

Investigating the Impact of Media Bias on News Readers

Keywords: News, Bias, LLMs, Online Experiment, Persuasion

Extended Abstract




Figure 1: An overview of our experiment pipeline. We collect multiple news articles on a given topic and extract all factual statements from them. Using GPT-4, we write new articles using different subsets of these factual statements and frame the articles in different ways. We then give each article to news readers and gauge its impact on their views about the topic.

How news media present facts surrounding an event can have a significant impact on readers’ perception of them. News sources on opposing sides of the political spectrum may spin their coverage in very different ways, leading to a divergence in the worldviews of their readership and increasing polarization in society. While the phenomenon of media bias has been extensively covered in various social science disciplines, including political science, communications, economics, and sociology [1, 2, 3, 4], the precise measurement of its impact on readers’ opinions, and which factors influence them the most, has received less attention.

In this project, we investigate the influence of this bias in news articles on readers’ perceptions of real world events. We deconstruct this bias into two components: selection and framing. Selection bias relates to which subset of topics or facts are presented to the reader, while framing bias relates to how those topics or facts are presented. To isolate the two components and study their impacts separately, we propose a novel pipeline that utilizes LLMs to extract facts from collections of news stories on a particular event and then generate new articles with precisely controlled amounts of selection and framing bias. This nuanced approach allows for a more detailed understanding of how bias in media presentations can shape public perception, ultimately contributing to the broader discourse on media influence and bias.

Our pipeline consists of multiple stages. First, we compile a dataset of news articles from various sources pertaining to specific events. This focus on particular events allows us to disentangle the impact of selection and framing of facts from the selection of topics. Sampling news articles on a broader subject, such as the economy in general, could result in disparate coverage across sources due to their focus on different aspects of it. For instance, The New York Times (NYT) might emphasize inflation, whereas The Wall Street Journal (WSJ) could prioritize discussing the job market. Focusing on singular events, such as the release of the monthly job market report, or a particular debate or speech in politics, ensures that all sources are addressing the same subject by definition. Additionally, this methodology helps to control for temporal variables that might influence the content of news articles. As a news story

develops and new information come to light, the selection and framing of facts might change over time for completely natural reasons. Therefore, we focus on coverage in the immediate aftermath of a given event by limiting the time window during which we collect articles. To collect a balanced sample of news coverage, we source news articles from a variety of outlets with varying political leanings.

With our collection of news stories in place, we use GPT-4 to extract all factual statements from these articles, and also ask it to label the stance of each statement with respect to the topic (positive, negative, or neutral). From this extensive collection of factual statements, we then select a representative subset, comprising 16 negative and 16 positive statements about the topic. To ensure their accuracy and relevance, the sampled statements and their identified stances were further subjected to verification through human labeling.

Next, we tasked GPT-4 with generating new articles from this set of factual statements under defined bias conditions. Our approach seeks to disentangle the effects of framing from those of selection by holding one variable constant while varying the other. The dimension of framing is managed through specific input prompts directing the model to craft the article with a positive, neutral, or negative tone. On the other hand, the selection dimension is governed by the assortment of facts presented to the model as input, including sets of exclusively 16 positive facts (positive selection), a balanced set of 8 positive and 8 negative facts (mixed selection), and the set consisting solely of 16 negative facts (negative selection). This methodology is visualized in a 3x3 experimental design in Figure 2. It is important to note that scenarios involving negative selection with positive framing and positive selection with negative framing are omitted, as they do not realistically reflect practical conditions. All generated articles were verified by human annotators to ensure their accuracy and consistency.

Finally, we conducted a randomized survey experiment to assess the impact of these biased articles on readers' perceptions of the topic. For the pilot experiment, the topic was confined to the Federal Reserve's announcement of interest rates in late 2023. To minimize experimenter demand effects, we prefaced the experiment with specific instructions highlighting that the article represents a selected perspective on the matter and is not necessarily indicative of the full spectrum of viewpoints. Participants were randomly allocated to one of seven experimental conditions with equal probability. Following their reading of the article, we posed five different questions regarding the economy. To determine the treatment effect, we utilized a regression model for each outcome, with the neutral condition (mixed facts with a neutral tone) serving as the baseline for comparison. Additionally, we factored in demographic indicators and news consumption habits as covariates. Due to the high correlation among the outcomes, we focus on presenting results for the initial outcome question: "How would you rate economic conditions in the United States today?" As depicted in Figure 3, the influence of negative framing and selection (p-value = 0.048) and negative selection alone (p-value = 0.046) were marginally significant at the 0.05 significance level. Moreover, a comparison between these negative conditions and the converse extreme condition, i.e., positive selection and framing, yielded statistically significant differences (p-value = 0.004, p-value = 0.005). These results suggest that even on a subject like the economy, where individuals may hold strong preconceived notions, a single exposure to a slightly biased version of the news can significantly shape opinions. Notably, this effect exhibits asymmetry, with negative bias proving more impactful than positive bias.

This study showcases the effect that selection and framing bias can have on readers' opinions of real world events. By utilizing LLMs to generate articles with varying levels of bias along both dimensions, we are able to investigate how important each component is in influencing news readers. In the future, we plan to expand on this study by scaling it up to cover a larger set of topics and investigate how framing and selection impact each differently.

References

- [1] Robert M Entman. Framing bias: Media in the distribution of power. *Journal of communication*, 57(1):163–173, 2007.
- [2] Tim Groseclose and Jeffrey Milyo. A measure of media bias. *The quarterly journal of economics*, 120(4):1191–1237, 2005.
- [3] Shanto Iyengar and Kyu S Hahn. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of communication*, 59(1):19–39, 2009.
- [4] Sendhil Mullainathan and Andrei Shleifer. The market for news. *American economic review*, 95(4):1031–1053, 2005.

1 Figures

	Positive Selection	Mixed Selection	Negative Selection
Positive Framing	PS&F	PF	
Neutral Framing	PS	N	NS
Negative Framing		NF	NF&S

Figure 2: The experimental design and different treatment conditions varying selection and framing intensity.




Figure 3: Treatment effect (compared with the Neutral condition) on first outcome question: "How would you rate economic conditions in the United States today?". Error bars show associated 95% confidence intervals.