

Social and Political Framing in Search Engine Results

Amrit Poudel, Tim Weninger

University of Notre Dame
{apoudel, tweninger}@nd.edu

Abstract

Search engines play a crucial role in shaping public discourse by influencing how information is accessed and framed. While prior research has extensively examined various dimensions of search bias—such as content prioritization, indexical bias, political polarization, and sources of bias—an important question remains underexplored: how do search engines and ideologically-motivated user queries contribute to bias in search results. This study analyzes the outputs of major search engines using a dataset of political and social topics. The findings reveal that search engines not only prioritize content in ways that reflect underlying biases but also that ideologically-driven user queries exacerbate these biases, resulting in the amplification of specific narratives. Moreover, significant differences were observed across search engines in terms of the sources they prioritize. These results suggest that search engines may play a pivotal role in shaping public perceptions by reinforcing ideological divides, thereby contributing to the broader issue of information polarization.

Introduction

“There is competition for those top ten seats. There is serious competition. People are trying to take away the top spots every day. They are always trying to fine-tune and tweak their HTML code and learn the next little trick. The best players even know dirty ways to “bump off” their competition while protecting their own sites.”

(Anderson and Henderson 1997)

Every day, billions of people turn to search engines like Google, Bing and DuckDuckGo for information, making these platforms indispensable in the digital age. With billions of queries submitted to search engines daily, they have become the modern-day equivalent of libraries—vast directories of human knowledge curated and presented at the click of a button. In many ways, they also serve as media companies (Goldman 2005). Google, one of the most widely used search engines, claims to rank results based on the following criteria¹: relevance, quality, usability, and context. These criteria align closely with the four components of search engine

quality assessment *i.e.*, index quality, results quality, search features quality, and search engine usability (Lewandowski and Höchstötter 2008). However, the impartiality of these algorithms has been increasingly called into question by political leaders, organizations, and scholars, raising concerns about the suppression and promotion of content (Kulshrestha et al. 2019; Gezici et al. 2021).

Nevertheless, users often trust search engine results (Pan et al. 2007) and treat top-ranked links as more credible. This gap between general awareness and everyday behavior highlights the need to examine how search systems reinforce or counter ideological framing.

Before the rise of search engines, traditional media such as newspapers, radio, and television served as central sources of information. Though once celebrated as pillars of democracy, these institutions have long been shaped by commercial and political interests (Introna and Nissenbaum 2000). Editorial decisions (*e.g.*, about what to print and how to print it) significantly influence how news is selected and framed, shaping public perception of political issues (Bush 1951). Search engines, while initially praised for democratizing access to information, are subject to similar biases (Poudel and Weninger 2024; Poudel et al. 2025). Ranking systems inherently prioritize one thing over another; and these decisions can affect user beliefs and behaviors. For example, search rankings have been shown to sway voting preferences by more than 20% among undecided voters (Epstein and Robertson 2015), a phenomenon known as the *Search Engine Manipulation Effect*. Although algorithms are designed to optimize relevance and usability, they are also influenced by commercial and political pressures (Introna and Nissenbaum 2000; Anderson and Henderson 1997). Despite periodic updates aimed at improving fairness, the inner workings of these systems remain opaque (Bozdog 2013), and users are often unaware of how search results are curated.

In addition, users bring cognitive biases to search. People often anchor on early results, seek confirmatory evidence, and overlook contradictions, especially in political searches (Azopardi 2021). These behaviors increase the impact of ranking, making it critical to study how political cues in queries affect what users see.

In the present work, we address these questions by comparing the results for politically inclined queries with those for neutral queries across three prominent news search engines:

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.google.com/intl/en/search/howsearchworks/how-search-works/ranking-results/>

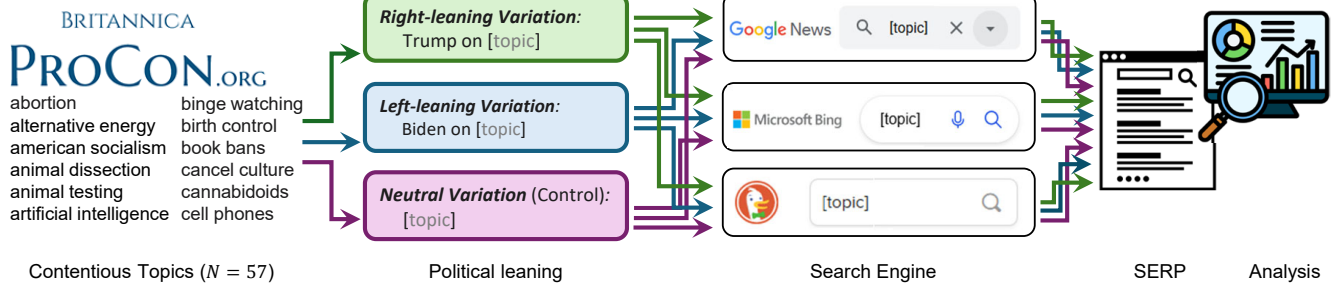


Figure 1: Flowchart illustrating the analysis of 57 debate topics from ProCon.org. Each query was tested in three variations: Right, Left, and Neutral and submitted to the news sections of three search engines. The retrieved results were aggregated and analyzed for potential variations in semantics, frames, source dynamics, ideological bias and factuality.

Google News, Bing News, and DuckDuckGo News. These sites were chosen due to their status as the most frequently visited search engines in the United States². We deliberately exclude Yahoo, as it is generally powered by Bing³, and focus on news search engine results pages (SERPs), as they often present a distinct framing of topics (Alam and Downey 2014). Compared to general web search, news verticals prioritize journalistic content and differ systematically in how they rank sources, including emphasis on certain outlets or viewpoints (Alam and Downey 2014). As such, they serve as a critical locus for studying how politically relevant information is curated and framed (Trielli and Diakopoulos 2022; Urman, Makhortykh, and Ulloa 2022). Specifically, we aim to address the following research questions:

1. **Semantic Polarity (RQ1)**. How do the news headlines returned for politically left- and right-leaning queries differ in semantic content from those generated by neutral queries across various search engines?
2. **Framing Dynamics (RQ2)**. How do search engines differ in their framing of politically salient topics across left-, right-, and neutral-leaning queries?
3. **Source Dynamics (RQ3)**. Do search engines disproportionately surface content from a limited number of dominant sources, and how does this tendency align with the political leanings of those sources?
4. **Bias and Factuality across Ideologies (RQ4)**. How do different search engines prioritize bias and factuality when curating top search results?

Findings. Through a series of experiments and data analysis detailed in the remainder of this paper, we have found that: (RQ1) Left-leaning and right-leaning queries return more ideologically inclined results. Neutral queries tend to bridge ideological divides. Results from Google News were less likely to reflect the ideology of the query; (RQ2) Bing and DuckDuckGo are more likely to return results having *Episodic*, *Human-interest*, and *Conflict* frames compared to

Google, with politically-oriented queries amplifying *Game/Strategy* and *Conflict* frames; (RQ3) search engines surface different sources disproportionately, with Bing and DuckDuckGo more closely aligned than Google; (RQ4) Google provides the most politically-neutral sources, while Bing and DuckDuckGo are more likely to return right-leaning sources; curiously, right-leaning queries unexpectedly shifted toward left-leaning sources on all platforms. Factuality is consistent across search engines, but sources are less factual for political queries.

Data Collection Methodology

We analyzed search engine results from the news verticals of three major platforms: Bing, Google, and DuckDuckGo. To evaluate political variation in content curation, we used 57 issue-based queries derived from ProCon.org⁴, a non-partisan source known for curating socially and politically contentious topics (Gezici et al. 2021; ProCon.org 2023). To introduce ideological cues, we appended Biden or Trump to each query, following prior work that uses political figure names to simulate partisan framing in search audits (Trielli and Diakopoulos 2022; Urman, Makhortykh, and Ulloa 2022). Importantly, these are queries about left- or right-leaning political figures, rather than indicators of any user’s own ideology.

Search Engine Data

Collecting search engine data presents unique challenges due to the dynamic and constantly evolving nature of search results. For this study, we utilized the SerpWow API⁵ to systematically gather search engine data within a single ten-hour window in Fall of 2024.

Mitigating Data Collection Bias To ensure a fair comparison across platforms, we synchronized the data collection process by retrieving results for each keyword from all three search engines simultaneously. This approach minimized temporal variations and allowed for a consistent evaluation of search engine outputs. Prior work has found no significant differences in search engine data collected within a 10–15

²<https://www.statista.com/forecasts/997254/most-used-search-engines-by-brand-in-the-us>

³<https://techcrunch.com/2015/04/16/microsoft-and-yahoo-renew-search-allian/>

⁴<https://www.britannica.com/procon>

⁵<https://trajectdata.com/serp/serp-wow-api/>

Search Engine	Right	Left	Neutral
Google News	3,809	3,542	4,697
DuckDuckGo News	4,749	4,900	5,825
Bing News	506	502	582

Table 1: Number of search results collected for politically right-leaning, left-leaning, and neutral queries across three search engines.

minute range (Gezici et al. 2021), further supporting the robustness of our synchronized data collection methodology. We specifically set the location to the United States and the language to English, retrieving results from up to the top five pages (if available); otherwise, the collection stopped when no further results were provided for a given query. Furthermore, the SerpWow API uses thousands of distributed proxy services to abate any locality-effects. For analysis, we focus on the top 20 results, as research indicates that users typically gather information from the top 10 results (Schultheiß and Lewandowski 2021) and rarely venture past the first 25 results (Glenski, Pennycuff, and Weninger 2017).

The total number of search results across different search engines for right-leaning, left-leaning, and neutral queries are presented in Table 1. A χ^2 test of independence was conducted to examine the relationship between search engine and political leaning. The test revealed a statistically significant association, $\chi^2(4, N = 29112) = 18.09, p = 0.0012$. However, the effect size, measured using Cramér’s V, was 0.018, indicating a very weak effect size in the number of results returned for search engines and political leaning.

Semantic Polarity

Headlines are the first point of engagement in a news article, shaping the reader’s expectations and framing their interpretation of the content (Konnikova 2014). They convey significance, emphasize gravitas, and reinforce authority (Papacharissi 2018). Due to this central role, we focus on headlines as the primary unit of analysis and ask: **(RQ1)** How do the news headlines returned for politically left- and right-leaning queries differ in semantic content from those generated by neutral queries across various search engines?

We examine the semantic polarity of search engines by comparing how their results for politically inclined queries differ from those for neutral queries. This analysis assesses the impact of political keywords on the content and framing of retrieved headlines. Previous studies have measured semantic polarization by analyzing linguistic differences in broadcast news and social media discourse (Mowshowitz and Kawaguchi 2002; Ding, Horning, and Rho 2023), often using contextualized language embeddings to quantify shifts in framing and context.

However, in search engine rankings individual results are not independent but rather part of an ordered list influenced by underlying ranking algorithms. This order is crucial to how the user views and interprets information. So, to analyze the alignment or divergence between result lists (e.g., comparing a neutral query against a politicized query), it is necessary




Figure 2: Cosine dissimilarity ($1 - \cos \theta$) between search result headlines, grouped by query leaning and search engine. Values reflect estimated marginal means from a linear mixed-effects model (LMM) with random intercepts for query key and fixed effects for search engine and query leaning. Error bars represent 95% confidence intervals. Results indicate that dissimilarity varies by query type and platform, with right-leaning queries often producing more internally divergent results.

to perform pairwise comparisons. Pairwise cosine similarity allows us to capture the semantic relationship between each pair of articles across result vectors.

Specifically, for each neutral query, we define left- and right-leaning variants. When these queries are submitted to a search engine, the engine returns results for each variant. We then used MPNetv2 (Reimers and Gurevych 2019), a pretrained sentence transformer, to encode each headline into to a 128-dimensional vector space and computed the pairwise comparison for 10,000 random headline-pairs using the cosine distance.

To account for dependencies across repetitions in the pairwise comparisons, we fit a linear mixed-effects model (LMM) predicting pairwise semantic dissimilarity using cosine distance between headline embeddings. The model included fixed effects for search engine and query leaning, and random intercepts for each headline to control for repeated items. This allowed us to examine how both platform design and political query framing shape the internal consistency of search results. We find that both factors significantly influence headline similarity, with some platforms returning more semantically cohesive results than others, and neutral queries generally producing more consistent content than political ones.

Figure 2 presents estimated marginal cosine dissimilarity scores between headlines, grouped by query leaning and search engine. We find that queries with political cues, especially left-leaning ones, tend to generate more divergent headline content, particularly on Bing and DuckDuckGo. By contrast, neutral (i.e., non-political) queries return more semantically similar headlines, suggesting broader alignment in framing. Google shows the most stable dissimilarity patterns across query types, indicating a greater degree of internal coherence in its retrieved results regardless of ideological input.

These results are not unexpected, but serves as a sanity check, demonstrating that the methodology is effectively cap-

turing meaningful semantic differences. The larger semantic distance between left-leaning and right-leaning content, compared to the neutral comparisons, validates the approach’s sensitivity to ideological variation in the data. Next, we will look how these differences impact how users see the news from various perspectives.

Framing Dynamics

The way information is presented significantly impacts how readers perceive and understand news stories. This concept, known as framing, originates from sociology (Goffman 1974), where Entman (1993) describes framing as the process of highlighting specific elements of a perceived reality to promote particular interpretations and evaluations, focusing on “selection” and “salience.” The crucial questions here are: Who chooses what to highlight, and why does it matter to the audience (Entman 1993)?

Primarily, journalists decide on news story headlines and ledes, effectively acting as news framers (Brüggemann 2014; Reese 2010; Baden 2019). However, this is not done in a vacuum; articles undergo extensive editorial processes to meet publishers’ standards before public release. This frame-building involves a dynamic interplay between journalists, editors, elites (Gans 1979; Tuchman 1978), and social movements (Cooper 2002).

Framing plays a pivotal role in shaping information and influencing decisions (Kahneman and Tversky 1984). Politicians and journalists often compete to control news narratives and guide public opinion (Riker 1986).

In this section, we address framing dynamics by asking: **(RQ2)** How do search engines frame politically salient topics across left-, right-, and neutral-leaning queries?

Framing theory distinguishes between issue-specific frames, which focus on particular topics, and generic frames, which apply across various stories (De Vreese 2005). Scholars typically classify frames into six archetypes: (1) conflict, (2) game/strategy, (3) thematic/issue, (4) economic consequences, (5) episodic, and (6) human interest (Baden 2019). For example, conflict frames emphasize competing viewpoints, while game/strategy frames highlight the tactics behind social and political struggles (Lawrence 2000). Although there may be several other types of frames and frame-ontologies, prior work has found that these six appear most frequently in news (De Vreese 2005); for a more thorough exploration of frames see the chapter by Baden (2019).

Our research builds on this framework by categorizing news headlines from our dataset into these specific frames. Through this, we aim to uncover how framing shapes public discourse and influences political perceptions.

Methodology To classify headlines into different frames, we employed GPT-4 as a labeling tool. Using a carefully designed prompt (provided in the appendix for reference), the model categorized headlines into one of five predefined frames. This approach enabled consistent and scalable labeling of a large dataset while preserving contextual nuance.

To assess stability, we examined all headlines that appeared multiple times. GPT-4 produced identical frame labels in

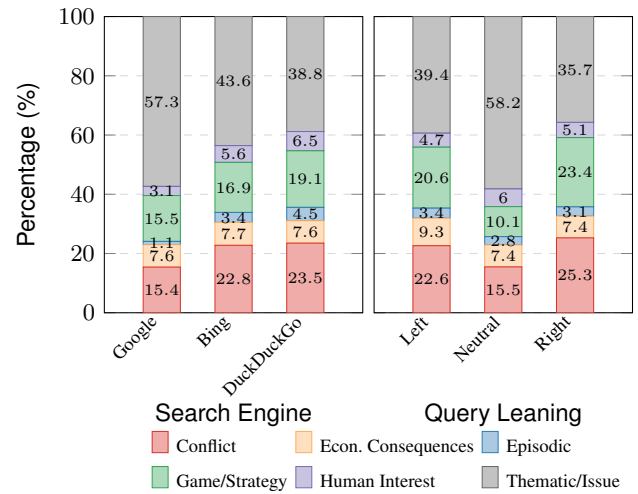


Figure 3: Percentage distribution of news frames across search engines and query leaning. Thematic/Issue frames dominate across all search engines and sources, while Episodic frames are the least common. DuckDuckGo exhibits a more balanced frame distribution, whereas partisan sources (left and right) show higher representation of Conflict and Game/Strategy frames.

93.75% of these cases (1,334 out of 1,423), demonstrating high internal consistency. We also conducted a manual validation study: human–human agreement was modest (Cohen’s $\kappa = 0.23$), reflecting the subjective difficulty of the task. Notably, in cases where human coders agreed, GPT-4 matched the majority label with substantial agreement (Cohen’s $\kappa = 0.69$, macro- $F_1 = 0.74$). Like previous work (Ding et al. 2022) shows, these results also suggest that GPT-4 may be more consistent and reliable than human annotators in this framing context.

Results

Before turning to LMM results, we begin with a descriptive overview of frame distributions across platforms and query leanings. As shown in Fig. 3, Thematic/Issue frames dominate across all platforms, with DuckDuckGo exhibiting a more balanced distribution compared to Google and Bing. Conflict and Game/Strategy frames are more prominent in partisan sources and right-leaning queries, while Episodic frames are the least frequent overall. A one-way ANOVA further reveals that average ranking position varies significantly by frame type ($F(5, 7895) = 3.54, p = .003$), indicating that certain frames tend to appear higher in search results regardless of political leaning.

To assess how framing prominence varies across political query leaning and search engine, we conducted two two-way ANOVAs using ranking position as the dependent variable. Then, because multiple results are returned for each query, we again fit linear mixed-effects models (LMMs) with random intercepts for query and fixed effects for frame type, search engine, query leaning, and ranking. This combination of analyses allows us to identify overall patterns while accounting

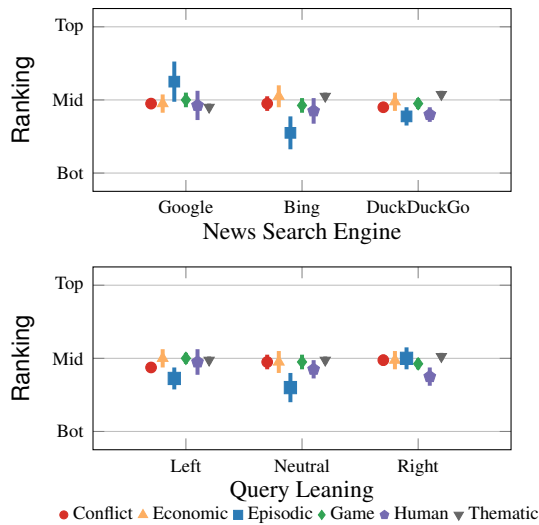


Figure 4: Tukey post-hoc comparisons of average ranking positions for each frame across search engines. These results highlight how ranking algorithms shape the visibility of narrative types across platforms.

for within-query dependencies. We also tested models with interaction terms, but these did not significantly improve model fit compared to the fixed-effects-only specification.

Frame vs Search Engine. We first examined whether search engines differ in how they prioritize different news frames. A two-way ANOVA revealed significant main effects of both frame type ($F(5, 8143) = 6.02, p < .001$) and search engine ($F(2, 8143) = 466.66, p < .001$), along with a significant interaction between the two ($F(10, 8143) = 5.64, p < .001$).

As shown in Fig. 4 (top), post-hoc Tukey comparisons reveal that Google systematically ranks Episodic frames higher (*i.e.*, earlier in the results), while DuckDuckGo and Bing return Episodic, and Human Interest frames in lower positions. These patterns suggest that search engines structurally shape the narrative emphasis users see, not just through content selection, but also through positional ranking.

Frame vs Query Leaning. We then tested whether the political orientation of the query influenced the ranking of frames (Fig. 4 (bottom)). A two-way ANOVA with frame and query leaning as predictors of normalized rank revealed a significant main effect of frame ($F(5, 5107) = 4.42, p < .001$), but no main effect of political leaning ($F(2, 5107) = 0.53, p = .59$). A modest but significant interaction between frame and leaning ($F(10, 5107) = 1.90, p = .041$) suggests that while overall frame rankings remain relatively stable, there are some subtle differences in frame prominence depending on the ideological slant of the query.

Together, these findings demonstrate that framing is not only unevenly distributed across platforms and political query types, but also systematically structured in terms of visibility and rank.

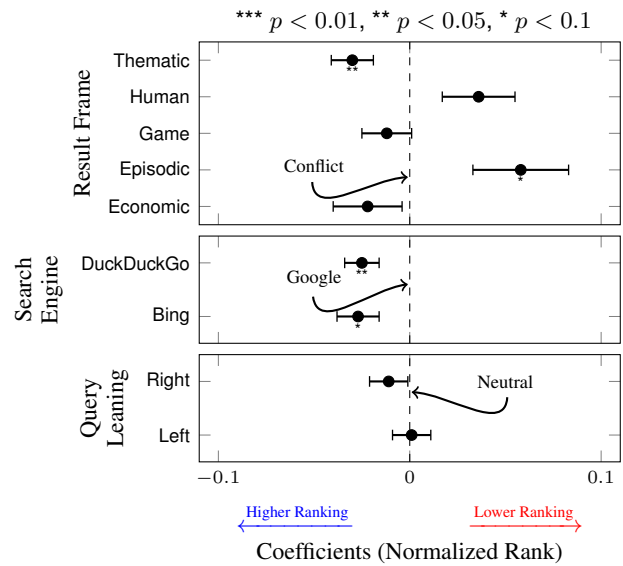


Figure 5: Estimated coefficients from a linear mixed-effects model predicting normalized rank of news results, with random intercepts by query. Episodic frames are ranked lower, while thematic frames are elevated across platforms. Query leaning has limited effect on result prominence. Error bars represent 95% confidence intervals.

Mixed Effects Model. To account for within-query dependencies, we fit an LMM predicting the normalized rank of search results from frame type, search engine, and query leaning, with random intercepts for each query key. Results are visualized in Fig. 5.

We find that frame type significantly influences result prominence. Compared to conflict frames, episodic frames are ranked lower, while thematic frames appear higher in the result list. This suggests that news search engines may elevate more issue-focused content and downrank personal or anecdotal narratives.

Query leaning had limited effect. There were no statistically significant differences in ranking between politically neutral queries and those framed as left- or right-leaning, suggesting that ideological slant in the query does not systematically affect where results appear in the ranking.

Together, these findings suggest that differences in frame visibility are more strongly shaped by platform architecture and narrative style than by query ideology. Google appears more consistent in its ranking behavior, while Bing and DuckDuckGo elevate more variable frame types.

Source Dynamics

Search engines are often seen providing broad access to information. However, the challenge lies in determining which data to prioritize for each query (Introna and Nissenbaum 2000). If only ten news articles can be returned for a query, which ten sources should be chosen? This challenge is amplified by the influence of special interests (McChesney 1998), raising concerns about whether certain sources dominate

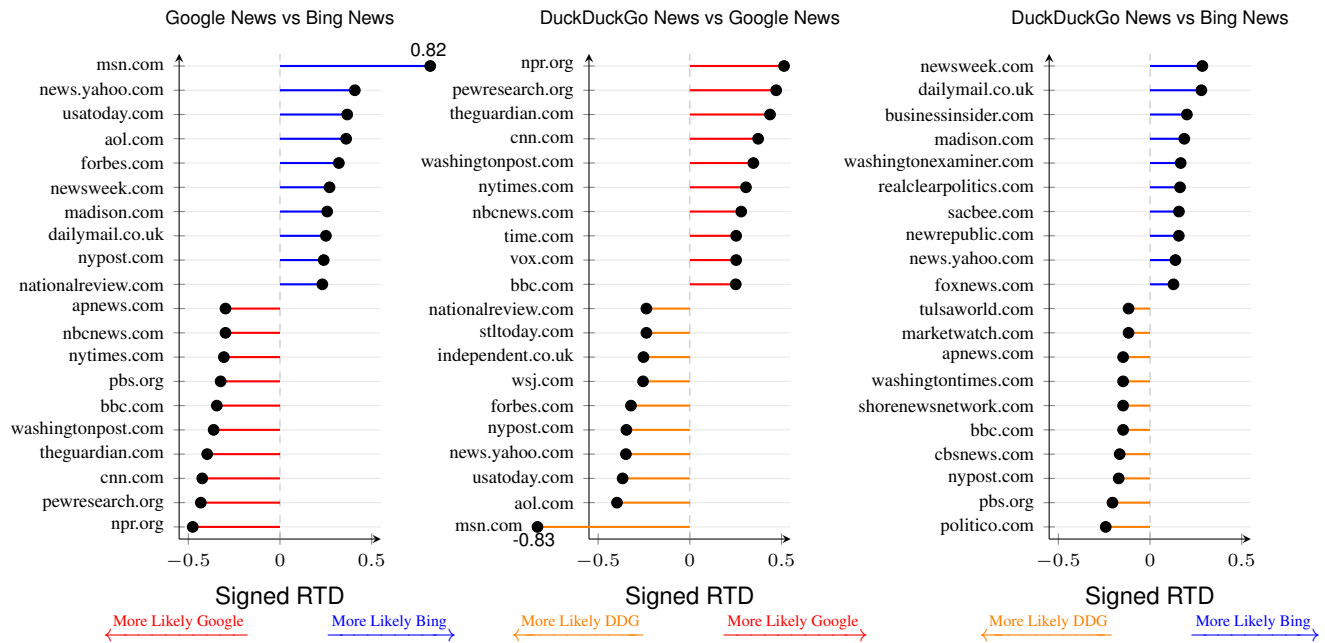


Figure 6: Divergence of sources across search engines. The figure illustrates the top 10 most divergent sources based on Rank Turbulence Divergence (RTD) for Google News, Bing News, and DuckDuckGo News. Notable differences in source prioritization are observed, with Google favoring NPR, Pew Research, CNN, and The Guardian, while Bing shows a strong preference for MSN. A similar pattern is seen when comparing Google to DuckDuckGo, while Bing and DuckDuckGo display less divergence, with subtle differences in surfacing sources.

search results, potentially marginalizing diverse or alternative viewpoints.

To understand this better, in this section we investigate source dynamics in search engines. Specifically, as we ask: **(RQ3)** Do search engines disproportionately surface content from a limited number of dominant sources, and how does this align with the political leanings of those sources?

Methodology To examine how source rankings differ across search engines, we use Rank Turbulence Divergence (RTD) (Dodds et al. 2023), a method that quantifies and visualizes variations in rankings, helping us understand whether alternative media is adequately represented.

We analyze the sources from the search engine results pages using traditional token analysis to determine any relative difference. While various statistical methods can compare these distributions (Cha 2007; Deza and Deza 2006), they often struggle with the Zipfian nature typical of text datasets (Gerlach, Font-Clos, and Altmann 2016; Dodds et al. 2023). To tackle this, we employ Rank Turbulence Divergence (RTD) (Dodds et al. 2023), which quantifies the disparity in frequency distributions of sources across different search engines, offering a clearer picture of information diversity in search results. See the Appendix for details on RTD.

Results Table 2 presents the RTD values between the source distributions across different search engines. Statistical tests were conducted by performing 1000 bootstrap

Search Engine Comparison		RTD Value
Google News	Bing News	0.5896
Google News	DuckDuckGo News	0.5781
Bing News	DuckDuckGo News	0.3141

Table 2: Rank Turbulence Divergence (RTD) between Source Distributions Across Search Engines

iterations of 100 samples each. 95% confidence intervals are on the order of 10^{-16} and are therefore omitted from the table. The RTD value between Bing and DuckDuckGo suggests a closer alignment in the types of sources these engines prioritize compared to Google.

Although these results are interesting, we are mostly interested in which sources are most divergent, *i.e.*, contribute most to the overall RTD value. Figure. 6 illustrates the most divergent sources for different search engines. Here we focus on the top 10 most divergent sources. The comparison between Google and Bing reveals notable differences in the prioritization of sources. Sources such as NPR, Pew Research, CNN, and The Guardian are more likely to appear in Google’s results, while Bing shows a strong preference for MSN, likely because MSN and Bing are both owned by the same parent company. For instance, MSN’s RTD score of +0.82 in Bing versus Google indicates it ranks substantially higher in Bing’s results. A similar pattern emerges when comparing Google

Search Engine	Left	Right	Neutral	Total
Bing News	172	159	191	522
DuckDuckGo News	273	249	308	830
Google News	222	198	274	694
Total	667	606	773	2046

Table 3: Unique Sources by Search Engine and Query Leaning

with DuckDuckGo. In contrast, the comparison between Bing and DuckDuckGo reveals less pronounced divergence. Despite this smaller gap, subtle differences in content surfacing can still be observed.

This variation underscores the distinct curation strategies employed by search engines, which can influence the diversity of information available to users depending on their chosen platform.

Table 3 presents the number of unique sources by search engine and political leaning⁶. A χ^2 test of independence found significant differences between the number of unique sources between search engines and political leanings $\chi^2(4, N = 2046) = 1.39, p = 0.85$. This suggests that there is no significant association in the number of unique sources between search engine and political leanings.

Bias and Factuality across Ideologies

While research on algorithmic transparency in social media and Web search has been extensive, the relationship between search engines and political ideology remains underexplored. This section investigates (RQ4) how ideological bias influences the prioritization of sources in top search results, examining whether and how search engines curate results based on political leanings and how this affects the balance between bias and factuality.

Source Ideological Bias

To assess the ideological leaning and factual reliability of retrieved sources, we rely on ratings from Media Bias/Fact Check (MBFC), an independent and widely used news source rating platform (Bozarth, Saraf, and Budak 2020). MBFC assigns each outlet a political bias label on a 7-point scale ranging from extreme left (-3) to extreme right (+3), and a factuality rating on a 6-point ordinal scale from very low to very high. Sources categorized as conspiracy, satire, or unknown were excluded from quantitative analysis. MBFC ratings have been widely utilized in previous studies to annotate bias and factual accuracy (Starbird 2017; Darwish, Magdy, and Zanouada 2017) and serve as a ground truth for predictive models (Dinkov et al. 2019).

Following prior work (Weld, Glenski, and Althoff 2021), we converted MBFC’s categorical ratings into numerical scores to enable statistical comparison across search engines and query types.

⁶Note that sources are not disjoint across conditions, so the totals in Table 3 do not represent unique sources across conditions

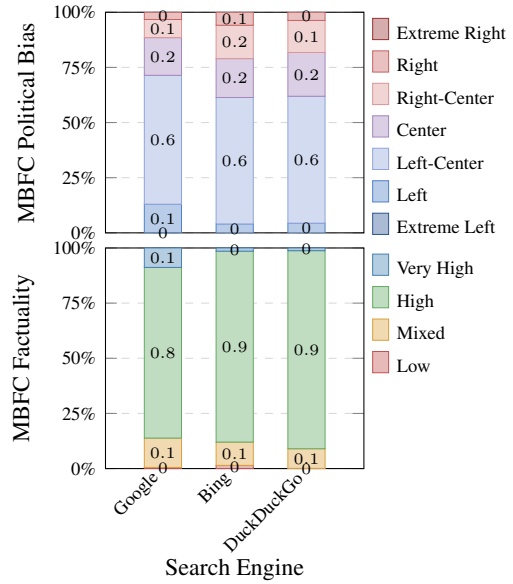


Figure 7: Stacked bar charts showing the distribution of MBFC Political Bias (top) and MBFC Source Factuality (bottom) across search engines. Google and DuckDuckGo favor center and left-leaning sources, while Bing returns more right-leaning content. All platforms predominantly surface high-factuality sources.

To evaluate differences in source bias and factuality, we applied a combination of statistical methods. We first conducted two-way ANOVA tests with search engine and query leaning as independent variables. Post-hoc Tukey HSD tests were used to assess pairwise differences. To account for the nested structure of our data—multiple results per query—we also fit linear mixed-effects models (LMMs) with random intercepts for each query and fixed effects for search engine, query leaning, and ranking position. This combination of analyses allows us to detect both overall patterns and within-query variation.

The key questions driving our analysis are: (1) Do search engines amplify or mitigate the ideological leaning of the sources behind the top-ranked content? (2) Are certain search engines more prone to ideological bias than others? (3) How does the interplay between search engine and query leaning affect the political orientation of the retrieved information?

Result Bias vs Search Engine and Query Leaning

We begin our analysis by examining how the political bias of retrieved news sources varies as a function of search engine and query leaning. Using the MBFC-coded bias scores described above, we fit a two-way ANOVA model with search engine, query leaning, and their interaction as predictors of source bias.

The ANOVA revealed significant main effects of both search engine ($F(2, 5143) = 30.69, p < .001$) and political query leaning ($F(2, 5143) = 15.01, p < .001$) on the bias scores of returned sources. The interaction between search engine and leaning was not significant, suggesting that the

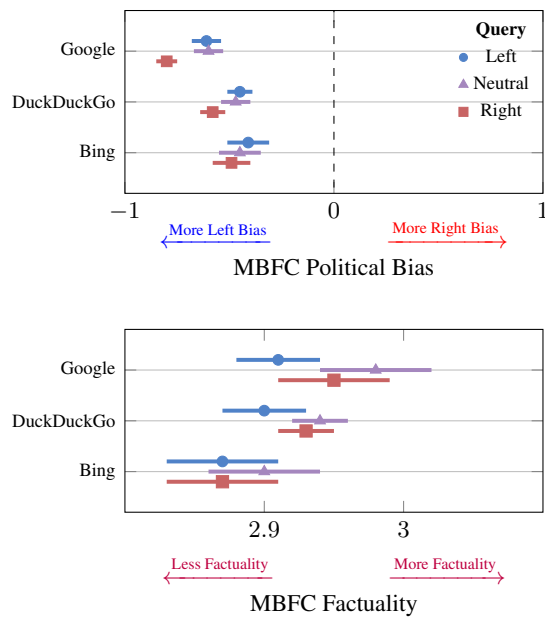


Figure 8: Mean source bias (top) and factuality (bottom) by search engine and query leaning. Bars show 95% confidence intervals. While DuckDuckGo and Bing return more right-leaning sources than Google, Tukey tests reveal that right-leaning queries unexpectedly retrieve more left-leaning content. Google also consistently returns higher factuality sources than Bing, with neutral queries producing the most reliable results across platforms.

effect of political leaning on source bias is consistent across platforms.

Post-hoc Tukey tests, illustrated in Fig. 8, showed that Google returns significantly less biased sources than both Bing ($p < .001$) and DuckDuckGo ($p < .001$), while Bing and DuckDuckGo do not significantly differ from one another. In terms of query leaning, right-leaning queries are associated with more ideologically extreme sources than both left-leaning ($p < .001$) and neutral queries ($p < .001$), while no significant difference was found between left-leaning and neutral queries ($p = .99$).

Overall, these results demonstrate that both platform and query ideology systematically shape the political slant of search results, with Google tending toward centrist sources and right-leaning queries drawing more partisan content across engines.

Mixed Effects Model The preceding ANOVA and Tukey tests assume that all observations are independent, but multiple results are returned for each query, introducing within-query dependencies. To account for this structure, we fit a linear mixed-effects model (LMM) with fixed effects for search engine, query leaning, and result ranking, and a random intercept for each query.

We compared a main-effects model to one that included interaction terms between search engine and query leaning. A likelihood ratio test found no improvement in fit from

the interaction terms ($\chi^2(4) = -10.78, p = 1.00$). We therefore report results from the more parsimonious main-effects model.

This model confirmed our earlier findings: Bing and DuckDuckGo return significantly more right-leaning sources than Google, and right-leaning queries are associated with more partisan content compared to neutral queries. Left-leaning queries do not differ significantly from neutral ones, and ranking position had no measurable effect on bias. A random intercept for each query accounted for within-query dependencies. Full model coefficients are presented in Fig. 9.

The ideological leaning of the user’s query consistently influenced the bias of retrieved sources. Across all models, **right-leaning queries were associated with a clear shift toward more left-leaning sources**—more so than either neutral or left-leaning queries. This counterintuitive pattern may reflect algorithmic efforts by search engines to balance or counteract the ideological slant of the query itself. Alternatively, it may be shaped by the dynamics of *negative partisanship*, wherein media—particularly left-leaning outlets—devote more coverage to political opponents than to ideological allies (Abramowitz and Webster 2018). For example, Republican-leaning sources may publish more critical content about Biden than Democratic sources publish about Trump.

Result Factuality vs Search Engine and Query Leaning

To assess how search engines and query leaning influence the reliability of retrieved information, we assigned ordinal scores to MBFC factuality ratings, ranging from 1 (low) to 5 (very high). A two-way ANOVA revealed significant main effects of both search engine ($F(2, 5143) = 9.05, p < .001$) and query leaning ($F(2, 5143) = 5.89, p = .003$), with no significant interaction.

Tukey post-hoc comparisons indicate that Google returns significantly higher-factuality sources than both Bing ($p < .001$) and DuckDuckGo ($p = .018$). Bing and DuckDuckGo also differ slightly ($p = .018$), though to a lesser degree. Among query types, neutral queries yielded more factual sources than left-leaning queries ($p = .003$), while differences between right-leaning queries and the other groups were not statistically significant.

Mixed Effects Model To again account for the nested structure of the data (multiple results per query), we also fit a linear mixed-effects model with fixed effects for search engine, query leaning, and ranking, and a random intercept for each query key. Results from this model mirror the ANOVA findings: Google surfaces the most factual content, followed by DuckDuckGo and Bing. Left-leaning queries produced marginally less factual results than neutral queries, while right-leaning queries showed a borderline-significant decrease. Ranking position had no measurable effect on source quality. Full model estimates are shown in Fig. 9 (bottom).

Taken together, these findings indicate that both platform and query framing subtly shape the factuality of retrieved sources. Neutral queries consistently yield the most reli-

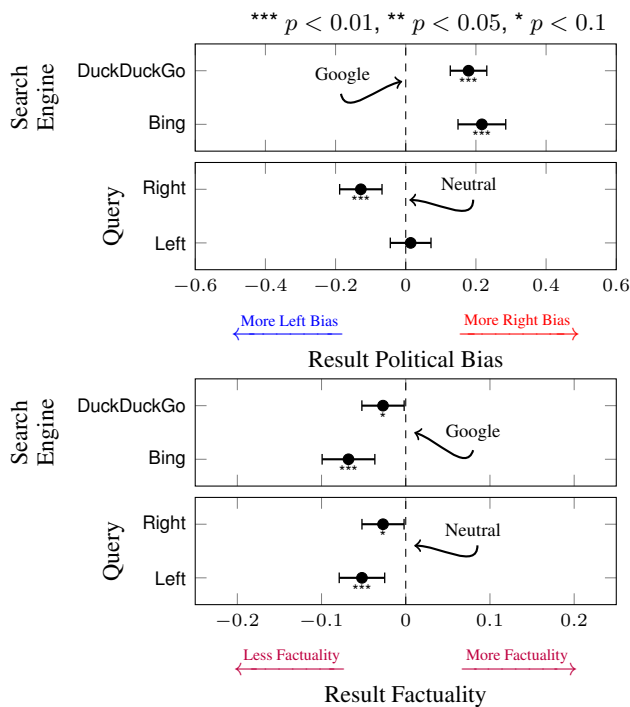


Figure 9: Fixed effects estimates from linear mixed-effects models predicting result bias (top) and result factuality (bottom). Error bars show 95% confidence intervals. Results confirm that Google returns more left-leaning and higher-factuality sources than other platforms. Right-leaning queries are associated with more partisan sources, though not lower factuality. Ranking position was not a significant predictor in either model.

able information across platforms. While differences across search engines are small, they are directionally consistent: Google provides a slight edge in factual accuracy, especially compared to DuckDuckGo. Importantly, politically charged queries (whether left- or right-leaning) tend to surface less factual sources than neutral queries, underscoring how user framing influences not just the ideological content but also the quality of results.

Discussion

Our study examines how major news search engines—Google News, Bing News, and DuckDuckGo News—mediate access to politically inclined queries. We analyze semantic polarity, framing effects, source selection, and ideological bias and factuality to understand their influence on the online news ecosystem and public discourse.

Neutral queries act as bridges, offering semantically diverse results that promote balanced perspectives crucial for democracy. In contrast, politically charged queries exhibit greater polarization, often reinforcing ideological silos. Notably, Google News mitigates this effect better than Bing News and DuckDuckGo News.

Framing analysis reveals that Bing and DuckDuckGo pri-

oritize conflict and strategy narratives, especially for right-leaning queries, potentially deepening partisan divides. These tendencies influence how users interpret and emotionally engage with news stories (Entman 1993), heightening perceptions of political contention while reducing focus on substantive issues.

Our source-level analysis shows that Google prioritizes outlets like NPR and CNN, while Bing and DuckDuckGo exhibit more overlap with each other and include a broader range of partisan and lesser-known domains. This divergence may stem from Google’s greater investment in algorithmic trustworthiness metrics or its different editorial standards. While it is unsurprising that different platforms surface different sources, the consequences are nontrivial: search engine design choices systematically shape what counts as visible, credible news in response to a user’s query.

One especially noteworthy finding is that right-leaning queries tend to retrieve more left-leaning sources. While this could reflect algorithmic efforts to balance perceived user bias, it also aligns with theories of negative partisanship, where partisan media disproportionately focus on critiquing the opposition. This suggests that even oppositional coverage can amplify attention to certain political actors, shaping perceptions indirectly.

In terms of factuality, DuckDuckGo returned consistently high-factuality sources across queries, while Google’s broader inclusion criteria produced more variability. This highlights a subtle tradeoff: platforms that emphasize diversity and breadth may also expose users to more mixed-quality information. Designing for both balance and credibility remains a persistent challenge.

Taken together, our findings reveal how search engines act not just as gatekeepers but as agenda-setters, shaping what users see, how it is framed, and whose voices are amplified. These effects are not accidental—they arise from the structure of search algorithms, source selection heuristics, and response to ideological cues embedded in queries.

Our results suggest three concrete implications. First, encouraging neutral or descriptive search behavior could reduce polarization in news exposure. Second, platform-level decisions about source and frame prioritization have downstream effects on users’ political understanding. Third, designers and policymakers should treat algorithmic transparency and ideological fairness not as abstract goals but as measurable, testable system behaviors that can be audited and improved.

Conclusions

In conclusion, we demonstrate that search engines not only mediate access to information but also actively shape the nature of political discourse by influencing semantic exposure, framing, source diversity, and factuality. The variations among Google News, Bing News, and DuckDuckGo News reflect differing algorithmic approaches with significant implications for ideological polarization and informed citizenship.

As digital platforms continue to be primary gateways to information, it is imperative for search engines to prioritize transparency, promote content diversity, and uphold high standards of factual accuracy. Stakeholders—including platform

developers, policymakers, and educators—must collaborate to ensure that the design and operation of these algorithms serve the public interest by facilitating access to reliable and diverse information. Only through such concerted efforts can we mitigate the risks of echo chambers and support a vibrant, informed, and democratic society.

Limitations While our study offers insights into how news search engines mediate politically inclined queries, several limitations merit attention. First, data were collected in late September 2024, during the lead-up to the U.S. presidential election. This timing is appropriate for capturing politically salient dynamics but may limit generalizability to less active political periods. In addition, we did not constrain result publication dates, so some headlines may predate the snapshot. Future research should adopt longitudinal designs to examine how platform responses evolve over time, especially around major events.

Second, our analysis focuses on English-language queries in the U.S., limiting its applicability to other regions and linguistic contexts. Expanding to multilingual and cross-national comparisons would offer a more global view of how search engines curate political information.

Third, our framing measures rely only on the headline that appears on the search-results page. This choice reflects the practical reality that these cues are all most users see before deciding whether to click. Nonetheless, restricting the analysis to surface text may miss frames or partisan nuance embedded in the full article. Headlines often compress or soften the ideological stance found deeper in the story, and some outlets craft headlines for engagement rather than comprehensive framing. Future work should incorporate fuller article context, e.g., first-paragraph or entire-text embeddings, where licensing permits, to assess when headline framing diverges from in-depth content and how that affects user interpretation.

Finally, although our research sheds light on the content presented by search engines, it does not measure the priming effects of framing, sources, or ideological bias on user opinions and public discourse. Understanding the influence of these factors on individual attitudes and societal polarization requires experimental approaches. Future work could involve controlled experiments or surveys to assess how exposure to specific frames or biased sources affects opinion shifts, providing deeper insights into the causal relationships between search engine content and public perceptions.

By addressing these limitations, subsequent research can build upon our findings to enhance the understanding of search engines' roles in shaping information access and the broader implications for democratic engagement and public discourse.

References

- Abramowitz, A. I.; and Webster, S. W. 2018. Negative partisanship: Why Americans dislike parties but behave like rabid partisans. *Political Psychology*, 39: 119–135.
- Alam, M. A.; and Downey, D. 2014. Analyzing the content emphasis of web search engines. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, 1083–1086.
- Anderson, P.; and Henderson, M. 1997. Hits to sales. *Retrieved December*, 12: 2001.
- Azzopardi, L. 2021. Cognitive biases in search: a review and reflection of cognitive biases in Information Retrieval. In *Proceedings of the 2021 conference on human information interaction and retrieval*, 27–37.
- Baden, C. 2019. Framing the news. *The handbook of journalism studies*, 229–245.
- Bozarth, L.; Saraf, A.; and Budak, C. 2020. Higher ground? How groundtruth labeling impacts our understanding of fake news about the 2016 US presidential nominees. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, 48–59.
- Bozdog, E. 2013. Bias in algorithmic filtering and personalization. *Ethics and information technology*, 15: 209–227.
- Brüggemann, M. 2014. Between frame setting and frame sending: How journalists contribute to news frames. *Communication theory*, 24(1): 61–82.
- Bush, C. R. 1951. The Analysis of Political Campaign News. *Journalism and Mass Communication Quarterly*, 28(2): 250.
- Cha, S.-H. 2007. Comprehensive survey on distance/similarity measures between probability density functions. *City*, 1(2): 1.
- Cooper, A. H. 2002. Media framing and social movement mobilization: German peace protest against INF missiles, the Gulf War, and NATO peace enforcement in Bosnia. *European Journal of Political Research*, 41(1): 37–80.
- Darwish, K.; Magdy, W.; and Zanouda, T. 2017. Trump vs. Hillary: What went viral during the 2016 US presidential election. In *Social Informatics: 9th International Conference, SocInfo 2017, Oxford, UK, September 13-15, 2017, Proceedings, Part I* 9, 143–161. Springer.
- De Vreese, C. H. 2005. News framing: Theory and typology. *Information design journal+ document design*, 13(1): 51–62.
- Deza, M.-M.; and Deza, E. 2006. *Dictionary of distances*. Elsevier.
- Ding, B.; Qin, C.; Liu, L.; Chia, Y. K.; Joty, S.; Li, B.; and Bing, L. 2022. Is gpt-3 a good data annotator? *arXiv preprint arXiv:2212.10450*.
- Ding, X.; Horning, M.; and Rho, E. H. 2023. Same Words, Different Meanings: Semantic Polarization in Broadcast Media Language Forecasts Polarity in Online Public Discourse. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, 161–172.
- Dinkov, Y.; Ali, A.; Koychev, I.; and Nakov, P. 2019. Predicting the leading political ideology of YouTube channels using acoustic, textual, and metadata information. *arXiv preprint arXiv:1910.08948*.
- Dodds, P. S.; Minot, J. R.; Arnold, M. V.; Alshaabi, T.; Adams, J. L.; Dewhurst, D. R.; Gray, T. J.; Frank, M. R.; Reagan, A. J.; and Danforth, C. M. 2023. Allotaxonomy and rank-turbulence divergence: A universal instrument for comparing complex systems. *EPJ Data Science*, 12(1): 37.
- Entman, R. M. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication*, 43(4): 51–58.

- Epstein, R.; and Robertson, R. E. 2015. The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences*, 112(33): E4512–E4521.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles/>.
- Gans, H. J. 1979. Deciding what's news: Story suitability. *Society*, 16(3): 65–77.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Gerlach, M.; Font-Clos, F.; and Altmann, E. G. 2016. Similarity of symbol frequency distributions with heavy tails. *Physical Review X*, 6(2): 021009.
- Gezici, G.; Lipani, A.; Saygin, Y.; and Yilmaz, E. 2021. Evaluation metrics for measuring bias in search engine results. *Information Retrieval Journal*, 24: 85–113.
- Glenski, M.; Pennycuff, C.; and Weninger, T. 2017. Consumers and curators: Browsing and voting patterns on reddit. *IEEE Transactions on Computational Social Systems*, 4(4): 196–206.
- Goffman, E. 1974. Frame analysis: An essay on the organization of experience. *Northeastern UP*.
- Goldman, E. 2005. Search engine bias and the demise of search engine utopianism. *Yale JL & Tech.*, 8: 188.
- Introna, L. D.; and Nissenbaum, H. 2000. Shaping the Web: Why the politics of search engines matters. *The information society*, 16(3): 169–185.
- Kahneman, D.; and Tversky, A. 1984. Choices, values, and frames. *American psychologist*, 39(4): 341.
- Konnikova, M. 2014. How headlines change the way we think. *The New Yorker*, 17.
- Kulshrestha, J.; Eslami, M.; Messias, J.; Zafar, M. B.; Ghosh, S.; Gummadi, K. P.; and Karahalios, K. 2019. Search bias quantification: investigating political bias in social media and web search. *Information Retrieval Journal*, 22: 188–227.
- Lawrence, R. G. 2000. Game-framing the issues: Tracking the strategy frame in public policy news. *Political Communication*, 17(2): 93–114.
- Lewandowski, D.; and Höchstätter, N. 2008. Web searching: A quality measurement perspective. In *Web search: Multidisciplinary perspectives*, 309–340. Springer.
- McChesney, R. W. 1998. Making media democratic. *Boston Review*, 23(3–4): 4–10.
- Mowshowitz, A.; and Kawaguchi, A. 2002. Assessing bias in search engines. *Information Processing & Management*, 38(1): 141–156.
- Pan, B.; Hembrooke, H.; Joachims, T.; Lorigo, L.; Gay, G.; and Granka, L. 2007. In Google we trust: Users' decisions on rank, position, and relevance. *Journal of computer-mediated communication*, 12(3): 801–823.
- Papacharissi, Z. 2018. The importance of being a headline. *Trump and the Media*, 71–78.
- Poudel, A.; Ding, Y.; Pfeffer, J.; and Weninger, T. 2025. Digital Gatekeepers: Google's Role in Curating Hashtags and Subreddits. *arXiv preprint arXiv:2506.14370*.
- Poudel, A.; and Weninger, T. 2024. Navigating the Post-API Dilemma. In *Proceedings of the ACM on Web Conference 2024*, 2476–2484.
- ProCon.org. 2023. ProCon.org - Pros and Cons of Controversial Issues. Accessed: 2023-09-30.
- Reese, S. D. 2010. Finding frames in a web of culture: The case of the war on terror. In *Doing news framing analysis*, 33–58. Routledge.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Riker, W. H. 1986. *The art of political manipulation*. Yale University Press.
- Schultheiß, S.; and Lewandowski, D. 2021. How users' knowledge of advertisements influences their viewing and selection behavior in search engines. *Journal of the Association for Information Science and Technology*, 72(3): 285–301.
- Starbird, K. 2017. Examining the alternative media ecosystem through the production of alternative narratives of mass shooting events on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11, 230–239.
- Trielli, D.; and Diakopoulos, N. 2022. Algorithmic agenda setting: The shape of search media during the 2020 US election#. *Special Issue Theme: AI and the News*, 45.
- Tuchman, G. 1978. Making news: A study in the construction of reality. *Free Press*.
- Urman, A.; Makhortykh, M.; and Ulloa, R. 2022. The matter of chance: Auditing web search results related to the 2020 US presidential primary elections across six search engines. *Social science computer review*, 40(5): 1323–1339.
- Weld, G.; Glenski, M.; and Althoff, T. 2021. Political bias and factualness in news sharing across more than 100,000 online communities. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 796–807.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
 - (e) Did you describe the limitations of your work? **Yes**

- (f) Did you discuss any potential negative societal impacts of your work? *NA*
 - (g) Did you discuss any potential misuse of your work? *NA*
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? *NA*
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? *Yes*
2. Additionally, if your study involves hypotheses testing...
- (a) Did you clearly state the assumptions underlying all theoretical results? *NA*
 - (b) Have you provided justifications for all theoretical results? *NA*
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? *NA*
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? *NA*
 - (e) Did you address potential biases or limitations in your theoretical framework? *NA*
 - (f) Have you related your theoretical results to the existing literature in social science? *NA*
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? *NA*
3. Additionally, if you are including theoretical proofs...
- (a) Did you state the full set of assumptions of all theoretical results? *NA*
 - (b) Did you include complete proofs of all theoretical results? *NA*
4. Additionally, if you ran machine learning experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? *NA*
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? *NA*
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? *NA*
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? *NA*
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? *NA*
 - (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? *NA*
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**

- (a) If your work uses existing assets, did you cite the creators? *Yes*
 - (b) Did you mention the license of the assets? *Yes*
 - (c) Did you include any new assets in the supplemental material or as a URL? *NA*
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? *NA*
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? *NA*
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? *NA*
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? *NA*
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? *NA*
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? *NA*
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? *NA*
 - (d) Did you discuss how data is stored, shared, and deidentified? *NA*

Appendix

Rank Turbulence Divergence (RTD)

Intuition. RTD measures how much two ranked lists disagree, with an emphasis on the highest-ranked items that users are most likely to see. An RTD of 0 means the lists are identical, whereas values approaching 1 signal that the items that rank highly in one list are largely absent (or much lower) in the other. The *signed* RTD we report assigns a positive sign when the top ranks in list R_1 shift downward in list R_2 (and a negative sign for the reverse), allowing us to say which system elevates or demotes content.

Formal definition. Let R_1 and R_2 be two ranked lists of tokens (here, news-source domains) and let $r_{\xi,1}$ and $r_{\xi,2}$ denote the rank of token ξ in those lists. Following Dodds et al. (2023), we first compute an element-level divergence

$$\left| \frac{1}{[r_{\xi,1}]^\alpha} - \frac{1}{[r_{\xi,2}]^\alpha} \right|^{\frac{1}{\alpha+1}} \quad (1)$$

where the control parameter $\alpha \in (0, 1)$ tunes how strongly early ranks are weighted (we use $\alpha = \frac{1}{3}$, a common default that balances head and tail items).

The aggregate RTD is then

$$RTD_\alpha^R(R_1 \parallel R_2) = \frac{1}{N_{1,2;\alpha}} \frac{\alpha + 1}{\alpha} \sum_{\xi \in R_{1,2;\alpha}} \left| \frac{1}{[r_{\xi,1}]^\alpha} - \frac{1}{[r_{\xi,2}]^\alpha} \right|^{\frac{1}{\alpha+1}} \quad (2)$$

where $N_{1,2;\alpha}$ is a normalization factor ensuring $0 \leq RTD_\alpha \leq 1$ and $R_{1,2;\alpha}$ is the union of tokens present in either list.

Concrete example. Consider the domain `msn.com`, ranked $r_{\text{msn},1} = 5$ in Bing News and $r_{\text{msn},2} = 20$ in Google News. With $\alpha = \frac{1}{3}$,

$$\delta_{1/3}(\text{msn}) = |5^{-1/3} - 20^{-1/3}|^{3/4} \approx 0.28.$$

Summing analogous terms over all domains and applying the normalisation yields a *signed* RTD of +0.82 for Bing vs. Google, indicating that Bing surfaces a markedly different—and, on average, higher-ranked—set of domains than Google for the same query. For reference, RTD values in our study range from 0.05 (near-identical rankings) to 0.86 (highly turbulent).

Interpretation scale.

- **0–0.2:** Lists are almost identical.
- **0.2–0.5:** Moderate re-ordering; some head items shift ranks.
- **0.5–1.0:** Substantial turbulence; head items in one list are low-rank or absent in the other.

Prompting Template for Framing Classification

We use the following prompt to classify the given headline into one of the six frames: Conflict, Game/Strategy, Thematic/Issue, Human Interest, Episodic, and Economic. Please find the descriptions to each of the frames in the prompt below.

instruction: You are an assistant trained to classify news headlines into one of the following generic frames based on their dominant focus. Below are the definitions of each frame:

1. Conflict Frame

- Description: Presents events as a conflict between competing actors, issues, or interpretations.

2. Game/Strategy Frame

- Description: Focuses on the efforts of actors to gain support, influence, or achieve specific goals.

3. Thematic/Issue Frame

- Description: Centers on the substantive content of public concerns and issues.

4. Human Interest Frame

- Description: Narrates events from the perspective of individuals affected by the issues or events.

5. Episodic Frame

- Description: Presents specific events or episodes without extensive context or connection to broader themes.

6. Economic Consequences Frame

- Description: Assesses the expected implications of events and policies on the economy.

Task:

Given a news headline, classify it into one of the above frames by selecting the most appropriate

single frame that best represents the headline's primary focus**. Respond only with the name of the frame (e.g., "Conflict Frame"). Additional Notes:

- Select Only One Frame: Assign only the dominant frame that best fits the headline, even if multiple frames seem relevant.

- Consistency: Use the exact frame names provided in the definitions for clarity and consistency.

- Clarity: Ensure that the classification is based solely on the headline's content without requiring external context.

Example:

- Headline: "Local Hero Rescues Family from Burning Building"

- Frame: Human Interest Frame

Now, classify the following headline: Headline: [headline]