–12689.

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

Gentzkow, M. and Shapiro, J. M. (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839.

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

Khosravi, M. and Yoganarasimhan, H. (2026). Impact of AI search summaries on website traffic: evidence from Google AI Overviews and Wikipedia. *Available at SSRN 6164926*.

Minici, M., Consonni, C., Cinus, F., and Manco, G. (2025). Auditing LLM editorial bias in news media exposure. *arXiv preprint arXiv:2510.27489*.

Padilla, N., Lam, H. T., Lambrecht, A., and Hollenbeck, B. (2025). The impact of LLM adoption on online user behavior. *Available at SSRN 5393256*.

Robertson, R. E., Jiang, S., Joseph, K., Friedland, L., Lazer, D., and Wilson, C. (2018). Auditing partisan audience bias within Google search. *Proceedings of the ACM on human-computer interaction*, 2(CSCW):1–22.

Steinacker-Olsztyn, N., Gosain, D., and Dao, H. (2026). Is misinformation more open? a study of robots. txt gatekeeping on the web. In *Proceedings of the ACM Web Conference 2026*, pages 7507–7516.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

Yang, K.-C. (2025). News source citing patterns in AI search systems. *arXiv preprint arXiv:2507.05301*.

Yang, K.-C. and Menczer, F. (2025). Accuracy and political bias of news source credibility ratings by large language models. In *Proceedings of the 17th ACM Web Science Conference 2025*, pages 127–137.

Zhang, P., Ye, Q., Peng, Z., Garimella, K., and Tyson, G. (2025). Source coverage and citation bias in LLM-based vs. traditional search engines. *arXiv preprint arXiv:2512.09483*.

Zhao, H. and Berman, R. (2026). Strategic response of news publishers to generative AI. *Available at SSRN 5992774*.

Zhu, Y. and Dukes, A. (2015). Selective reporting of factual content by commercial media. *Journal of Marketing Research*, 52(1):56–76.

Appendix

A MBFC Rating Scale Thresholds

In our main analysis we aggregate bias into five groups based on MBFC ratings and categorization: Far Left $[-10, -5]$, Left-Center $(-5, -2]$, Least Biased $(-2, +2)$, Right-Center $[+2, +5)$, and Far Right $[+5, +10]$; and factual reporting into four groups: Very High/High $[0, 2)$, Mostly Factual $[2, 4.5)$, Mixed $[4.5, 6.5)$, and Low/Very Low $[6.5, 10]$. Below we summarized these thresholds and their descriptions for each dimension.

Table 7: MBFC Factual Reporting score ranges

Category	Numeric Range	Description
Very High	[0, 0.1)	Consistently factual, uses credible information, no failed fact checks
High	[0.1, 2.0)	High factual, minor sourcing issues, reasonable fact-check record
Mostly Factual	[2.0, 4.5)	Generally reliable but may have occasional fact-check failures, transparency, and sourcing issues
Mixed	[4.5, 6.5)	Reliability varies; multiple fact-check failures, poor sourcing, lack of transparency, one-sidedness
Low	[6.5, 8.5)	Often unreliable; frequent fact-check failures and significant issues with sourcing, transparency, propaganda, conspiracies, and pseudoscience promotion
Very Low	[8.5, 10]	Consistently unreliable, heavily biased, with intentional misinformation likely

Table 8: MBFC Slant Rating score ranges

Category	Numeric Range	Description
Extreme Left	$[-10, -8)$	Extreme left bias
Far Left	$[-8, -7]$	Far left bias
Left	$(-7, -5]$	Left bias
Left-Center	$(-5, -2]$	Left-center bias
Least Biased	$(-2, +2)$	Minimal partisan bias
Right-Center	$[+2, +5)$	Right-center bias
Right	$[+5, +7)$	Right bias
Far Right	$[+7, +8)$	Far right bias
Extreme Right	$[+8, +10]$	Extreme right-wing bias

B Sample Distributions

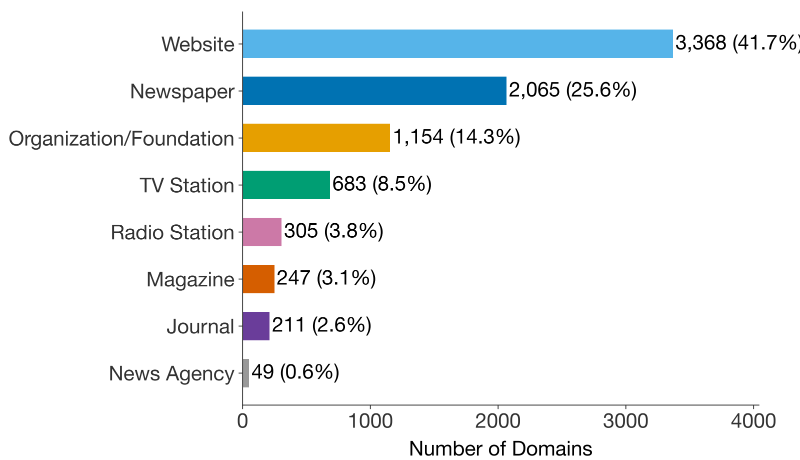

Figure 11: Distribution of sample domains by media type.

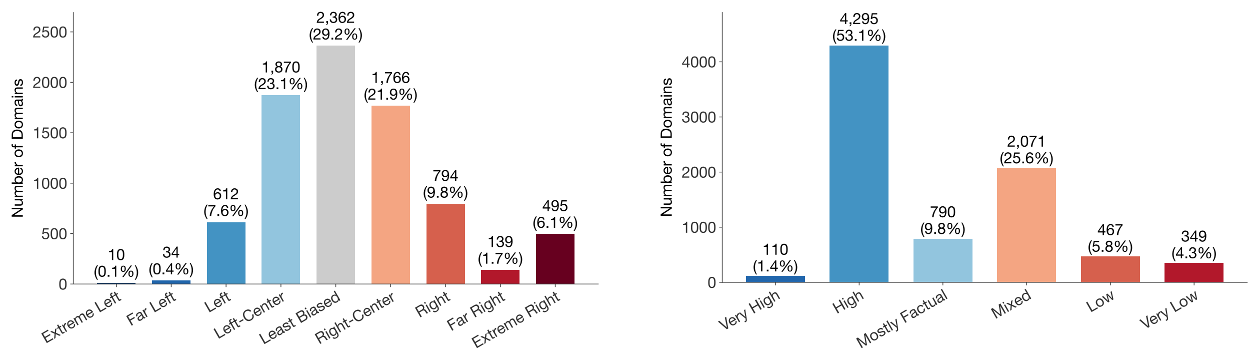

(a) Distribution by political slant group. (b) Distribution by factual reporting group.

Figure 12: Distribution of sample domains by political slant and factual reporting group.

C AI Bot Ban Definitions

From the daily `robots.txt` snapshots described in Section 3.3, we construct a monthly panel: for each domain × agent × month, the ban indicator is the maximum across all daily snapshots in that month (banned if any day shows a disallow directive). Months with no snapshot are forward-filled from the last observed value, since `robots.txt` is a persistent configuration. Table 9 lists the bot-to-platform mapping for the five tracked platforms.[14]

[14]† Google documents `google-extended` as governing both model training and grounding (using content to generate Gemini responses), so we include it in Gemini's access-ban definition as a content-access control: `https://developers.google.com/search/docs/crawling-indexing/overview-google-crawlers`. [a] OpenAI bot documentation: `https://developers.openai.com/api/docs/bots`. [b] Anthropic bot documentation: `https://support.claude.com/en/articles/8896518-does-anthropic-crawl-data-from-the-web-and-how-can-site-owners-block-the-crawler`. [c] No official documentation found for xAI crawler tokens; ban counts for Grok are substantially lower than for other platforms. We include all identified xAI crawler variants in the ban definition for completeness.

Table 9: Tracked AI bots and platform-specific access-ban definitions

Platform	Bot	Bot Category	Access Ban
ChatGPT	`chatgpt-user`	User Agent	✓
ChatGPT	`oai-searchbot`	Search Indexing	✓
ChatGPT	`gptbot`	Model Training[a]	
Claude	`claude-user`	User Agent	✓
Claude	`claude-searchbot`	Search Indexing	✓
Claude	`claudebot`	Model Training[b]	
Gemini	`google-notebooklm`	User Agent	✓
Gemini	`googlebot`	Search Indexing	✓
Gemini	`google-extended`	Model Training & Grounding	✓†
Perplexity	`perplexity-user`	User Agent	✓
Perplexity	`perplexitybot`	Search Indexing	✓
Grok	`grok`, `grok-crawler`, `grokbot`, `xai-grok`, `grok-deepsearch`, `grokai`, `grokcrawler`, `xai`, `xai-bot`	xAI crawler variants	✓[c]

Notes: A ✓ marks bots whose blocking contributes to a platform-specific access ban (see Section 3.3 for the access-ban definition).

D Peer-Ban Pre-Trends (Perplexity)

For model specification (4), under both TWFE and Sun–Abraham, these coefficients are near zero before the ban (Figure 13), consistent with parallel trends.

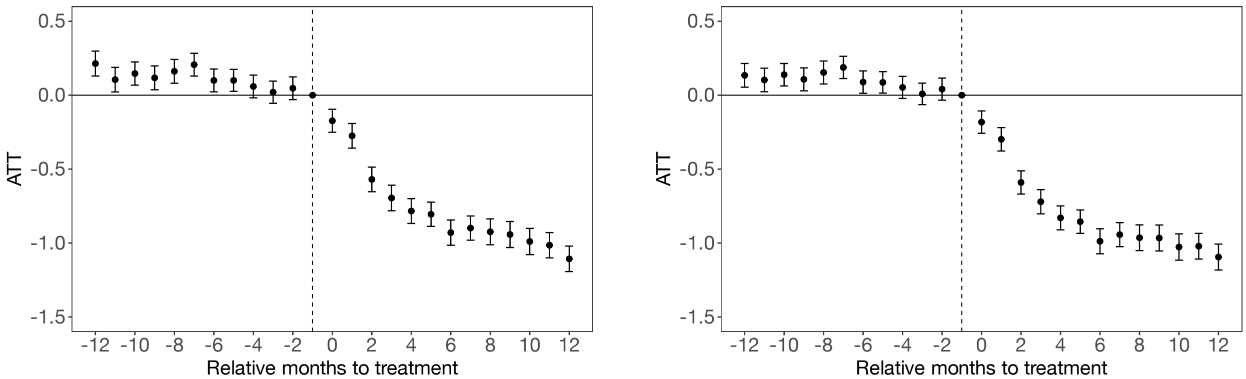

(a) TWFE (b) Sun–Abraham

Figure 13: Each panel plots period-by-period ATT with 95% CIs from equation (4). Each model includes domain and month fixed effects. Standard errors are clustered at the domain level.

For model specification (5), under both TWFE and Sun–Abraham, these coefficients are near zero before the ban (Figure 14), consistent with parallel trends.

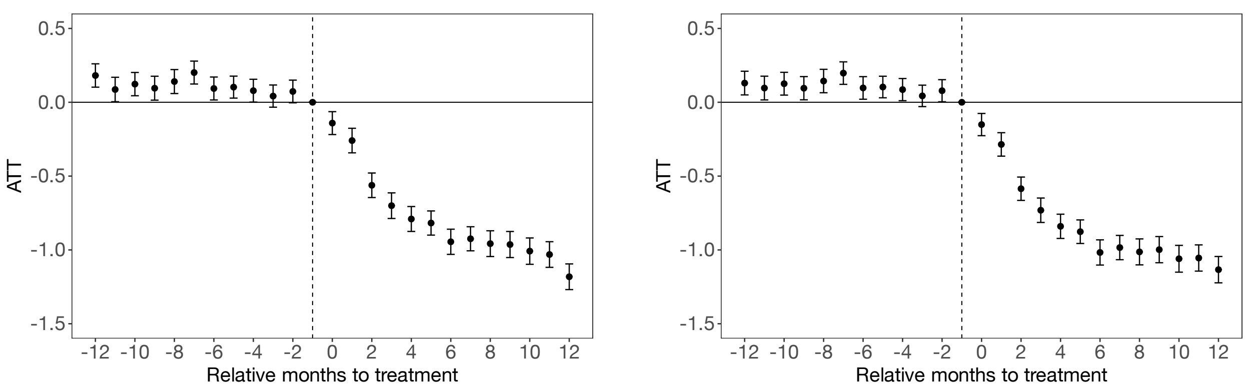

(a) TWFE (b) Sun–Abraham

Figure 14: Each panel plots period-by-period ATT with 95% CIs from equation (5). Each model includes domain and month fixed effects. Standard errors are clustered at the domain level.

E Peer-Ban ATT (Perplexity)

Table 10 reports overall ATT for model specification (4) under both TWFE and Sun–Abraham.

Table 10: Overall ATT from the joint peer-exposure regression.

Estimator	ATT	(SE)	95% CI
TWFE	-0.78^{***}	(0.03)	[−0.85, −0.71]
Sun–Abraham	-0.85^{***}	(0.03)	[−0.91, −0.78]

Notes: $N = 284{,}732$ domain–months. Standard errors clustered at the domain level are reported in parentheses. *Significance levels*: $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

Table 11 reports the overall ATT for model specification (5) under TWFE and Sun–Abraham.

Table 11: Overall ATT, heterogeneous-effects specification.

Estimator	ATT	(SE)	95% CI
TWFE	-0.79^{***}	(0.03)	[−0.86, −0.72]
Sun–Abraham	-0.89^{***}	(0.04)	[−0.96, −0.82]

Notes: $N = 284{,}732$ domain–months. Standard errors clustered at the domain level are reported in parentheses. *Significance levels*: $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.