

Public versus Less-Public News Engagement on Facebook: Patterns Across Bias and Reliability

Alireza Mohammadinooshan
Linköping University
Sweden
alireza.mohammadinooshan@liu.se

Niklas Carlsson
Linköping University
Sweden
niklas.carlsson@liu.se

Abstract

The rapid growth of social media as a news platform has raised significant concerns about the influence and societal impact of biased and unreliable news on these platforms. While much research has explored user engagement with news on platforms like Facebook, most studies have focused on publicly shared posts. This focus leaves an important question unanswered: how representative is the public sphere of Facebook’s entire ecosystem? Specifically, how much of the interactions occur in less-public spaces, and do public engagement patterns for different news classes (e.g., reliable vs. unreliable) generalize to the broader Facebook ecosystem?

This paper presents the first comprehensive comparison of interaction patterns between Facebook’s more public sphere (referred to as *public* in paper) and the less public sphere (referred to as *private*). For the analysis, we first collect two complementary datasets: (1) aggregated interaction data for all Facebook posts (public + private) for 19,050 manually labeled news articles (225.3M user interactions), and (2) a subset containing only interactions with public posts (70.4M interactions). Then, through discussions and iterative feedback from the CrowdTangle team, we develop a robust method for fair comparison between these datasets.

Our analysis reveals that only 31% of news interactions occur in the public sphere, with significant variations across news classes. Engagement patterns in less-public spaces often differ, with users, for example, engaging more deeply in private contexts. These findings highlight the need to examine both public and less-public engagement to fully understand news dissemination on Facebook. The observed differences hold important implications on content moderation, platform governance, and policymaking, contributing to healthier online discourse.

CCS Concepts

• Information systems → World Wide Web.

ACM Reference Format:

Alireza Mohammadinooshan and Niklas Carlsson. 2025. Public versus Less-Public News Engagement on Facebook: Patterns Across Bias and Reliability. In *Proceedings of the 17th ACM Web Science Conference 2025 (Websci '25), May 20–24, 2025, New Brunswick, NJ, USA*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3717867.3717912>



This work is licensed under a Creative Commons Attribution 4.0 International License. *Websci '25, May 20–24, 2025, New Brunswick, NJ, USA*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1483-2/25/05
<https://doi.org/10.1145/3717867.3717912>

1 Introduction

Social media has emerged as an increasingly important news source, with Facebook maintaining a prominent position among social platforms [8, 31]. For example, according to a 2024 Pew Research Center report [8], 33% of U.S. adults regularly get news on Facebook, similar to YouTube (32%), and notably ahead of Instagram (20%) and Twitter (12%). This extensive reach highlights the potential impact of news sharing on Facebook in shaping both individual opinions and broader societal discourse [39].

Unfortunately, not all news are reliable. With an increasing amount of misinformation being circulated on Facebook, it is crucial to understand how users engage with news of varying reliability [12]. Understanding of these dynamics are expected to benefit media researchers, journalists, content moderators, and policymakers, whose choices based on these insights, in turn, are expected to influence public opinion and the behavior of regular users [20, 30].

Previous research has extensively examined public user engagement dynamics with different news content and the factors affecting this engagement [2, 7, 10, 26, 42]. However, with a growing privacy inclination of Facebook users and a rise in user engagement in private spaces of Facebook [9, 24], it is becoming increasingly important to understand these dynamics also with regards to private sharing. Yet, the current literature is centered around more public posts, leaving gaps in our understanding of engagement differences between these realms. Most importantly, prior research has not studied the dynamics of public vs. private interaction.

To address this gap, in this paper we present the first comprehensive comparison of the news article sharing and user interaction patterns seen on Facebook’s more public sphere (referred to as *public* in paper) versus in the less public sphere (for simplicity referred to as *private*). Specifically, we investigate differences in how news articles written with different *bias* and *reliability* are shared and engaged with, as well as the depth of these interactions. In this context, *bias* refers to a tendency for news articles to exhibit partiality or favoritism towards particular groups or ideas (e.g., left or right on the political spectrum), while *reliability* refers to the accuracy and credibility of the information presented (e.g., fake or true). While most previous research on news engagement has primarily focused on either bias or reliability (and on public engagement), we consider both dimensions, as they have been found related, but yet each has its own ability to influence news article sharing behaviors and affect the quality and diversity of information that users encounter on social media [15, 34].

Furthermore, by analyzing whether interactions are *deep* or *shallow*, we can determine if users engage more deeply with certain content types (e.g., fake news) in private or public spaces. This distinction is crucial, as deeper engagement—such as sharing or commenting—can significantly enhance an article’s impact compared to shallow actions like pressing Facebook’s “like” button [31].

Research questions and methodology: Guided by the goal of better understanding how representative the public sphere of Facebook’s ecosystem (typically studied in prior works) is of the private sphere and the ecosystem as a whole, we designed our study to address the following research questions:

- RQ1** How do the patterns of public and private sharing of news articles on Facebook differ for articles with varying levels of bias and reliability?
- RQ2** If and how does the depth of user interactions with news articles differ within public and private sharing contexts for articles with varying levels of bias and reliability?

We took several steps to address these previously unaddressed questions in as controlled matter as possible. First, in contrast to most prior works that use publisher-level labeling (e.g., Adelson et al. [10]; Horne et al. [18]), we employ article-level labeling of bias and reliability. By doing so, we capture that not all articles by a publisher have the same bias or reliability, recognizing that the publisher (source) is only one factor [36] in the bias/reliability. In particular, we consider the bias and reliability of articles ($N = 19,050$), meticulously selected from a substantial pool of publishers ($N = 1,121$) and 30K+ manually labeled news articles. Here, the original labeling was provided by Ad Fontes Media [1], which provides bias and reliability of news articles. Second, through active back-and-forth discussions and feedback with the Crowdtangle team, we developed and implemented a methodology using their Chrome add-on that allows us to obtain simultaneous interaction statistics for a representative sample of publicly shared Facebook posts linking these articles, as well as across all Facebook posts (including both public and private posts) linking these articles. Finally, using this unique dataset, we perform a comparison of interaction dynamics when users share news articles with different levels of bias and reliability, both publicly and privately.

Empirical example findings: Our analysis reveals several key insights into how users interact with news articles of varying bias and reliability on Facebook. For instance, we show that users tend to engage more deeply in private discussions than in public ones, irrespective of the news class. When considering the news class, we highlight that users exhibit relatively higher deep interaction levels with highly-unreliable content. Our results also show that reliable news content has significantly lower private interaction shares compared to the highly-reliable or even unreliable content.

Example beneficiaries: Our methodology and findings contribute to a better understanding of news sharing and interaction dynamics on Facebook, offering valuable insights for various stakeholders. Media researchers and journalists can use our analysis of how bias and reliability shape engagement in public and private spheres for improved content creation and distribution strategies. For content moderators, the analysis offers data-driven guidance on prioritizing efforts to curb the spread of problematic content. For

policymakers, we highlight how privacy settings influence engagement patterns, with highly unreliable news, for example, garnering more engagement in the private sphere. These results underscore how platform design choices can influence selective exposure and engagement, potentially exacerbating issues like ideological polarization and misinformation spread.

Roadmap: Sect. 2 explains our research design, including news article selection, labeling, data collection, and the processing steps taken to ensure fair comparisons. The resulting dataset is summarized in Sect. 3. Sect. 4 analyzes public vs. private dynamics both from a high-level aggregated perspective (Sect. 4.1) and then using a detailed statistical analysis (Sect. 4.2). In Sect. 5, we turn our attention to differences in the depth of interactions. Topic and top-publishers analysis are provided in Sects. 6 and 7. Finally, related works (Sect. 8) and conclusions (Sect. 9) are presented.

2 Research Design, Data Collection and Limitations

At a high level, our research design has three parts.

Part 1. Article selection and labeling: We first obtained bias and reliability scores for the articles evaluated by Ad Fontes Media [1]. Subsequently, each article was categorized into one of three bias classes and one of four reliability classes based on these scores.

Part 2. Collection of interactive data: After careful preprocessing of the URLs, we used the CrowdTangle browser extension to obtain interaction data for two sets of posts linked to these URLs (1) *public* and (2) *combined* (public + private). To address some limitations of the API and to ensure fair comparison of the sets, we leave a four-month gap between the latest labeled article (Nov. 1, 2022) and the primary data collection (Mar. 2023), as well as apply some additional post-processing (e.g., examining the actual posts), filtering (based on thresholds determined via discussions with the CrowdTangle team), and collect some complementing data directly from CrowdTangle (to address limitations of the extension API).

Part 3. Data analysis: Finally, we use the final dataset (capturing the two sets) to compare the properties of the *public* vs. *combined* (private + public) sets and identify significant statistical differences between the interaction patterns of public and private posts. We next provide details of the initial two steps, before presenting our analysis results (Part 3) in the subsequent sections.

2.1 News Article Selection and Labeling

Selection of articles: There are several independent initiatives that assess the bias and/or reliability of individual news articles and/or news sources. Examples of such evaluation efforts include Media Bias Fact Check [27], Ad Fontes Media, AllSides [3], and NewsGuard [32]. Among these, we opted to use data from Ad Fontes Media for the following main reasons: (1) it has been widely used in previous research [16, 19, 22], (2) they evaluate individual news articles (not only the news publishers), (3) each assessed article receives a score for both bias and reliability, (4) they offer a transparent evaluation methodology, which is published and explained in a white paper [33], and finally (5) the dataset comprises a large set of news contents randomly selected from a diverse set of publishers.

We obtained all evaluated news from Ad Fontes Media as of Nov. 1, 2022. This included 31,446 news article URLs from 1,121 publishers. After filtering out URLs serving as event reporting pages (e.g., <https://www.nola.com/news/hurricane>) and articles with updated ratings (per Ad Fontes Media’s documentation), the final dataset included 31,408 articles.

Labeling of articles: Each article in the Ad Fontes Media dataset is assessed for both bias and reliability by a minimum of three human analysts, representing a mix of right, left, and center self-reported political perspectives. Bias scores from Ad Fontes Media range between -42 and $+42$, where more negative values suggest a stronger left-ward bias, and positive values indicate a right-leaning bias. As for reliability, scores vary from 0 to 64, with 64 representing the most-reliable news.

For our analysis, we categorized the bias and reliability scores into distinct classes. While previous work mainly focuses on binary classification of news (e.g., Left vs. Right or Fake vs. True), we opted for a more granular approach by defining three classes of bias (*Left*, *Right*, and *Center*) and four classes of reliability (*Most-unreliable*, *Unreliable*, *Reliable*, and *Most-reliable*). Figures 1a and 1b show the Cumulative Distribution Function (CDF) of the bias and reliability scores of the articles included in our analysis (after additional filtering steps, explained later in this section) together with the threshold values and class labels used to define each class.

With our labeling, anything with a bias score below -6 is called *Left*-biased, and anything above 6 is called *Right*-biased, with the range $(-6,6)$ capturing the *Center* class. We note that Ad Fontes Media also refers to this range as “Middle or Balanced Bias” [1]. As observed, this class represents most samples in our dataset, representing 58% (11,094 out of 19,050 articles in the final dataset). With this split, the two biased classes have similar, close to 20%, shares.

For the reliability classes, we assigned the *Most-unreliable* class to cover the lowest range (0, 32). According to Ad Fontes Media’s terminology [1], these values typically correspond to articles “which contain inaccurate and misleading info, selective stories or have opinion and wide variation in their reliability”. Similarly, the range (48-64), chosen for the *Most-reliable* class, encompasses articles characterized by “thorough and original fact reporting” due to Ad Fontes terminology. Finally, the threshold of 40, between *Reliable* and *Unreliable*, matches the mid-point between these two classes as well as where Ad Fontes make their split between “wide variation in reliability” and “mix of fact reporting and analysis”.

2.2 Collection of Interaction Data

Preprocessing of URLs: Before collecting the interaction data associated with each article (URL), we (1) expanded link shorteners and (2) converted URLs to canonical forms. During the conversion process, we primarily removed unnecessary URL parameters after “?” except those essential to canonical forms (e.g., “id” parameters are sometimes part of the canonical form). We also replaced web archive links in the Ad Fontes Media’s dataset with original links.

Primary data collection for the two sets of posts: Prior to its shutdown in Aug. 2024, CrowdTangle (owned by Facebook) provided access to interaction statistics for public posts, which it indexed in a database widely used by researchers. Additionally,




Figure 1: CDFs of scores for (a) bias and (b) reliability.

their Chrome extension offered access to the life-time interaction statistics available at the time, encompassing all posts (public + private). In this study, we utilized the extension to collect statistics for two datasets: *combined* and *public*, where the *combined* dataset includes all interactions from all posts (public + private), and the *public* dataset represents a subset of the combined dataset, containing only interactions with publicly accessible posts.

To collect the data for each URL, we used the “Download” option within the extension’s interface when browsing each article. The downloaded CSV file contained the interaction data for both sets. As expected from a privacy standpoint, only the aggregated statistics are provided for the *combined* set. In contrast, every post and its individual interaction data are provided for the public set.

At a high level, the *combined* category aggregates the interactions (e.g., likes, shares, comments, etc.) of all Facebook posts referencing the article’s URL, regardless of the post’s privacy settings. This includes interactions on posts with limited privacy settings, such as “Friends only” or “Only Me”. In contrast, the *public* category, includes all interactions with the public posts tracked by CrowdTangle, which, due to heavy-tailed characteristics and significant coverage, captures most interactions with public Facebook posts. As supporting examples, all posts of all US-based public groups with 2K+ members are indexed, and in 2021 over 99% of all Facebook pages with at least 25K likes were indexed [38].

Timeline and limitations: We collected the interaction data in Mar. 2023, ensuring at least a four-month gap between the publication date of the newest article in our dataset and our interaction data collection. This interval ensured capturing most posts and interactions for the URLs. This four-month window is notably conservative, as research has shown that most interactions occur within days of publication [35, 40]. Although this conservative window may have overlooked interactions with content that was later removed or deleted, we are unable to assess the extent of this effect—if it exists—due to CrowdTangle’s shutdown in Aug. 2024.

As mentioned in [11], the information from the CrowdTangle add-on originated from two different sources: (1) public interaction data is obtained from the CrowdTangle database, while (2) combined interaction data is derived from Facebook’s Graph API. Consequently, variations in data capturing times might cause some URLs to exhibit higher public than combined interactions.

Furthermore, as confirmed by the CrowdTangle team, Facebook’s aggregation data restarts the interaction count whenever websites update their connection schemes (e.g., switching from HTTP to HTTPS) after publishing a URL, while CrowdTangle continued accumulating data. Although the discrepancy is not substantial (9% in our initial dataset), we aimed to mitigate the effect of these differences where possible (described next).

Table 1: Statistics based on Bias and Reliability.

	Class	Articles no	Combined interactions	Public interactions
Bias	Left	4,516	79,038,246	20,307,107
	Center	11,094	90,084,398	32,109,924
	Right	3,440	56,160,101	18,011,285
Reliability	Most-unreliable	3,520	45,299,475	13,206,643
	Unreliable	3,795	46,712,927	14,199,322
	Reliable	9,308	89,596,137	30,639,973
	Most-reliable	2,427	43,674,206	12,382,378
Total		19,050	225,282,745	70,428,316

Further filtering and data processing for fairer comparisons: After consulting with the CrowdTangle team regarding all numerical discrepancies that we observed, we identified two additional reasons for the discrepancies. First, as stated in [28], for privacy reasons, for the combined data (that include private posts) “the values are intentionally not precise, but you can be confident they accurately reflect user engagement with a URL.” Discussing the identified discrepancy for such cases, the CrowdTangle team confirmed that combined interaction data with counts below 100 is not reliable; however, that data with counts above 100 provides a reliable estimate of public + private interactions. For this reason, we removed all URLs in our dataset with combined interactions below 100. Although this reduced our dataset to 19,505 articles, it is important to note that this group of URLs accounted for less than 4% of public interactions in our dataset. Furthermore, we note that the public interaction data for this group still remains reliable.

Second, the CrowdTangle team informed us that the Facebook Graph API only considers posts with the URL attached to the post (i.e., those displaying a preview of the URL) when counting interaction with the URL, while CrowdTangle’s algorithm includes any post containing a linked URL, regardless of attachment status. Taking this into account, in our final step toward enhancing interaction data quality, we chose not to rely on the aggregated data from the CrowdTangle extension (which sums up the interactions of all posts mentioned in the extension). Instead, we examined each post to determine whether the related URL was attached or not, and if so, we included them in the sum of public interactions for that URL.

This resulted in a more accurate comparison between combined and public interactions. After these processing steps, we had only 455 URLs for which public interactions exceeded combined interactions, which were excluded from our final dataset. The primary causes of these discrepancies were beyond the scope of our research to address (e.g., changes in URL, publisher protocol schemes, or different data capturing times as mentioned above).

Enhancing the data with additional posts: Another limitation of the CrowdTangle extension was its restriction to retrieving data for up to 500 posts per URL. To address this, we retrieved additional data from CrowdTangle for URLs with over 500 posts in the *public* interaction category. While some of the pre- and post-processing steps described above require significant effort, they help ensure an accurate and fair comparison between the public and combined sets in such a way that we can provide conclusive insights into the relative sharing patterns of public vs. private posts.

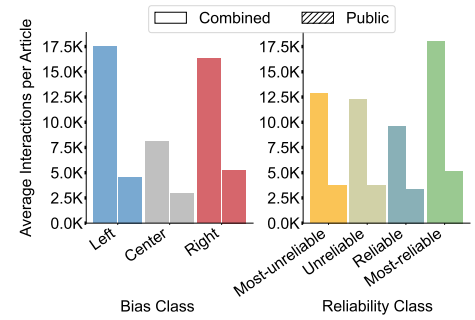


Figure 2: Interactions per article

3 Dataset

High-level summary: After applying the aforementioned filtering steps, we have in total 19,050 articles from 1,121 news outlets remaining in our dataset. These articles have been shared in 253,350 posts, which combined are responsible for 225,282,745 interactions (out of which 70,428,316 are public).

Bias data: Table 1 shows the number of articles and the total combined and public interactions for each bias and reliability class. With the selected bias thresholds, our dataset includes 4,516 *Left*-biased articles and 3,440 *Right*-biased articles, with the remaining 11,094 articles falling in the *Center* category. These articles are in turn responsible for 79,038,246 (*Left*), 56,160,101 (*Right*), and 90,084,398 (*Center*) interactions. While the *Center* category includes 58% of the articles, it is interesting to note that it is responsible for a significantly smaller fraction of the total interactions (40%). Instead, the *Left*-biased and *Right*-biased articles see relatively higher interaction rates (studied in the next section). For example, the *Left*-biased articles account for only 24% of the articles but 35% of the interactions, and the *Right*-biased articles are responsible for only 18% of the articles but 25% of the interactions.

Reliability data: For the reliability classes as shown in Table 1, the relative differences are smaller. Here, the *Reliable* articles make up the largest share (9,308 articles and 89,596,137 interactions), followed by *Unreliable* (3,795 articles and 46,712,927 interactions), *Most-unreliable* (3,520 articles and 45,299,475 interactions), and *Most-reliable* (2,427 articles and 43,674,206 interactions).

4 Public vs. Combined Interactions

To understand how engagement patterns differ between public and private spheres on Facebook, we present two complementary analyses: (1) an aggregated comparison of public vs. combined (public + private) interactions (Sec. 4.1), and (2) a statistical analysis of their differences (Sec. 4.2). By comparing public interactions against combined interactions, where private interactions constitute the majority (as we show later), we can effectively study how engagement in private spaces differs from public ones.

4.1 Aggregated (Macro) Analysis of Interactions

Figure 2 presents the number of combined and public interactions per article for each class. Here, two metrics are used:

Average combined interactions: (left bars for each class) represents the average number of combined (private+public) interactions per article in each class. More specifically, considering $I_{c,i}^{\text{comb}}$ as the combined interaction for the i^{th} article in class c , the average combined interactions for this class (\bar{I}_c^{comb}) is calculated as:

$$\bar{I}_c^{\text{comb}} = \frac{\sum_{i=1}^{N_c} I_{c,i}^{\text{comb}}}{N_c}, \quad (1)$$

where the numerator ($\sum_{i=1}^{N_c} I_{c,i}^{\text{comb}}$) corresponds to the “Combined interactions” column in Table 1 and the denominator (i.e., the total number of articles N_c for class c) is found in the “Articles no” column of Table 1. For example, for the *Center* class in the *Bias* category, the average combined interactions per article is computed by dividing the total interactions (90,084,398) by the articles count (11,094), yielding $\bar{I}_{\text{Center}}^{\text{comb}}$ to be 8,120 combined interactions per article (gray bar in Figure 2).

Average public interactions per article: (right bars) is defined in a similar manner but considering only the public interactions. More specifically, considering $I_{c,i}^{\text{public}}$ as the public interactions for the i^{th} article in a class c , the average public interactions for this class ($\bar{I}_c^{\text{public}}$) is calculated as follows:

$$\bar{I}_c^{\text{public}} = \frac{\sum_{i=1}^{N_c} I_{c,i}^{\text{public}}}{N_c}, \quad (2)$$

where the numerator (i.e., $\sum_{i=1}^{N_c} I_{c,i}^{\text{public}}$) corresponds to the “Public Interactions” column in Table 1. Now, taking the *Center* class as an example again, the average public interactions per article is calculated by dividing the total public interactions (32,109,924) by the article count (11,094), resulting in an $\bar{I}_{\text{Center}}^{\text{public}}$ of 2,894 interactions per article (striped gray bar in Figure 2).

As seen in the figure, there are some very clear and interesting differences in the trends observed for the public vs. the combined statistics. First, for all three bias classes and for all four reliability classes, the interaction rates are significantly lower for the public posts than for the combined set. These aggregate rate differences suggest that users are more likely to interact with private posts, regardless of bias and reliability class, and underscores the importance of considering the total (combined) interactions, not only the (typically studied) public posts. We next look closer at the relative interaction levels of each class and the impact of using the public vs. combined sets for such interaction comparisons.

Bias comparisons: Second, examining the relative differences among the bias classes, the *Left* bias class has the highest average combined interactions per article (17,501), followed by the *Right* (16,325) and *Center* (8,120) classes. For public interactions; however, the *Right* bias class has the highest average public interactions per article (5,236), with the *Left* (4,497) and *Center* (2,894) classes trailing behind. This shows that biased articles see even greater relative interaction rates in private than in public posts.

Reliability comparisons: Third, for the reliability classes, the *Most-reliable* class has the highest average combined and public interactions per article (17,995 and 5,101). The *Unreliable* and *Reliable* classes have relatively similar average public interactions per article (3,741 and 3,291, respectively). Moreover, the *Most-unreliable* class has a higher average combined interactions per article (12,869)




Figure 3: Aggregated public interactions shares

and public interactions per article (3,751). It is interesting to see that the two extreme classes (i.e., the *Most-reliable* and *Most-unreliable*) have the highest average interactions per article in both combined and public contexts, implying that users are more likely to engage with articles on both ends of the reliability spectrum.

Aggregated public interaction share: Motivated by most prior works only studying the public posts tracked by CrowdTangle, we next evaluate what fraction of the total interactions this set captures and compare the relative differences in the public interaction share of different categories. For this analysis, we calculate the (aggregate) public interaction share as the ratio of public interactions to total interactions for each class c : $\bar{I}_c^{\text{public}} / \bar{I}_c^{\text{comb}}$.

Figure 3 summarizes these ratios. We note that the *Center* class has the highest aggregated public interaction share, and that among the biased groups, there is a clear gap between the left and right parties, as the *Right* class has a significantly (25% extra) higher public interaction share (lower private interaction share). When it comes to the reliability classes, the relative gap between the classes is smaller. However, it should be noted that the two extreme classes (*Most-reliable* and *Most-unreliable*) have the lowest (and almost similar) aggregated public interaction shares.

Takeaway: Compared to right-biased articles, left-biased articles receive higher private interactions share. In terms of reliability, extremely reliable/unreliable news articles receive higher levels of private interactions share.

4.2 Granular Statistical (Micro) Analysis of Interactions

Having presented aggregated analysis and insights, we next present a more detailed statistical comparison between different classes in which we give equal consideration to every article in a class, regardless of its total interaction count. For this analysis, we utilize the distributions of the per-article level interactions for each class c : $\mathcal{D}_c^{\text{comb}} = \{I_{c,i}^{\text{comb}}\}$ and $\mathcal{D}_c^{\text{public}} = \{I_{c,i}^{\text{public}}\}$, where $I_{c,i}^{\text{comb}}$ is the total number of interactions across all posts associated with the i^{th} article of class c ($1 \leq i \leq N_c$) and $I_{c,i}^{\text{public}}$ is the corresponding number of interactions with public posts.

The next two subsections compare distributions of the total (combined) interactions (Sect. 4.2.1) and study the distributions of the relative public interaction shares (Sect. 4.2.2), respectively. In both subsections, the distributions are compared using multiple

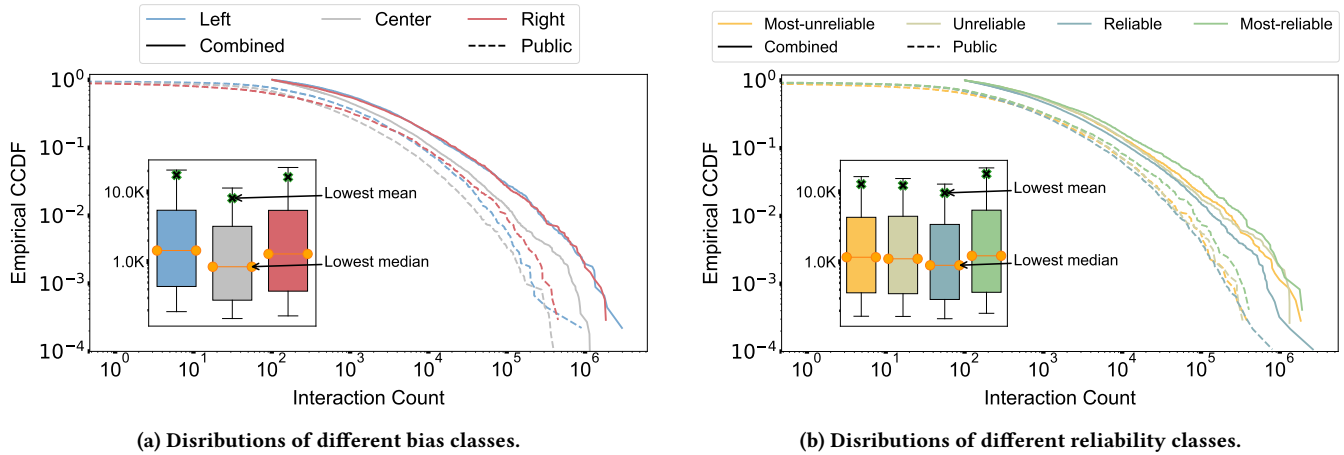


Figure 4: Empirical complementary cumulative distribution functions (CCDFs) for combined and public interactions, along with inset boxplots showing the distribution of combined interactions. Both CCDFs emphasize tail characteristics of distributions, while the boxplots provide granular statistical summaries, including medians, quartiles, and means.

statistical tests; statistical significance is reported if a p-value is below 0.001, and supporting p-values are reported for the main findings.

4.2.1 Distribution of Combined Interactions. Comparison of Bias Classes: Figure 4a shows the CCDFs of both the set of combined interactions ($\mathcal{D}_c^{\text{comb}}$) and the set of public interactions ($\mathcal{D}_c^{\text{public}}$) for each bias class c on a log-log scale. Accompanying the CCDF is a boxplot, highlighting key statistical percentiles for the distribution of the combined interactions ($\mathcal{D}_c^{\text{comb}}$). Specifically, the 10th percentile (bottom marker), 25th percentile (bottom of box), median (middle marker), 75th percentile (top of box), 90th percentile (top marker), and the mean (i.e., \bar{i}_c^{comb}) (circle). To allow higher resolution, outliers are not shown in the boxplot.

Upon examining the figure, we can make several observations. First, as articles with fewer than 100 combined interactions were excluded, the minimum value for $\mathcal{D}_c^{\text{comb}}$ starts at 100. Second, the shape of the of the CCDFs when plotted on log-log scale underscores the “heavy-tailed” nature of the interactions ($\mathcal{D}_c^{\text{comb}}$ and $\mathcal{D}_c^{\text{public}}$) of all bias classes c . This suggests that a small subset of the articles are responsible for a most of the interactions. Third, there are discernible variations in the distributions. For example, articles classified as biased (both *Left* and *Right*) consistently see higher interaction levels (i.e., right-shifted curves compared to the *Center* class), echoing the trends we observed in the aggregated analysis. To better understand these disparities and their implications, we next provide a more rigorous statistical analysis.

Statistical Tests for Bias-related Differences: To evaluate the differences in distributions of the set of combined interactions $\mathcal{D}_c^{\text{comb}}$ across all bias classes c , we employed a multi-step statistical testing approach:

1. Overall Differences: We first used the Kruskal-Wallis test to identify overall differences between the distributions between the bias classes. This test revealed significant disparities among the classes, with a p-value of $1.66 \cdot 10^{-65}$.

2. Pairwise Comparisons: Second, we used the Dunn test to identify specific pairs of classes that exhibited differences. This test confirmed that the two biased classes *Left* and *Right* are different than the *Center* class, but the difference between the two biased classes themselves was not significant. Here, the most significant difference were observed between the *Left* and *Center* class (i.e., $\mathcal{D}_{\text{Left}}^{\text{comb}}$ and $\mathcal{D}_{\text{Center}}^{\text{comb}}$), registering a p-value of $2.06 \cdot 10^{-55}$.

3. Median Comparisons: Third, we used the Mann-Whitney U test to compare the medians of the three classes: *Left* (1,440), *Center* (842), and *Right* (1,286). Also, here, we observed statistically significant pairwise differences when comparing each of the two biased classes (*Left* and *Right*) with the *Center* class, but not between each other. Here, the two significant cases obtained p-values of $8.4 \cdot 10^{-57}$ and $9.1 \cdot 10^{-28}$, respectively.

4. Comparing Means: Finally, due to the violations of assumptions intrinsic to a t-test (namely, normality and homogeneity of variances), we instead use bootstrapping with 100K iterations and a 99% confidence interval to compare the means. Again, the pairwise differences in the means between *Left* vs. *Center* and between *Right* vs. *Center* are statistically significant, but not between the two biased classes themselves (i.e., *Left* vs. *Right*).

Takeaway: Statistical analysis further strengthens the observation that biased news consistently garners higher combined engagement levels per article than center-aligned news. Interestingly, there is no statistically significant difference between the *Left* and *Right* class.

Comparison of Reliability Classes: Figure 4b shows the CCDFs for the different reliability classes, broken down for both combined and public interactions, accompanied by a boxplot offering a clearer perspective on the percentile values for the combined interactions. The CCDFs again demonstrate heavy-tailed distribution characteristic, and the relative shifts of the distributions are consistent with the aggregate (average values) previously discussed (and seen in Figure 2), with the *Most-reliable* class typically getting the most




Figure 5: Distribution of public interaction ratios

interactions, followed by the unreliable classes, and subsequently by the *Reliable* class. While comparing the tail parts is straightforward, comparing the whole distributions requires some care. For example, here, the core of the boxplot does not distinctly set apart the *Most-reliable* from the unreliable classes. To derive robust conclusions, we, therefore, again performed a sequence of statistical tests on the combined interactions distributions.

First, the Kruskal-Wallis test was applied to discern differences among the reliability classes. While we observed statistical differences (p-value of $7.9 \cdot 10^{-20}$), the magnitude of distinction among reliability classes is more nuanced than in the bias classes (which had a p-value of $1.66 \cdot 10^{-65}$). Further analysis employing the Dunn’s test confirmed that the *Reliable* class statistically diverges from the other three. The least significant p-value here, $7.63 \cdot 10^{-8}$, is attributed to the comparison between the *Reliable* and *Unreliable* classes. The most significant, on the other hand, emerges from the *Reliable* versus *Most-reliable* comparison with the p-value of $7.05 \cdot 10^{-15}$. Except for these distributions, none of the other pairwise distribution comparisons resulted in statistically significant differences. Third, the Mann-Whitney U test reinforced our findings, particularly highlighting the statistical significance when comparing the medians of the *Reliable* class (median of 891) with the others: *Most-unreliable* (1,162), *Unreliable* (1,104), and *Most-reliable* (1,218). Finally, we apply bootstrapping to compare the means. While only the *Reliable* and *Most-reliable* classes comparison is significant at the 99% confidence level, we note that all three comparisons against the *Reliable* class are significant at the 95% confidence level.

Takeaway: Comparing to the other reliability classes, the *Reliable* articles demonstrate lower interaction values, especially when compared to the *Most-reliable* articles.

4.2.2 Distribution of Public Interaction Shares. The preceding analysis centered on the distribution of combined interactions. Like the aggregated analysis, we next shift our focus to the proportion of interactions that are public. Specifically, for each class, we analyze the distribution of public interaction shares at the granularity of individual articles. For this analysis, we define the set of public interaction ratios (*PIR*) of class c as $\mathcal{D}_c^{PIR} = \left\{ \frac{I_{c,i}^{public}}{I_{c,i}^{comb}} \right\}$, where $1 \leq i \leq N_c$ is the index of articles included in class c .

Figure 5 presents the public interaction ratios using box-plots broken down per bias and reliability category. Here, we again show the 10th percentile, 25th percentile, median, 75th percentile, and 90th percentile, as well as the average. Like for the aggregate analysis, we observe some notable differences, which we can now support using statistical analysis. First, the Kruskal-Wallis test shows that there are statistically significant differences between both the biases (p-value of $2.68 \cdot 10^{-22}$) and reliability classes (p-value of $8.74 \cdot 10^{-33}$).

Motivated by the higher significance (smaller p-value) for the reliability classes, we consider these differences first. Using the Dunn test, we find statistically significant differences for all pairwise distribution comparisons except between the *Most-unreliable* and *Most-reliable* classes. Similarly, the Mann-Whitney U test shows that all pairwise relative differences in the medians (observed in Figure 5), except for between the *Most-unreliable* and *Most-reliable* classes, are significant (with the largest p-value among the these pairwise cases being $2.82 \cdot 10^{-14}$). Finally, bootstrapping confirms that all pairwise differences in the means observed in the figure are statistically significant, except for the case of the *Most-unreliable* vs. the *Most-reliable* class. It is also worth noting that our observation that the two extreme classes have lower public interaction ratios (higher private interaction ratios) aligns with our findings from the aggregated analysis.

Takeaway: The two extreme classes, *Most-reliable* and *Most-unreliable*, exhibit similar private interaction ratios that are statistically higher compared to the other classes.

We now shift our attention to the public interaction ratios across the bias classes, where we make some interesting observations regarding the classes’ relative order. Specifically, we now observe the first difference in the relative order of the classes when comparing the average (and median) public shares calculated on a per-article basis (Figure 5) with the aggregated results (Figure 3). More specifically, comparing with the aggregate results, the *Right* class has switched ranking with the *Left* class, becoming the one with smallest public interaction share. Furthermore, comparing the distributions seen for the *Right* and *Left* class is statistically significant: Dunn test (p-value of 3.25×10^{-14}), Mann-Whitney U test (p-value of 6.16×10^{-17}) and bootstrap results being significant at 99% level.

While these results may appear surprising at first, they are due to some interesting differences in the distributions. First, note that the aggregate analysis only calculates the overall ratio $\frac{\sum_{i=1}^{N_c} I_{c,i}^{public}}{\sum_{i=1}^{N_c} I_{c,i}^{comb}}$ compared to the granular analysis, where the average (for example) is calculated as $\frac{1}{N_c} \sum_{i=1}^{N_c} \frac{I_{c,i}^{public}}{I_{c,i}^{comb}}$. Second, taking a closer look at the distributions, \mathcal{D}_{Left}^{PIR} and $\mathcal{D}_{Right}^{PIR}$ differ substantially. Part of this can be observed in Figure 4a, where we can see that *Left* class demonstrates notable disparities between public and combined interactions in the tail, while the *Right* class exhibits relatively larger differences in the head of the distribution. This suggests that articles with higher combined number of interactions in the *Right* class may exhibit higher normalized public interaction shares,




Figure 6: Deep interaction ratios (*DIRC* and *DIRP*) for different bias classes and a baseline.




Figure 7: CDFs of *DIRC* (combined) and *DIRP* (public) for different bias classes.




Figure 8: CDFs of *DIRC* (combined) and *DIRP* (public) for different reliability classes.

and articles with few combined interactions may exhibit relatively smaller public shares (compared to the *Left* class).

To confirm this conjecture and substantiate our claim that the above differences are due to aggregated statistics being more affected by larger public interaction shares associated with the tail, we compared the public interaction ratios for articles associated with the tail and head of the distributions of the combined number of interactions. For this analysis, we split the total sets (\mathcal{D}_{Left}^{PIR} and $\mathcal{D}_{Right}^{PIR}$) into two to five equally sized subsets (based on the total combined number of interactions $I_{c,i}^{comb}$ associated with each entry) and then compared the medians and averages of the public interaction ratios seen in the first and last bucket of the *Right* class with the corresponding values seen for the *Left* class. In all cases, the *Left* dominates the *Right* for the head and is dominated by the *Right* for the tail. For example, with three equal-sized buckets, the two classes have median ratios of 0.34 vs. 0.13 for the first bucket (head) and 0.39 vs. 0.46 for the last bucket (tail). Furthermore, we observed statistically significant differences (99% confidence level) as per the above conjecture with both the Mann-Whitney U test (medians) and Bootstrap tests (averages), validating these differences. This confirms that the articles in the *Right* class with the highest interactions have elevated normalized public interaction ratios, and those in the *Left* class with the lowest interactions have similarly elevated ratios.

Takeaway: In terms of bias, the *Right* class sees the smallest public interaction share among the least popular content to interact with, while the *Left* class sees relatively less overall public interaction share (due to its most popular articles to interact with not seeing as big public interaction share).

5 Deep vs. Shallow Interactions

In this section, we investigate the prevalence of *deep* vs. *shallow* interactions and explore their variations among different reliability and bias classes. For this analysis, we categorize all (emoji-based) reactions (e.g., likes, loves, sads) as shallow interactions while comments and shares are considered as deep interactions. This distinction, supported by the previous works [2, 21], is motivated by the idea that comments and shares typically involve more cognitive effort and engagement from users. For example, comments involve formulating thoughts and opinions, while shares actively endorse and disseminate content.

To compare the depth of interaction seen for subsets of public and private posts, we define the following two per-article metrics:

- *Deep Interactions Ratio for Combined (DIRC)*: For each article, this ratio measures the proportion of the combined interactions that are classified as deep interactions.
- *Deep Interactions Ratio for Public (DIRP)*: For each article, this ratio measures the proportion of the public interactions that are classified as deep interactions.

To compute *DIRC*, we divide the number of deep interactions by the total interactions for each URL. For *DIRP*, we first identify the deep interactions (shares and comments) occurring on public posts associated with the URL. Summing these deep interactions across all public posts and dividing by the total number of public posts’ interactions for the URL yields the *DIRP*. For example, an article with 1,000 total interactions, 500 of which are deep, has a *DIRC* of 0.5. If 400 interactions are public, and 100 of those are deep, the *DIRP* is 0.25. We next analyze the distribution of *DIRC* and *DIRP* values for all articles in each class.

Comparisons of bias classes: We first compare the deep interaction patterns of the bias classes. Figure 6 shows the distributions of *DIRC* and *DIRP* for the different classes. As a baseline, we also include the distributions for the general population (rightmost box pair), considering all samples.

First, referring to the figure, for each bias class, we observe substantial differences between the distributions of *DIRC* and *DIRP*. For example, for every percentile and for the means, the *DIRC* distributions (combined) have significantly larger values than the *DIRP* distributions (public), indicating that users engage in more substantive interactions in less-public settings. This suggests that deeper engagement—through comments and shares—tends to *comparatively* have higher ratios in spaces with fewer visibility constraints, where users may feel more comfortable expressing opinions or endorsing content. This pattern is consistent across all bias (Figure 7) and reliability classes (Figure 8), reinforcing the notion that private interactions are an essential component of online engagement dynamics. This again highlights the difference between the public and combined interactions and the importance of considering both the private and public spheres in future research. This observation leads us to the following key insight.

Takeaway: Irrespective of the news class, users tend to engage more deeply in private discussions compared to public ones.

Now, and perhaps even more importantly, to better understand to what extent the relative level of interaction seen for different bias classes depends on which dataset was used we apply the same statistical testing methodology as discussed in previous sections for the two sets independently. First, the Kruskal-Wallis test provides evidence that the distributions differ significantly among the bias classes, regardless of using the *DIRC* and *DIRP*, with p-values of $7.14 \cdot 10^{-23}$ and $2.08 \cdot 10^{-9}$, respectively. Comparing the p-values, we note that the differences appear more significant using the combined set (*DIRC*) than using only the public data (*DIRP*). Second, consistent with visual comparison in the figure, the post-hoc Dunn test only supports statistically significant differences between the *Right* group (which has the biggest fraction of deep interactions regardless of using the combined or public sets) and the other two groups but not among the *Center* and *Left* group. Finally, the Mann-Whitney U test and bootstrapping support the same pattern for the median and means, with the only significant differences again being between the *Right* class and the other two bias classes. For the bias classes, the main results regarding the relative depth of the interactions among the classes are, therefore, consistent regardless of the dataset.

To illustrate the tendency of the posts with *Right*-biased articles seeing a higher deep interaction share, regardless of whether publicly shared or not, we include the CDFs of *DIRC* (combined) and *DIRP* (public) for all three bias classes in Figure 7. We note that in both cases, the *Right* class exhibits a noticeable shift to the right, supporting the following insight.

Takeaway: Content from the *Right* bias class in both public and combined contexts gets deeper user interactions.

Comparison of reliability classes: We now shift our focus to the reliability classes and their deep interaction ratios. Figure 8 summarizes these results. While we again observe deeper interactions for the combined set (suggesting shallower interactions with public posts), the relative differences among the reliability classes themselves are less apparent compared to those observed among the bias classes. We also observe some smaller differences in which relative differences are significant, highlighting the value of also considering the combined data (not only the public data).

First, starting with the combined dataset (*DIRC*), we observe the presence of significant differences among the distributions (Kruskal-Wallis test), with the post-hoc Dunn test (distribution) and Mann-Whitney U test (median) revealing that only the *Most-unreliable* class differs statistically from the others. This observation is further visually supported in Figure 8, which clearly shows a shift to the right in the *DIRC* distribution of the *Most-unreliable* class.

Switching our focus to the *DIRP* distributions, containing only public interactions, we interestingly observe a slightly different pattern, as we here also observe statistical differences between some additional classes. First, as suggested by the *DIRP* CDFs in Figure 8, both extreme classes (*Most-unreliable* and *Most-reliable*) are statistically different from the other two classes (Kruskal-Wallis test followed by Mann-Whitney test, followed by the post-hoc Dunn test). Second, no statistical significance was observed between the *Most-reliable* and *Most-unreliable* classes themselves or between

the *Reliable* and *Unreliable* classes. We base the following insight based on our observations across the two datasets.

Takeaway: Users exhibit higher deep interaction levels with content from the *Most-unreliable* class.

6 Topic-Based Analysis

In this section, we extend our investigation by conducting a topic-based analysis of news articles to assess whether engagement patterns vary across different content themes. For this analysis, we first extracted the textual content from the dataset news articles, a process that succeeded for 14,996 items. We then employed the ChatGPT 4o-mini model to classify each article into one of 14 predefined topics: *Politics*, *Health & Medicine*, *Crime*, *Business & Finance*, *Environment & Climate*, *Entertainment*, *Education*, *Sports*, *Science*, *Lifestyle & Leisure*, *Religion*, *Technology*, *Arts & Culture*, and *Food*. This specific set of topics was selected to capture the diverse range of subject matters prevalent in contemporary news consumption, while also reflecting societal interests and trends.

Notably, the ordering of topics here mirrors the distribution observed in our dataset—ranging from the most frequent (*Politics*, with 7,956 articles) to the least frequent (*Food*, with 57 articles). However, it is important to note that these frequencies are influenced by two key factors: (1) the initial article selection criteria of Ad Fontes Media, and (2) the timeframe of our dataset, which spans a period heavily impacted by the COVID-19 pandemic, leading to an elevated presence of *Health & Medicine* content (the second most prevalent topic). Finally note that for articles where none of these predefined topics were clearly identifiable, the model was instructed to leave them uncategorized. In total, 14,650 articles (97.7%) were successfully classified into one of the 14 categories.

First, to assess whether our earlier findings hold within specific topics, we conducted a stratified analysis across topic groups. Our results indicate that the previously observed patterns remain, with statistical significance for several of the classes with enough samples. For example, in the case of the *Politics* topic the insights reported in both Sects. 4 and 5 remain statistically robust when considering the reliability dimension. As another example, for *Health & Medicine* (the 2nd largest topic), all the conclusions in Sects. 4 and 5 for both the bias and reliability dimensions are still statistically significant. These confirm the general validity of our broader findings within many of these topics.

Second, to explore how engagement varies across topics, we analyzed the distributions of the public interaction ratio for each topic. This analysis was conducted independent of the articles' bias and reliability, as some topics had too few samples to allow for a breakdown by those categories. Figure 9 shows the distribution of the public interaction ratios (as a boxplot) for each of the topics, sorted based on their medians. From this figure, we can first see that *Entertainment*-related articles exhibit the highest public interaction ratio, indicating that content from this category is more likely to be engaged with in public discussions. This suggests that entertainment news, which often includes celebrity updates and viral stories, is inherently more shareable and widely disseminated in public spaces. On the other hand, *Technology*-related articles




Figure 9: Boxplots of the public interaction ratio for different topics (sorted by median).




Figure 10: Heatmap of Kruskal-Wallis test results for topics' public interaction ratios.

exhibit the lowest public interaction share, meaning they receive a disproportionately high share of engagement in private spaces. This may be attributed to the technical nature of these articles, which are often discussed within niche communities or shared in professional or interest-based private groups.

It is important to note that not all observed differences in public interaction ratios between topics in Figure 9 are statistically significant. To validate the statistical significance of these observed differences, we conducted pairwise Kruskal-Wallis tests between topic groups. Figure 10 presents a heatmap illustrating the statistical significance of differences in public interaction ratios across topics. Orange-colored cells here indicate significant differences ($p < 0.1$), while gray cells indicate non-significant comparisons.

This analysis confirms that *Entertainment*-related content has a significantly higher public interaction share compared to all other topics, whereas *Technology*-related content exhibits significantly lower public engagement compared to most categories. Conversely, topics such as *Food*, we could not find statistically significant differences in public interaction ratios compared to other topics (except for *Entertainment*). It is worth noting that we selected a p-value

threshold of 0.1 for our analysis, as stricter thresholds (e.g., 0.01) did not yield many statistically significant differences between topics. This more lenient threshold was chosen primarily to compensate for the limited sample size for certain topics in our dataset.

7 Top Publishers Analysis

To further contextualize our findings, we analyze the public interaction ratio (*PIR*) and the deep interaction ratio (*DIRC*) for the top publishers in our dataset. In this context, “top publishers” refers to those with the highest number of articles in our (and Ad Fontes Media) dataset. For this analysis, we excluded Yahoo, as it primarily functions as a news aggregator and does not produce original content.


We begin by examining the public interaction ratio across the top publishers. Figure 11a presents a boxplot of the public interaction ratio for each outlet, sorted by median values. For easier interpretation, the color of each “box” indicates the bias class (i.e., *Left*, *Center*, or *Right*) associated with each publisher (based on the labeling of the articles they have published). It is worth noting that some outlets, such as *Politico* and *NBC News*, may shift between the *Left* and *Center* classifications over time due to editorial changes or evolving media landscapes. From this figure, we observe that *New York Post* exhibits the lowest public interaction ratio, suggesting that a larger proportion of its engagement occurs in private spaces. This could be attributed to the nature of its audience, which may prefer engaging with its content in more private settings, such as closed groups. Additionally, the sensationalist reporting of the *New York Post* might encourage discussions that users feel more comfortable having in less-visible online spaces.

Next, we investigate which of these publishers generate the deepest user engagement, as measured by the deep interaction ratio (*DIRC*). Figure 11b illustrates the *DIRC* distributions for the top publishers, again sorted by median values. Notably, *New York Post* and *Fox News*, both from the *Right* class, exhibit the highest median *DIRC* values, followed by *CNN* and *New York Times* from the *Left* class. At the lower end of the spectrum, we find *NPR* and *Reuters*, both from the *Center* class. In general, we observe that publishers with the highest *DIRC* values tend to be biased ones, while those from the *Center* class have the lowest values. The probability of this pattern occurring if we randomly sort the 10 publishers is approximately 2.22%. This observation may suggest that users are more likely to engage deeply with content from publishers that have a clear political leaning, compared to those that maintain a more neutral stance.


8 Related Works

This research most closely relates to studies exploring user engagement dynamics with various social media content. Our main contribution lies in the novel analysis of interaction disparities between public and private news article sharing on Facebook, particularly concerning articles with varying bias and reliability. Examining interactions depth adds another layer to our analysis and findings.

There exist studies that partially address some dimensions highlighted in this paper. However, as discussed in the introduction, the vast majority of these works, particularly pertaining to Facebook,



(a) Boxplots of the public interaction ratios (sorted by median).



(b) Boxplots of the deep interaction ratios (sorted by median).

Figure 11: Comparison of (a) public interaction ratio and (b) deep interaction ratio for the top-10 outlets.

predominantly concentrate on the platform’s public sphere. To contextualize our study within the context of the existing literature and to highlight the novelty of our contribution, we next describe the most closely related works.

One line of this prior research has studied the differences in the engagement levels of different bias or reliability (or both) classes on different platforms including Facebook [4, 10, 17], Reddit [5, 44], and Twitter [29, 43]. For example, Edelson et al. [10], while focusing on the public sphere, conducted a large-scale study on user engagement with 7.5 million posts from 2,551 publishers (Facebook pages) across both bias and reliability dimensions. Their results show that individual posts from non-misinformation news outlets tend to attract lower median engagement than misinformation, aligning with our findings. Weld et al. [44] carried out a similar investigation on Reddit rather than on Facebook. When considering the reliability parameter, Vosoughi et al. [43] analyzed a dataset of approximately 126,000 stories tweeted by over 3 million people more than 4.5 million times, finding that false news stories reached more people than true news and diffused significantly farther, faster, deeper, and more broadly than the truth across various categories of information. Notably, during the critical period of the COVID-19 pandemic and by doing a cross-national study, Altay et al. [4] delineated the surge in online news consumption, with credible news outlets witnessing a significant boost. Samory et al. [37] characterized the social media news sphere through user co-sharing practices by focusing on 639 news sources, both credible and questionable, and divided them into 4 clusters, and characterizing them according to the audience that shares their articles on Twitter and how the stylometric features used by each cluster is successful in getting users engagement. Lamot et al. [23] examined how news headlines are remediated on Facebook and how this affects user engagement. Finally, Boukes et al. [6] compared user-content interactivity and audience diversity across news and satire, finding differences in online engagement between satire, regular news, and partisan news. Considering the bias dimension, on the other hand, Wischnewski et al. [45] discovered that users are more inclined to share hyperpartisan news articles that coincide with their own political views. The bias towards the right party in sharing and engagement with news is another thesis that has been studied by previous works on both Twitter [14] and Facebook [13].

Research has also examined other dimensions of user engagement, with, for example, Aldous et al. [2] showing that content

topics influence engagement. Studies have found that emotional content generates higher Facebook engagement, with Maier et al. [25] demonstrating this effect while controlling for author influence. Finally, some works have considered the time factor and analyzed temporal patterns of engagement with different news classes [41].

In summary, while previous research has partly explored some dimensions of user engagement with news, our study is the first to directly compare the dynamics of public and private sharing of news articles on Facebook.

9 Conclusions

In conclusion, this study makes several important contributions to the understanding of news engagement on social media. By developing a robust methodology through collaboration with the CrowdTangle team, we present the first comprehensive comparison of engagement dynamics between Facebook’s publicly tracked content and overall platform interactions. Our analysis of over 19K news articles reveals that public engagement patterns often fail to reflect platform-wide behavior, with significant implications for research methodology and content moderation strategies.

Our findings challenge several assumptions about news engagement on Facebook. First, users consistently engage more deeply with content in less public spaces, regardless of the content’s bias or reliability classification. Second, particular attention should be paid to right-biased and least reliable news content, which generate notably deeper engagement across both public and less public spheres. Third, we highlight some important limitations of focusing solely on Facebook’s public sphere. For example, the extremities of the reliability spectrum (*Most-reliable* and *Most-unreliable*) exhibited a tendency to elicit deeper interactions from users within the public sphere. However, this narrative shifted when examining combined interactions, indicating that only the *Most-unreliable* class triggered statistically significant deeper engagement from Facebook users. This finding suggests that research conclusions drawn solely from public data may not accurately represent overall platform engagement patterns, as not only the magnitude of engagement differs between public and private but also the relative ordering of how different content types engage users.

These insights carry significant implications for multiple stakeholders. For researchers, our findings underscore the importance

of considering engagement disparities between public and less public spheres when studying social media behavior, particularly for news content at different ends of the bias and reliability spectrums. For content moderators, our results suggest the need to reevaluate strategies based solely on publicly visible interactions. For policy-makers, our findings highlight the importance of considering both public and less public spheres when developing policies for social media governance.

Future research should investigate the underlying factors driving these engagement patterns, particularly through user interviews, behavioral studies, and analysis of post comments. While challenges remain in understanding social media dynamics, our findings illuminate significant differences in how users interact with news across Facebook's diverse spheres, highlighting the complexity of news-sharing behavior on social media platforms.

Acknowledgments

The authors express their gratitude to CrowdTangle for providing the Facebook data. They also extend their thanks to the anonymous reviewers for their insightful comments that helped improve the paper. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

- [1] Ad Fontes Media. 2024. Ad Fontes Media Bias Chart. <https://adfontesmedia.com/interactive-media-bias-chart/>.
- [2] Kholoud Khalil Aldous, Jisun An, and Bernard J. Jansen. 2019. View, Like, Comment, Post: Analyzing User Engagement by Topic at 4 Levels across 5 Social Media Platforms for 53 News Organizations. In *Proc. ICWSM*, Vol. 13. 47–57.
- [3] AllSides. 2023. AllSides. <https://www.allsides.com/>.
- [4] Sacha Altay, Rasmus Kleis Nielsen, and Richard Fletcher. 2022. Quantifying the “infodemic”: People turned to trustworthy news outlets during the 2020 coronavirus pandemic. *J. Quant. Descr. Digit. Media* 2 (Aug. 2022).
- [5] Robert M Bond and R Kelly Garrett. 2023. Engagement with fact-checked posts on Reddit. *PNAS Nexus* 2, 3 (01 2023).
- [6] Mark Boukes, Xiaotong Chu, MF Abdulqadir Noon, Rufe Liu, Theo Araujo, and Anne C Kroon. 2022. Comparing user-content interactivity and audience diversity across news and satire: Differences in online engagement between satire, regular news and partisan news. *J. Inf. Technol. Politics* 19, 1 (2022), 98–117.
- [7] Giovanni Brena, Marco Brambilla, Stefano Ceri, Marco Di Giovanni, Francesco Pierri, and Giorgia Ramponi. 2019. News Sharing User Behaviour on Twitter: A Comprehensive Data Collection of News Articles and Social Interactions. In *Proc. ICWSM*, Vol. 13. 592–597.
- [8] Jacob Liedke Christopher Aubin. 2024. Social Media and News Fact Sheet. *Pew Research Center* (2024).
- [9] Ratan Dey, Zubin Jelveh, and Keith Ross. 2012. Facebook users have become much more private: A large-scale study. In *Proc. IEEE PerCom*. 346–352.
- [10] Laura Edelson, Minh-Kha Nguyen, Ian Goldstein, Oana Goga, Damon McCoy, and Tobias Lauinger. 2021. Understanding Engagement with U.S. (Mis)Information News Sources on Facebook. In *Proc. IMC*. 444–463.
- [11] Matt Garmur. 2022. Chrome Extension: Why do the items add up to more than the totals? <https://archive.is/mrNPU>
- [12] Christine Geeng, Savanna Yee, and Franziska Roesner. 2020. Fake news on Facebook and Twitter: Investigating how people (don't) investigate. In *Proc. CHI*.
- [13] Sandra González-Bailón, David Lazer, Pablo Barberá, Meiqing Zhang, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Deen Freelon, Matthew Gentzkow, Andrew M Guess, et al. 2023. Asymmetric ideological segregation in exposure to political news on Facebook. *Science* 381, 6656 (2023), 392–398.
- [14] Sandra González-Bailón, Valeria d'Andrea, Deen Freelon, and Manlio De Domenico. 2022. The advantage of the right in social media news sharing. *PNAS Nexus* 1, 3 (07 2022), pgac137.
- [15] Andrew Guess, Brendan Nyhan, and Jason Reifler. 2018. Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances* 4, 1 (2018).
- [16] Ehsan-Ul Haq, Yang K Lu, and Pan Hui. 2022. It's All Relative! A Method to Counter Human Bias in Crowdsourced Stance Detection of News Articles. *Proc. ACM on Human-Computer Interaction* 6, CSCW (2022), 1–25.
- [17] Dan Hiaeshutter-Rice and Brian Weeks. 2021. Understanding Audience Engagement with Mainstream and Alternative News Posts on Facebook. *Digital Journalism* 9, 5 (2021), 519–548.
- [18] Benjamin D Horne, Jeppe Nørregaard, and Sibel Adali. 2019. Different spirals of sameness: A study of content sharing in mainstream and alternative media. In *Proc. ICWSM*, Vol. 13. 257–266.
- [19] Ferenc Huszár, Sofia Ira Ktena, Conor O'Brien, Luca Belli, Andrew Schlaikjer, and Moritz Hardt. 2022. Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences* 119, 1 (2022).
- [20] Jason J Jones, Robert M Bond, Eytan Bakshy, Dean Eckles, and James H Fowler. 2017. Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 US presidential election. *PLoS one* 12, 4 (2017).
- [21] Cheonsoo Kim and Sung-Un Yang. 2017. Like, comment, and share on Facebook: How each behavior differs from the other. *Public Relations Review* 43, 2 (2017).
- [22] Brian Knutson, Tiffany W Hsu, Michael Ko, and Jeanne L Tsai. 2024. News source bias and sentiment on social media. *PLoS one* 19, 10 (2024).
- [23] Kenza Lamot, Tim Kreutz, and Michaël Oppenhaffen. 2022. “We Rewrote This Title”: How News Headlines Are Remediated on Facebook and How This Affects Engagement. *Social Media + Society* 8, 3 (2022).
- [24] Danielle Lottridge and Frank R. Bentley. 2018. Let's Hate Together: How People Share News in Messaging, Social, and Public Networks. In *Proc. CHI*. 1–13.
- [25] Scott R. Maier. 2015. Compassion Fatigue and the Elusive Quest for Journalistic Impact: A Content and Reader-Metrics Analysis Assessing Audience Response. *Journalism & Mass Communication Quarterly* 92, 3 (2015), 700–722.
- [26] Alice E Marwick. 2018. Why do people share fake news? A sociotechnical model of media effects. *Georgetown law technology review* 2, 2 (2018), 474–512.
- [27] Media Bias Fact Check. 2024. Media Bias Fact Check. <https://mediabiasfactcheck.com/>.
- [28] Meta. 2023. URL - Graph API Reference. <https://developers.facebook.com/docs/graph-api/reference/v21.0/url>. Accessed: 2024-09-11.
- [29] Alireza Mohammadinodooshan and Niklas Carlsson. 2024. Understanding Engagement Dynamics with (Un) Reliable News Publishers on Twitter. In *Proc. ASONAM*. 36–47.
- [30] Seungahn Nah and Masahiro Yamamoto. 2018. Communication and Citizenship Revisited: Theorizing Communication and Citizen Journalism Practice as Civic Participation. *Communication Theory* 29, 1 (07 2018), 24–45.
- [31] Nic Newman, Richard Fletcher, Kirsten Eddy, Craig T. Robertson, and Rasmus Kleis Nielsen. 2023. Digital News Report 2023.
- [32] NewsGuard. 2024. NewsGuard. <https://www.newsguardtech.com/>.
- [33] Vanessa Otero. 2021. Ad Fontes Media's Multi-Analyst Content Analysis White Paper. <https://adfontesmedia.com/white-paper-2021>
- [34] Gordon Pennycook, Tyrone D Cannon, and David G Rand. 2018. Prior exposure increases perceived accuracy of fake news. *J. Exp. Psychol. Gen.* 147, 12 (2018).
- [35] Jürgen Pfeffer, Daniel Matter, and Anahit Sargsyan. 2023. The Half-Life of a Tweet. *Proc. ICWSM (2023)*, 1163–1167.
- [36] Srihaasa Pidikiti, Jason Shuo Zhang, Richard Han, Tamara Lehman, Qin Lv, and Shivakant Mishra. 2020. Understanding how readers determine the legitimacy of online news articles in the era of fake news. In *Proc. ASONAM*. 768–775.
- [37] Mattia Samory, Vartan Kesiz Abnousi, and Tanushree Mitra. 2020. Characterizing the social media news sphere through user co-sharing practices. In *Proc. ICWSM*, Vol. 14. 602–613.
- [38] Tess. 2021. What data is CrowdTangle tracking? <https://archive.is/ItEuA>.
- [39] Nik Thompson, Xuequn Wang, and Pratiq Daya. 2020. Determinants of News Sharing Behavior on Social Media. *J. Comput. Inf. Syst.* 60, 6 (2020), 593–601.
- [40] Elin Thorgren, Alireza Mohammadinodooshan, and Niklas Carlsson. 2024. Temporal Dynamics of User Engagement on Instagram: A Comparative Analysis of Album, Photo, and Video Interactions. In *Proc. WebSci*. 224–234.
- [41] Luca Vassio, Michele Garetto, Emilio Leonardi, and Carla Fabiana Chiasserini. 2022. Mining and modelling temporal dynamics of followers' engagement on online social networks. *Social Network Analysis and Mining* 12, 1 (2022), 96.
- [42] Jessica Vitak, Paul Zube, Andrew Smock, Caleb T Carr, Nicole Ellison, and Cliff Lampe. 2011. It's complicated: Facebook users' political participation in the 2008 election. *CyberPsychology, behavior, and social networking* 14, 3 (2011), 107–114.
- [43] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.
- [44] Galen Weld, Maria Glenski, and Tim Althoff. 2021. Political Bias and Factualness in News Sharing across more than 100,000 Online Communities. In *Proc. ICWSM*, Vol. 15. 796–807.
- [45] Magdalena Wischniewski, Axel Bruns, and Tobias Keller. 2021. Shareworthiness and Motivated Reasoning in Hyper-Partisan News Sharing Behavior on Twitter. *Digital Journalism* 9, 5 (2021), 549–570.