Understanding Engagement Dynamics with (Un)Reliable News Publishers on Twitter

Alireza Mohammadinodooshan and Niklas Carlsson

Linköping University, Sweden
alireza.mohammadinodooshan@liu.se, niklas.carlsson@liu.se

Abstract. According to the Pew Research Center, a majority of X (formerly Twitter) users in the U.S. (55%) regularly consume news through the platform, exceeding the ratio for all other major social media platforms. Still, the current literature falls short in providing insights into the relative interaction patterns seen for different classes of news on this platform. To address this gap, this study provides a large-scale analysis of user interactions with different news classes, emphasizing both the bias and reliability of the publishers. To this end, we have compiled a robust dataset comprising more than 75 million tweets posted over 56 months by 2,041 labeled U.S. news publishers. Using this dataset, we study the engagement patterns across news categories, identifying several statistically significant variances. Understanding these dynamics is crucial for developing informed strategies for news dissemination, audience targeting, and content moderation. Accordingly, this study offers data-driven insights to support such strategy development.

1 Introduction

Social media have transformed people from passive consumers to active participants in news dissemination [12]. In fact, through their interactions (i.e., likes, shares, and comments), users play a key role in the successful dissemination of news, including unreliable content, making it crucial to understand these interaction patterns. While prior research has extensively studied user interaction patterns with different news publishers on Facebook [4] and Reddit [20], the dynamics of these interactions on X remain relatively underexplored; despite a higher portion of X's U.S. user base (55%) regularly accessing news via its service than seen on competing services (e.g., Facebook (43%), Reddit (38%) [19]).

To address this gap, we present a large-scale analysis of user interactions with a wide range of U.S. news publishers (outlets) on X that addresses the following research question (**RQ**): *How do user engagement patterns on X differ across news publishers from varying bias and reliability classes?*

Specifically, we gathered and analyzed a dataset containing 75 million tweets from 2,041 publishers, spanning 56 months (June 2018 to Feb. 2023), along with the 5.2 billion interactions associated with these tweets. Like some other recent studies (e.g., [4]), we categorize U.S. publishers based on both their bias ("Left",

This is the author's version. It is posted here for your personal use, not for redistribution. The definitive version is published in Proceeding of International Conference on Advances in Social Networks Analysis and Mining (ASONAM) 2024: https://doi.org/10.1007/978-3-031-78548-1_4.

"Left-Center", "Least Biased", "Right-Center", and "Right") and reliability ("Unreliable" or "Reliable"). Using the large labeled dataset, we then provide a comprehensive study into how users engage with more or less reliable news across the political spectrum. Here, we explore key factors that shape user engagement, including the interplay between a publisher's bias, reliability, followership, content views, and different interaction types. We next outline our main contributions.

First, the paper provides insights into differences in the level of user engagement across content types, both at an aggregate level (Sec. 4.2) and when statistically analyzed (Sec. 4.3). This analysis clarifies which news publishers attract more interactions, the patterns in content volume, and the influence of followership size on publishers' success dynamics. As a concrete example, our empirical insights reveal that across all biased publisher categories, *Unreliable* publishers consistently exhibit higher interactions compared to their *Reliable* counterparts. However, a deeper examination of the follower base nuances refines these findings, emphasizing that only *Unreliable* publishers from the most extreme bias categories outperform *Reliable* counterparts in terms of engagement.

Second, by making comparisons based on what type of interactions (i.e., likes, retweets, replies, and quotes) users most engage in with different content, we provide insights into subtle differences in how users interact with different outlets (Sec. 6). A noteworthy observation is that right-leaning publishers are more effective in prompting dialogue-based interactions (i.e., replies) compared to other publishers. Moreover, users tend to engage in shallow interactions like "liking" with *Unreliable* publishers, whereas deeper forms of engagement, such as quoting and replying, are more likely to occur with *Reliable* news sources.

Third, by leveraging that X provided research access to view counts (impressions) from Dec. 2022, the paper presents a pioneering analysis of impressions and engagement rates (i.e., the portion of impressions translating into user interactions) across news classes (Sec. 7). Our analysis here shows that *Unreliable* publishers, regardless of bias, excel in capturing user attention per view. While the first two dimensions of our study complement prior research (done on other platforms than X), this third dimension introduces the first analysis of its kind.

While our target audience primarily is other researchers, we note that our findings also may help publishers refine their strategies, assist moderators in safeguarding information integrity, and provide policymakers with valuable input to address the challenges posed by *Unreliable* publishers.

Outline: After discussing the related works (Sec. 2), we present our methodology (Sec. 3), results and discussions (Secs. 4-7), and our conclusions (Sec. 8).

2 Related Works

This work aligns with works focused on user engagement with news content of varying political biases and levels of reliability. As the most similar work to our research in this group, we can mention [4]. Considering both bias and reliability dimensions and collecting an 8-month panel of 7.5M posts, they try to analyze the users' engagement with different classes of news on Facebook. Similar to what we have recorded for X, their results show that posts from *Unreliable* news outlets receive consistently higher median engagement than *Reliable*. Similar to [4], a study of 416 news organizations on Facebook from 2012 to 2017 found that the most engaging sources often are the most ideologically extreme [8], a trend that part of our results confirmed on X. Weld et al. [20] also extended this line of inquiry to Reddit by examining the effects of bias and reliability there.

In terms of reliability, the seminal work by Vosoughi et al. [18] also stands out, demonstrating the faster and broader spread of unreliable news compared to the truth. However, during the COVID-19 pandemic, Altay et al. [3] observed an upsurge in news consumption, with credible outlets experiencing increased traffic. On the other hand, and investigating the impact of bias, users' bias towards sharing right-leaning news on X has been reported by [6].

When it comes to X, separate from the influence of political bias and reliability, some researchers have explored user engagement with news publishers more broadly. One line of research here is Aldous et al.'s research [2, 1], whose studies have a limited scope in terms of news publishers studied (8 and 53 news publishers, respectively). Another related work here is [5], where the authors study the user interactions with 232 news sources through 432K tweets, explicitly mentioning or retweeting those sources.

Our research differentiates itself from prior studies by providing the largest scale study (2,041 news sources, more than 75M tweets, and 5.2B interactions) on X, conducting a combined analysis of political bias and reliability, and providing a pioneering analysis of impressions and engagement rates for different news classes. Finally, we should mention that while our focus is on the U.S. news ecosystem, studies like [17] have considered other news ecosystems.

3 Methodology

Dataset Compilation: The dataset used for our analysis was created using a multi-step approach involving six primary steps. Our study commenced with compiling a comprehensive list of U.S. news publishers sourced from Media Bias Fact Check (MBFC)[10], an independent organization that evaluates the bias and reliability of media sources.

Second, using MBFC's labels and bias-meter icons assigned to each publisher, we identified the bias class for each publisher. For publishers missing one of the five labels of interest ("Left", "Left-Center", "Least Biased", "Right-Center", and "Right"), we utilized Robertson et al.'s bias scores [14] to train five kernel density estimators and determine the bias class of those publishers.

Third, we focused on classifying publishers as either *Unreliable* or *Reliable*, using the iffy index (https://iffy.news), which relies on MBFC labels and is widely employed in previous works [7]. Fourth, we identified and linked the publishers to their corresponding X accounts, mainly by visiting their websites.

Fifth, we collected all tweets posted by these accounts, along with associated user interactions (likes, retweets, quotes, and replies) from June 2018 to the end of Jan. 2023. We refer to this collection of data as the "Interactions dataset."

Table 1: Summary statistics categorized by bias and reliability type. Here (B = Billion, M = Million, K = Kilo) and (Rel. = Reliable, Unrel. = Unreliable)

Bias	Left		Left-Center		Least Biased		Right-Center		Right		Total
Reliability	Rel.	Unrel.	Rel.	Unrel.	Rel.	Unrel.	Rel.	Unrel.	Rel.	Unrel.	
Outlets	227	15	519	10	773	9	235	9	118	126	2,041
Followers	177.0M	19.8M	437.1M	13.8M	117.4M	1.1M	75.1M	13.7M	16.7M	73.4M	945.0M
Tweets	5.9M	637.9K	19.8M	134.7K	33.3M	118.4K	9.9M	367.7K	2.6M	2.3M	75.0M
Interactions	1.2B	287.2M	1.6B	101.2M	475.8M	6.1M	233.1M	92.4M	314.6M	938.2M	5.2B

For tweets published after Dec. 15, 2022, we were also able to gather impressions data, which we call the “Impressions dataset." While API access limitations prevented complete data collection of 97 accounts, we were able to complete the collection for 2,041 accounts. By prioritizing statistical analysis over aggregated analysis, we mitigate the impact of this limitation on our overarching conclusions.

Finally, our analysis required the follower count of publishers' Twitter accounts, including how these numbers had changed over time. We here utilized the Internet Archive's Wayback Machine to manually obtain historical snapshots and reconstruct a timeline of the follower growth for each account.

Statistical Tests: While the ANOVA family of tests at first might seem fitting for our purposes, they face challenges due to the heavy-tailed distributions under study (and non-normal errors also after log-transform) and correlations between bias and reliability factors. We thus used the non-parametric Kruskal-Wallis test to compare the medians and overall distributions across groups without assuming normality. In our analysis, p-values below 0.01 are considered significant. For similar reasons, Spearman's rank is used for correlation analysis.

4 High-level Category Comparisons

4.1 Dataset Overview

Table 1 presents an overview of our dataset, categorized by bias and reliability dimensions. For each bias and reliability category (columns), the table shows (split over four rows): (1) the number of outlets in that category, (2) the combined number of followers as of Feb. 2023, (3) the aggregate number of tweets in the dataset, and (4) the combined sum of all interaction types in the dataset, including likes, retweets, quotes, and replies. In total (last column in the table), our dataset comprises 5.2B interactions from 75M tweets across 2,041 news outlets, engaging a combined following of 945M users. This extensive engagement underscores the significant reach and influence of news outlets on X.

4.2 Aggregate Category Comparison

Focusing first on the *Reliable* publishers, a distinct pattern unfolds: outlets affiliated with the left party (aggregating *Left* and *Left-Center*) have the largest

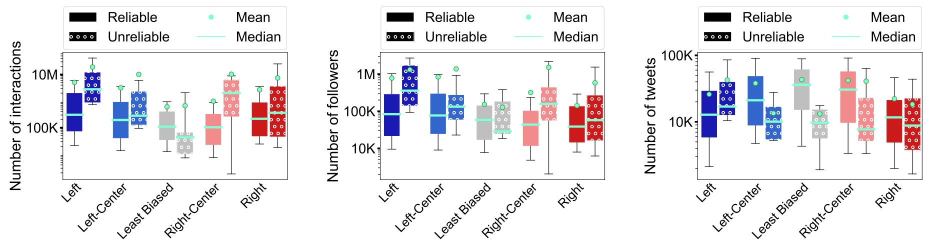

(a) Number of interactions (b) Number of followers (c) Number of tweets

Fig. 1: Distributions comparisons across different outlets (interactions dataset).

follower base at 614.1M (177M+437.1M), with a total of 2.8B (1.2B+1.6BM) interactions. Interestingly, the numerous *Least Biased* outlets do not exhibit the same level of engagement, indicating a user preference for engaging with outlets expressing a clear bias rather than those perceived as neutral. This inference is reinforced when examining the average interactions per outlet: 3.8M/outlet for left-party (*Left* and *Left-Center*), 0.6M/outlet for *Least Biased*, and 1.5M/outlet for right-party (*Right* and *Right-Center*), confirming that the same trend holds for the right-party-affiliated outlets. In addition to seeing the lowest average interaction rate per outlet, the *Least Biased* group is noteworthy for its substantial disparity in the number of *Unreliable* versus *Reliable* publishers (9 vs. 773).

Furthermore, looking at the interactions per outlet for the *Unreliable* outlets reveals substantial differences between biased classes and *Least Biased*: 15.5M/outlet for the 25 left-party outlets, 0.67M/outlet for *Least Biased*, and 7.6M/outlet for the 135 right-party outlets. These observations imply that neutral publishers may not yield the same influence as biased ones.

Considering the *Unreliable* publishers, the *Right* class outperforms other classes in both outlet count and total engagement, with 126 outlets amassing 73.4M followers and 938.2M interactions. Interestingly, the *Left-Unreliable* publishers though fewer, still attract a significant 287.2M user interactions and more interactions per outlet, indicating stronger individual engagement despite fewer outlets (for both *Reliable* and *Unreliable* spheres). These observations warrant further statistical analysis, which attributes equal weight to each outlet. We next provide such per-publisher-based analysis.

4.3 Per-publisher-based Comparison

To account for variations within each category and more fairly weight the outlets (publishers) within each category, we next compare their relative distribution statistics. Fig. 1 shows box plots for the number of (a) interactions, (b) followers, and (c) tweets per outlet across the ten outlet categories. Here, the interquartile range (IQR), shown as a box, spans from the 25^{th} to the 75^{th} percentile, and the whiskers specify the 10^{th} to the 90^{th} percentile. Finally, the median is shown using a horizontal line and the mean using a marker. We note that the y-axis is plotted on a log scale, and the averages are substantially larger than the median values, indicating that the distributions are highly skewed and heavy-tailed.

Higher engagement with *Unreliable*: Comparing the total number of interactions (Fig. 1a), with the exception for the *Least Biased* class, the group of *Unreliable* outlets of each bias sees higher distribution statistics than the corresponding *Reliable* outlets. This suggests that neutral news may be consumed and interacted differently with than biased news.

Biased news more engaging: When considering the bias, biased publishers consistently achieve higher interaction levels compared to the *Least Biased* class. This difference is statistically significant for both the right and left parties, irrespective of reliability labels. Within the *Unreliable* class, it is unexpected that *Right-Center* party outperforms *Right*. Several factors contribute here, including the larger number of publishers in the *Right* class, intensifying competition, while the total user engagement potential in this group is limited. Another explanation for the higher engagement levels of *Left Unreliable* and *Right-Center Unreliable* may be their larger follower base (Fig. 1b), which we will discuss next.

Key Takeaway: For all four biased classes, the *Unreliable* outlets receive significantly more interactions than their *Reliable* counterparts.

Not all followers contribute equally to interactions: Comparing Figures 1b and 1a reveals a correlation between the two metrics. While accounts with more followers are expected to have greater interactions due to increased visibility, the *Right-Center Unreliable* group shows exceptionally higher interaction levels that surpass follower count expectations. To understand the uneven levels of interactions per follower, Sec. 5 explores normalized interaction levels.

Left-biased outlets delivering *Reliable* content has most followers: In addition to the above correlations, we notice a decreasing trend in followers from left bias to right bias among *Reliable* publishers. This suggests that while inherent variations exist in follower counts across biases; the reliability dimension may not be the primary factor influencing follower counts. Instead, bias plays a more significant role. It is worth noting that while this trend is statistically significant, mean values influenced by outliers do not follow this pattern.

Outlets spreading *Unreliable* content often achieve high engagement with relatively fewer tweets: Examining the number of tweets in Fig. 1c reveals a divergent pattern. Most *Unreliable* classes, excluding *Left*, exhibit notably lower tweet output. Combined with our previous observation that these classes often see higher interaction levels, this implies that for these groups, the tweet frequency does not have a straightforward correlation with the interaction levels. Perhaps the nature of content, selective attention, and audience capacity play more significant roles here.

Higher tweet levels for the *Least Biased Reliable* outlets: For the *Reliable* classes, the *Least Biased* outlets exhibit a notably higher volume of tweets than their more politically biased counterparts. The trend for the *Reliable* classes, that the more biased publishers generate fewer tweets than the less biased groups, is significant among all five groups (all six possible pairwise comparisons matching this statement have p-value < 0.001). Note that in these comparisons, we compared all four biased groups against the *Least Biased* class, and for each party, the more biased one against the more centered one.

The heightened tweet levels of the *Least Biased* outlets publishing *Reliable* content see less user engagement: This phenomenon may be attributed to *Least Biased* outlets attempting to cover a broader spectrum of users and topics, potentially leading to content fatigue among followers. Additionally, the less polarizing nature of reliable and unbiased content may inherently attract less engagement compared to more controversial or partisan material. However, the number of followers also significantly influences interaction levels, potentially obscuring these effects. Therefore, and to provide a refined understanding, in the next section, we examine interaction levels normalized by follower count.

5 Followers and Interactions

The prior section suggested a correlation between the number of followers of an outlet and its interaction levels. To understand this link, we conduct a correlation analysis to quantify its impact. Here, a strong correlation is observed across all classes, quantified by a Spearman coefficient of 0.84 (p-value < 0.001), emphasizing a significant positive relationship between followers and interactions.

Second, we compare the (normalized) distributions of the interactions per follower seen across the outlets. Here, we make an extra effort to (1) consider that the number of followers changed over the measurement period, and (2) give equal weight to each post made during the measurement period. For this purpose, rather than simply calculating the ratio between the total interactions (Fig. 1a) and the number of followers (Fig. 1b), for each outlet, we calculated a weighted per-post-based metric as $\sum_{i=1}^{N} \frac{I_i}{f_i}$, where I_i is the number of interactions received for post i (out of the N posts observed for that outlet), and f_i is the outlet's (estimated) follower count at the time of that tweet post. Note that the above type of weighting is particularly important for our study periods, during which the follower base of many of these outlets has been found to be highly dynamic (especially during the COVID period [11]).

The box plot with this normalized metric (Fig. 2) reveals that some differences between classes diminish when considering per-follower interactions. Notably, the *Right-Center* class shows minimal distinctions between *Unreliable* and *Reliable* classes. While statistically significant median differences exist among *Reliable* classes, they lack a discernible pattern as in raw interaction (Fig. 1a) and follower counts (Fig. 1b). However, a consistent trend persists in *Unreliable* categories, emphasizing that political bias remains a key factor in interactions, even when normalized for follower count. Notably, extreme bias classes (*Left* and *Right*) within *Unreliable* garner higher per-follower interactions, highlighting challenges in content moderation. Despite having a larger audience, these outlets excel at engaging individuals, reflecting a higher interaction commitment per follower. We provide a deeper analysis of this trend in Sec. 7.

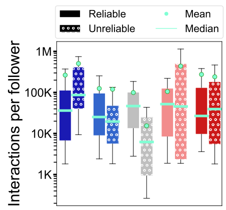

Fig. 2: Interactions per follower

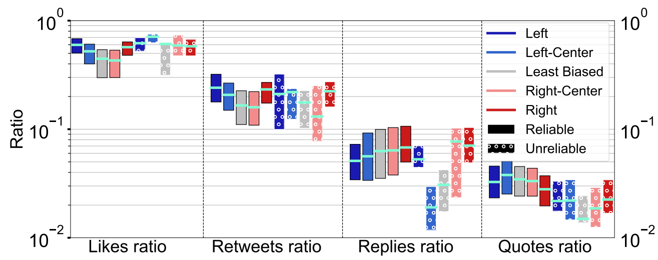

Fig. 3: Distribution of interaction types ratios

Key Takeaway: Within the *Left* and *Right* political extremes, *Unreliable* publishers exhibit statistically significant higher interactions per follower, despite having higher levels of followers.

6 Interaction Types Analysis

Different interaction types (likes, retweets, replies, and quotes) reflect varying levels of user engagement, from simple agreement to active discussion or critique. Motivated by these differences, we present a fine-grained analysis of user interactions across different outlet classes. Fig. 3 shows the distribution ratios of interaction types for each category, calculated as the proportion of each interaction type relative to total interactions per publisher. For similar arguments as before, for each publisher, we calculate the mean of these ratios over all its posts (not taking the ratio of the totals) and exclude tweets without any interactions. Here, a preference for less effort-intensive interactions is observed, where we see a logarithmic decrement in the interaction levels seen across various actions: likes (69% of all interactions), retweets (19%), replies (8%), and quotes (4%). With likes dominating interactions, the likes ratio mirrors many trends observed earlier. These patterns are often reflected in retweet ratios as well.

Replies: The replies ratio offers a fresh perspective different than those seen for likes and retweets. In particular, the right-party classes exhibit a higher replies ratio than the left-party classes, suggesting a greater inclination for users to engage in conversations initiated by right-leaning publishers. For the *Reliable* classes, we observe a significant gradient: beginning with the *Left* class, which presents the lowest replies ratio, there is a marked and statistically significant increase through the *Left-Center*, *Least Biased*, and *Right-Center*, culminating with the *Right* class. This trend suggests that as we move toward the right of the political spectrum, user interactions become increasingly interactive. This higher replies ratio of right-party publishers repeats in the *Unreliable* sphere also.

Quote usage: As we progress from the simpler interactions of likes and retweets to the more involved replies, the advantage of *Unreliable* over *Reliable* counterparts diminishes, although certain groups deviate from this trend. For quotes, this trend intensifies. Here, *Unreliable* classes consistently exhibit lower

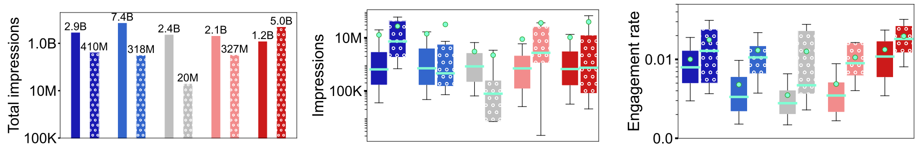

(a) Total impressions (b) Impressions distribution (c) Eng. rate distribution

Fig. 4: Impressions and engagement rates in the impressions dataset.

ratios across all classes compared to *Reliable* groups. This implies that, despite widespread sharing and liking, *Unreliable* is quoted less frequently, suggesting users may exercise greater selectivity or thoughtfulness when engaging with content requiring additional commentary or context.

Key Takeaway: The replies ratio favors the right party, and the quotes ratio favors *Reliable* publishers.

7 Impressions Study

We next study user engagement rates associated with the different classes. Here, we take advantage of X, having made available impression statistics for the last 48 days of our data collection period (Dec. 15, 2022 - Jan. 31, 2023).

Impressions statistics: The impressions dataset includes 1.9M tweets, constituting 2.5% of the main dataset. These tweets have amassed 136.6M interactions, with likes making up the majority at 72% (99.0M likes) and an (impressive) total of 22.0B impressions (distributed as per Fig. 4a). In the *Reliable* classes, *Left-Center* dominates with 7.4B views, while in the *Unreliable* classes, *Right* leads with 5.0B impressions, correlating with its higher account count.

Number of impressions per outlet: To normalize for class size, we examine the distribution of interactions per outlet for each class in Fig. 4b. Aside from a marginal distinction between *Least Biased* and *Right-Center* classes (p-value=0.05), the *Reliable* classes show no statistically significant differences in impression levels. This suggests a level playing field in content exposure among these classes. In contrast, *Unreliable* groups align with trends observed in the interactions study (Fig. 1a). Biased news outlets generally command higher impressions than the *Least Biased* class. This disparity can be attributed, in part, to the marked differences in interaction levels between *Unreliable* classes, reflected in impression counts. X's algorithm, recently published [16], incorporates interaction metrics as a core component, corroborated by a Spearman correlation coefficient of 0.90 between impressions and interactions, emphasizing the algorithm's reliance on interactions to shape content visibility.

Engagement rate: Given the above correlation, it becomes pertinent to examine the interaction levels relative to impressions. Following X's [15] and various other scholarly works convention [9], we refer to this metric as the "engagement rate". While this normalized metric has long been argued to be the most important barometer of user engagement [13], it has previously not been possible to be studied on X. We are, therefore, happy to present a head-to-head compassion of the engagement rates seen across the different classes. Fig. 4c summarizes these results, allowing us to make several notable observations. First, across all political biases, *Unreliable* groups consistently achieve significantly higher engagement rates compared to their *Reliable* counterparts, amplifying trends observed in interaction numbers with more pronounced deviations. Interestingly, *Unreliable* surpasses the engagement rate of *Reliable* (even within the *Least Biased* class), with all comparative values being statistically significant (p-values < 0.01), except for the *Left* class (p-value $= 0.08$). This draw to interacting with *Unreliable* publishers can be attributed to the provocative nature of the content they share, which often matches their followers' preferences, as well as the more isolated and condensed networks they form.

Conversely, within *Reliable* classes, significant engagement rate differences are observed between classes. This highlights a consistent pattern where polarized classes (*Left* or *Right*) exhibit higher engagement rates than their respective center-aligned classes (*Right-Center* and *Left-Center*), which, in turn, outperform the *Least Biased* class. The same pattern is observed for *Unreliable* classes. This suggests that content with a clear political stance, particularly aligning with the more extreme ends of the spectrum, tends to be more engaging than neutral content, possibly influenced by the journalistic techniques and focused topics employed by these groups.

Key Takeaway: Regardless of political bias and compared to *Reliable* publishers, *Unreliable* publishers consistently achieves higher engagement rates. Furthermore, publishers with more (left- or right-) pronounced political biases receive greater engagement rates.

8 Conclusions

This paper presents a large-scale analysis of interactions between X users and news publishers, offering a comprehensive view of engagement levels and rates' dynamics across various publishers, when considering their bias, reliability, and follower base. As an example, studying the engagement rate (0.62% on average, indicating roughly one interaction per 160 content views), a statistically significant variance among publisher classes was observed. Alarmingly, the class of *Unreliable* publishers consistently achieves higher engagement rates than *Reliable* publishers. More than reliability, we also consider the influence of political bias. For instance, for the *Reliable* class, an ascending trend in reply ratios was observed from left to right-leaning publishers, suggesting that political alignment may significantly affect user engagement, especially in dialogue-centric interactions. The analysis extends beyond the roles of reliability and bias to also consider the follower count's impact on the engagement with different publishers.

In summary, this paper contributes to scholarly and practical discussions on social media engagement, capturing engagement dynamics with diverse news publishers. It offers valuable guidance for publishers, strategists, and policymakers dedicated to fostering a well-informed, critically engaged online audience.

References

1. Aldous, K., An, J., Jansen, B.: View, like, comment, post: Analyzing user engagement by topic at 4 levels across 5 social media platforms for 53 news organizations. In: Proc. ICWSM (2019). https://doi.org/10.1609/icwsm.v13i01.3208
2. Aldous, K.K., An, J., Jansen, B.: Measuring 9 emotions of news posts from 8 news organizations across 4 social media platforms for 8 months. Trans. Soc. Comput. **4** (2022). https://doi.org/10.1145/3516491
3. Altay, S., Kleis Nielsen, R., Fletcher, R.: Quantifying the "infodemic": People turned to trustworthy news outlets during the 2020 coronavirus pandemic. J. Quant. Descrip. Digit. Media. **2** (2022). https://doi.org/10.51685/jqd.2022.020
4. Edelson, L., Nguyen, M.K., Goldstein, I., Goga, O., McCoy, D., Lauinger, T.: Understanding engagement with U.S. (mis)information news sources on Facebook. In: Proc. IMC (2021). https://doi.org/10.1145/3487552.3487859
5. Glenski, M., Weninger, T., Volkova, S.: How humans versus bots react to deceptive and trusted news sources. In: Proc. ASONAM 2018 (2018)
6. González-Bailón, S., d'Andrea, V., Freelon, D., De Domenico, M.: The advantage of the right in social media news sharing. PNAS Nexus **1** (2022)
7. Hanley, H.W.A., Kumar, D., Durumeric, Z.: A golden age: Conspiracy theories' relationship with misinformation outlets, news media, and the wider internet. ACM Hum.-Comput. Interact. **7** (2023). https://doi.org/10.1145/3610043
8. Hiaeshutter-Rice, D., Weeks, B.: Understanding audience engagement with mainstream and alternative news posts on Facebook. Digital Journalism **9** (2021)
9. Hwong, Y.L., Oliver, C., Van Kranendonk, M., Sammut, C., Seroussi, Y.: What makes you tick? the psychology of social media engagement in space science communication. Computers in Human Behavior **68** (2017)
10. Media Bias Fact Check: https://mediabiasfactcheck.com/ (2023)
11. Mohammadinodooshan, A., Carlsson, N.: Temporal dynamics of user engagement with U.S. news sources on Facebook. In: Proc. SNAMS (2022)
12. Mukerjee, S., Yang, T., Peng, Y.: Metrics in action: how social media metrics shape news production on Facebook. J. Commun. **73** (2023)
13. Niciporuc, T.: Comparative analysis of the engagement rate on Facebook and Google Plus social networks. In: Proceedings of International Academic Conferences. No. 0902287, International Institute of Social and Economic Sciences (2014)
14. Robertson, R.E., Jiang, S., Joseph, K., Friedland, L., Lazer, D., Wilson, C.: Auditing partisan audience bias within google search. Hum.-Comput. Interact. **2** (2018)
15. Twitter Help Center: (2023), https://help.twitter.com/en/managing-your-account/using-the-post-activity-dashboard
16. Twitter's Algorithm: (2023), https://github.com/twitter/the-algorithm
17. Vicario, M.D., Gaito, S., Quattrociocchi, W., Zignani, M., Zollo, F.: News consumption during the Italian referendum: A cross-platform analysis on Facebook and Twitter. In: Proc. DSAA (2017). https://doi.org/10.1109/DSAA.2017.33
18. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. Science **359** (2018). https://doi.org/10.1126/science.aap9559
19. Walker, M., Matsa, K.E.: News consumption across social media in 2021. Pew Research Center (2021)
20. Weld, G., Glenski, M., Althoff, T.: Political bias and factualness in news sharing across more than 100,000 online communities. In: Proc. ICWSM (2021)