Temporal Dynamics of User Engagement with U.S. News Sources on Facebook

Alireza Mohammadinodooshan and Niklas Carlsson
Department of Computer and Information Science, Linköping University, Sweden
alireza.mohammadinodooshan@liu.se, niklas.carlsson@liu.se

***Abstract*—Recently, researchers have modeled how reliability and political bias of news may affect Facebook users' engagement, as measured using interaction metrics such as the number of shares, likes, etc. However, the temporal dynamics of Facebook users' engagement with news of varying degrees of bias and reliability is less studied. In light of the COVID-19 pandemic, it is also important to quantify how the pandemic changed user engagement with various news. This paper presents the first temporal study of Facebook users' interaction dynamics, accounting for both the bias and reliability of the publishers. We consider a dataset of 992 U.S. publishers, and the study spans the period from Jan. 2018 to July 2022. This allows us to accurately assess the effect of the covid outbreak on the temporal dynamics of Facebook users' interactions with different classes of news. Our study examines these two parameters' effect on Facebook user engagement using both per-publisher and aggregated statistics. Several findings are revealed by our analysis, including that publishers in different bias and reliability classes experienced significantly different levels of engagement dynamics during and following the covid outbreak. For example, we show that the least reliable news exhibited the most considerable growth of followers during the *covid* period and the most reliable news sources exhibited the greatest growth rate of followers during the *post-covid* period. We also show that the interaction rate (number of interactions normalized over the number of followers) with Facebook news posts during the *post-covid* period is smaller than it was even before the outbreak. Furthermore, we demonstrate how the COVID-19 outbreak caused statistically significant structural breaks in the temporal dynamics of engagement with several types of news, and quantify this effect. With social media becoming a popular news source during crises, the observed temporal dynamics provide important insights into how information was consumed over the recent years, benefiting both researchers and public sectors.**

***Index Terms*—Facebook interaction, COVID, bias, reliability**

I. Introduction

The proliferation of coronavirus (COVID-19) has impacted many aspects of people's lives on a global scale. Their daily consumption of news is not an exception [1]. For example, Facebook, a source of news for 32% of Americans [1], has seen significant growth in the number of members and traffic over this time period. Yet, little is known about how people's interactions with various types of news on Facebook have changed during the *covid* period.

The primary factors considered by previous works profiling the news consumption ecosystem [2], [3] are news reliability and bias. Here, reliability typically captures the degree of factual reporting, and bias captures the tendency of journalists to (intentionally or not) favor one political side or the other.

The recent pioneering work by Altay et al. [3] studies the temporal dynamics of user interactions among different reliability classes. However, the work only considered the reliability aspect and only presents aggregated results in which they combine all news published by any news outlet associated with a reliability class. In this paper, we also study the temporal dynamics of Facebook user engagement, but quantify the impact of news sources labeled based on both their bias and their reliability. Furthermore, we consider both per-publisher results (capturing the impact on all publishers in a bias/reliability category) as well as the aggregated results.

Our work is also the first to study and compare the temporal dynamics of the number of followers for different publisher (outlet) classes based on both the outlets' perceived reliability and political bias, as well as the interactions associated with the news published by each such publisher. The number of followers plays a central role in the Facebook recommender system, impacting how many people are exposed to different kinds of news, but also captures how the popularity (of subscribing) to different classes of news outlets may have changed over time. With our analysis spanning the 4.5 years of Jan. 2018 to July 2022, we place particular focus on the impact that the covid pandemic has had on these different classes.

A. Approach and Research Questions

To study the temporal dynamics for different classes of publishers, we compiled a list of all U.S news publishers evaluated by Media Bias/Fact Check (MBFC) [4] that had an official Facebook page. Using the Facebook Crowdtangle platform, we then compiled longitudinal panel data for each such publisher, spanning roughly 4.5 years (Jan. 2018 to July 2022). By applying a combination of statistical and time-series analysis to the above panel data, we address the following research questions: How have the temporal dynamics of user engagement (including both interactions and followers) changed pre-, post-, and during the pandemic? For example, when did patterns shift? How has COVID contributed to the situation? How have the temporal dynamics of user engagement changed over the studied period in relation to the different classes of bias and reliability? We first ask these questions from the perspective of the individual publishers associated with different classes of political bias and reliability (**RQ1**) and then in the aggregate context where the individual news outlets are ignored, but only their classification is accounted for (**RQ2**).

The definitive version was published in *Proc. International Conference on Social Networks Analysis, Management and Security (SNAMS)*, Milan, Italy, Nov/Dec. 2022. https://doi.org/tbd.

Here, it is important to note that the two research questions are highly complementing, with the first question focusing on the full spectrum of publishers, whereas the second (due to a very high skew in popularity) gives almost all weight to the most popular publishers. For example, in our dataset, the top 14% of the publishers are responsible for 80% of the interactions. The first question is, therefore, better at capturing what an average publisher sees, and the second question is better at capturing what an average user may see. Finally, we note that we are the first to address RQ1 and that prior work only has addressed reliability aspects of RQ2 (for RQ2 we, therefore, focus primarily on the bias).

B. Motivating Example

Fig. 1 shows the relative growth in the number of followers of the official Facebook pages of The New York Times, Fox News, as well as the average growth seen by all publishers studied here. All numbers presented here are relative to the number of followers they had in Jan. 2018 (expressed as a percentage). In the figure, we highlight three distinct zones, referred to as the pre-covid, covid, and *post-covid* time periods. (These regions were identified using time-series analysis, with the details of determining the time thresholds for these time periods being described in Section II-D.) We note that the general population (blue curve with confidence interval at $\alpha = 0.1$ shown in shaded blue) grew more rapidly during the *covid* period compared to the *pre-covid* period and then subsequently (during the *post-covid* period) returned to its *pre-covid* growth rate. Looking at the two example publishers, we note that Fox News (biased towards the right on the political spectrum) sees a much greater covid boost than The New York Times (biased towards the left). Our research questions aim to address to what extent these observations are generalizable and to better understand which observations are statistically significant. We also consider a broad spectrum of biases and reliability classes, and provide fine-grained insights into these dynamics.

Outline: Section II describes how we collected and analyzed the data. Section III-A presents per-publisher results (**RQ1**) and Section III-B presents the aggregate results (**RQ2**). Related work is discussed in Section IV before Section V concludes the paper.

II. Methodology and Dataset

Here, we describe the methodology and dataset used in our analysis. First, Section II-A describes how we selected and labeled the publishers studied in this paper. As noted in section I, we examine, among the several outlets of a publisher, the temporal dynamics of the interactions with its Facebook page outlet. Therefore, Sections II-B and II-C explain how the Facebook pages of these publishers were collected and how the temporal statistics for each outlet were compiled, respectively. Next, Section II-D describes how we identified the three time periods considered (*pre-covid*, *covid*, *post-covid*). Finally, Section II-E describes the statistical tests used in the different parts of our analysis.

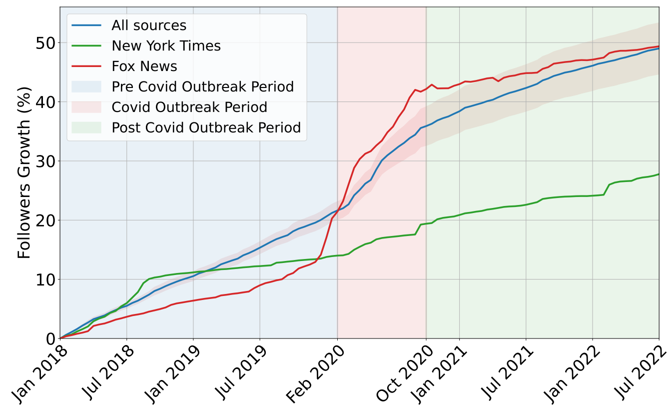

Fig. 1. Growth dynamics of followers

A. Outlet selection and labeling

Not all news sources have the same political bias or publish equally reliable news. Today, there are multiple independent organizations that analyze the bias and/or reliability of news sources. Such example initiatives include the Media Bias Fact Check (MBFC) [4], Ad Fontes Media [5], AllSides [6], and News-Guard [7]. Out of these, for the labeling used here, we picked MBFC for the following reasons: (1) Their evaluations consider both the bias and reliability of each outlet. (2) They provide a fine-grained source classification along both bias (eight levels) and reliability dimensions (six levels). (3) With the exception of News-Guard (not free), their dataset is larger than the other sources, and their labeling provides us with sufficient ($n \geq 30$) samples for all classes considered here except one. (4) Their ratings are publicly accessible for research purposes, greatly simplifying the reproduction of our (and others') results.

For the evaluation, we chose all U.S. media sources examined by MBFC until June, 21, 2022. Limiting ourselves to U.S. publications ensures that all sources more naturally can be mapped to the same political bias spectrum and that reliability is judged more similarly than if including sources from multiple countries. We note that most countries have vastly different political systems and that each country comes with its own partisan divisions.

To label the news sources, we use a variety of values provided by MBFC, including the "bias", the "accurate reporting" measure (referred to in this study as "reliability"), and the "traffic" volume associated with the source.

Bias: Our samples span five classes of political bias classes (from left to right), as labeled by MBFC: "Left," "Left-Center," "Least Biased," "Right-Center," and "Right". These five classes are widely used in previous research [2], [3] and will be the classes considered when presenting the results. In addition to these five classes, MBFC also includes the bias labels "Pro-Science", "Conspiracy-Pseudoscience", and "Questionable Sources". While the focus of our study, like that of past research [2], [3], is on the five political bias classes, we offer the analysis of the other ones for the benefit of researchers who are interested via the repository accompanying this paper.

Reliability: MBFC provides six sorted reliability classes: "very high," "high," "mostly factual," "mixed," "low", and "very low" reliability. However, as there are only 3 outlets in the "very low" class in our dataset, we moved those outlets to the "low" reliability class. Therefore, the "low" reliability class in our study encompasses both the "low" and "very low" ranked outlets by MBFC.

Traffic volume: MBFC uses Similar Web [8] statistics on page views to estimate the traffic volume. They then divide the outlets into three classes: those with more than 2.5M page views (high traffic), those with 150K to 2.5M page views (medium traffic), and those with fewer than 150K page views (minimal traffic)(e.g., local newspapers). To prevent this group from skewing our results, we restricted the research to medium and high traffic news publishers. This pruning resulted in a dataset with 1,622 publishers.

B. Compiling Facebook Pages

We next collected the Facebook pages of the selected publishers. Here, we first identified the webpage listed on MBFC's profile page of each publisher. In this step, we carefully reviewed the validity of each publisher's listed webpage and fixed incorrect links (e.g., for the "Addicting Info"). We then acquired the publisher's Facebook profile from the homepage of the publisher's website. If there was no link to the publisher's official Facebook page on the homepage, we manually searched for the publisher's official page using the Facebook search option (e.g., for "The New York Times"). In the case that we could find a verified Facebook page that matched the profile, we used this as the publisher's official page. After excluding the publishers that did not satisfy the above requirements, our final list contained 1502 publishers.

C. Temporal Dynamics of Users' Interactions

We studied the temporal dynamics of user engagement over the time period Jan 1, 2018, to July 15, 2022. For the analysis, we split the time period into 109 roughly two-week-long time intervals, which we call time buckets. More specifically, we used the 1st and 15th of each month to define the start date of each time bucket and calculated the statistics for each such time bucket based on the user engagement seen during the corresponding time interval. We next describe the calculations of these per-bucket statistics.

Here, we let $N_{p,j}^{\text{followers}}$ denote the number of followers for each page p, at the beginning of timebucket j and $N_{p,j}^{\text{interactions}}$ denote the total number of user interactions with the page p during the time interval associated with bucket j. By user interaction, we refer to all types of user reactions permitted by Facebook users, including likes, shares, comments, and emoji reactions, such as Like(s), Comment(s), Wow(s), Sad(s), Angry(s), Love(s), and Haha(s). To compute $N_{p,j}^{\text{interactions}}$, all interactions with all posts $\Pi_{p,j}$ of publisher p during time interval j are considered.

In addition, we tracked interaction rate $R_{p,j}^{\text{interaction}}$ (for page p and buckets j) as follows:

$$R_{p,j}^{\text{interaction}} = \frac{N_{p,j}^{interactions}}{|\Pi_{p,j}| \cdot N_{p,j}^{followers}} \times 100. \tag{1}$$

Note that this rate normalizes the total number of observed interactions compared to both the number of posts and the number of followers. All the above statistics are recorded for all pages $p \in \mathcal{P}$ and buckets $j \in \mathcal{J}$, where $\mathcal{P}$ and $\mathcal{J}$ are the set of publishers and buckets considered, respectively.

To calculate the above statistics, we used per-publisher data that we extracted using the CrowdTangle platform [9]. CrowdTangle, which is owned by Facebook, indexes the posts and engagement data for around 7M public pages, including all pages "with more than 50K likes", all public Facebook groups with 95K+ members, all US-based public groups with 2K+ members, all verified profiles, as well as any pages added to a CrowdTangle list by those with access to it. When compiling the statistics across all pages and time buckets, we removed the pages that were either created after Jan. 2018 or that we could not get all the discussed statistics needed for the given bucket thresholds. Finally, for the per-publisher analysis, we discarded all pages with fewer than 50K followers by Jan. 2018. These smaller publishers more frequently see large fluctuations in the interaction rates and the number of followers, negatively impacting the per-publisher results (Section III-A). Following this step, for the final per-publisher analysis, we have 992 publishers. Tables I and II show the number of outlets in the final dataset for each of the bias and reliability classes. Finally, we note that in the aggregated results (presented in section III-B), all pages are included regardless of the number of followers they have.

Following this step, we have a panel dataset for each of the three variables of interest (i.e., the number of followers, the interaction rate, and the total number of interactions), including values for all outlets and all 109 time bins. Using these per-outlet data, we then compile per-class (e.g., the "right center" class) panel datasets by computing the means over all the outlets belonging to a class and over each of the 109 time bins. These per-class panel datasets are then used for reporting the per-class results and the statistical tests.

D. Identification of Breakpoints

As shown in Fig. 1, we selected Feb. 2020 (the start of the *covid* period) as the first breakpoint. We identified the other breakpoint (Oct. 2020), which we refer to as the end of the outbreak period, by minimizing the total sum of squared errors if the time series is approximated using the mean of each section. This optimization can easily be solved using exhaustive search over all possible breakpoints.

While we have found that the optimal breakpoints differ slightly depending on which panel dataset is used, we selected to use the breakpoints based on the average followers' growth rate parameter when including all the 992 outlets. Choosing to use the same time period for the different parameters simplified head-to-head comparisons.

TABLE I
NUMBER OF SAMPLES IN EACH BIAS CLASS

Class	Left	Left-Center	Least Biased	Right-Center	Right	Pro-Science	Conspiracy-Pseudoscience	Questionable Sources
Outlets #	97	218	408	112	37	50	24	46

TABLE II
NUMBER OF SAMPLES IN EACH RELIABILITY CLASS

Class	Very High	High	Mostly Factual	Mixed	Low
Outlets #	28	757	65	111	31

E. Employed Statistical Tests

Three types of statistical tests are used in this study. First, we use the *two-sample t-test* to be able to compare the means between two independent groups and *paired t-test* when comparing two different means of the same group (e.g., the average interaction rates of a class before and after the covid outbreak). Second, we use the *t-test statistics* t^*_{exp} as detailed in [10] to compare the slopes of each two regressions lines in each period. The OLS implementation of the Statsmodels package [11] is used to model the regression lines. Third, we use the commonly employed in econometrics Chow test [12] to check for the presence of a structural break in a single time series. This test can be used to determine whether a structural break exists at a date supposed to be known beforehand.

When referring to any of the aforementioned test findings, we set the significance level (α) to 0.1 and deem the result significant if $p\text{-}value$ is less than this α.

F. Limitations

We next present the limitations of the dataset. First, our dataset only includes outlets evaluated by MBFC. However, of the services labeling publishers, MBFC was the only one that satisfied our criteria (discussed in section II-A) and has been utilized in several other studies [2]. Second, we excluded publishers without an official Facebook page and publishers with less than 50K followers to their official Facebook page and the ones we could not get the needed statistics on the bucket thresholds. While these exclusions mean that we leave some publishers outside the study, these choices helped remove outlier events associated with small publishers and potential sources of unreliabilities. Third, the reader is encouraged to take into consideration that the "very high" reliability class has a size of less than thirty when interpreting the results, even though this has been taken into consideration for the statistical tests and many of the them could have been passed by this class because of their large deviations from the averages. Finally, we note that due to the space restrictions, we only present panel data using mean values (of all outlets belonging to a class) and used this for our statistical tests. Moreover, we restrict ourself to window sizes of 5 when presenting the moving averages. For the interested reader, we also publish the material of the other aggregating level trends in the repository accompanying this paper, available at https://github.com/alireza-mon/snams2022.

III. RESULTS AND DISCUSSION

We split the analysis into two parts. We first present per-publisher results (Section III-A) that weight each publisher of a class equally and then an aggregated analysis (Section III-B) that combines all interactions associated with any publisher associated with a class (including those of smaller publishers).

A. Publisher Engagement

Our focus here is on the temporal patterns observed by a typical publisher associated with each class. For this analysis, we calculate and report both the increase in the number of followers a class has obtained and the interaction rates observed for that class. Therefore, first, for each page p and time bucket j, we calculate the normalized increase in the number of followers since Jan. 2018 as formulated in (2).

$$\Delta^{\text{followers}}_{p,j} = \frac{(N^{followers}_{p,j} - N^{followers}_{p,1})}{N^{followers}_{p,1}} \times 100 \qquad (2)$$

Then, for each class c we report average statistics as calculated over all publishers of a class (i.e., $\forall p \in \mathcal{P}_c$, where $\mathcal{P}_c$ is the set of publishers labeled to be in class c). Second, to investigate the temporal patterns of users' interaction rates with various types of news, we use the interaction rates $R^{\text{interaction}}_{p,j}$ associated with each publisher p and time bucket j (as defined above) and compute the mean of all publishers in a class for each j.

Bias-based trends in the growth of followers: Fig. 2 shows the average increase in the number of followers $\frac{1}{|\mathcal{P}_c|}\sum_{p\in\mathcal{P}_c}\Delta^{\text{followers}}_{p,j}$ as a function of time. Several observations can be made from Fig. 2. First, the "right-center" class experienced the greatest growth (82%) throughout the entire study period, while the "right" class experienced the least growth. In contrast, both the "left" and the "left-center" classes experienced significantly smaller growth than the overall average (49% as shown in Fig. 1). Yet, like the "right-centered" class outperforming the (most) "right" one, the "left-center" class outperformed the (most) "left" class over the full-time period.

We now turn our focus to the *covid* time period. First, it should be noted that the extreme ("left" and "right") classes experienced the lowest growth during this period which is significantly lower than the others. For example, the "right-center" class experienced a 22% increase in followers in this time period, while the "left" class experienced a 6% gain.

Bias-based trend break comparisons: Even though the breakpoints were determined on the overall aggregated data, all classes experience statistically significant structural breaks between both the *pre-covid* vs. *covid* and *covid* vs. *post-covid* time periods. When looking at the absolute changes, we again see that the "right-center" class gained the most during the *covid* period. For example, the slope of the regression

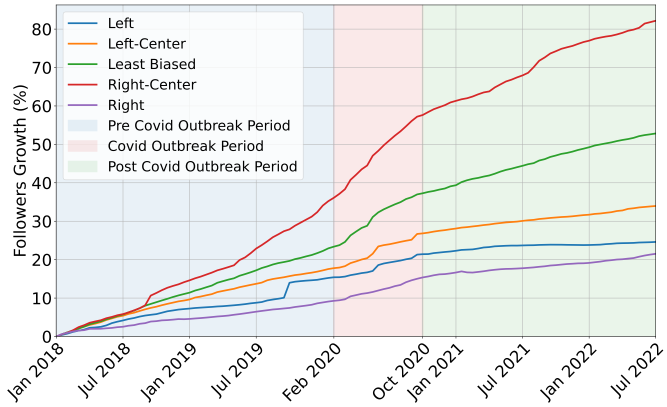

Fig. 2. Followers growth dynamics for different bias classes

TABLE III
SLOPE OF THE FOLLOWERS GROWTH REGRESSION LINE FOR EACH BIAS CLASS AND TIME PERIOD (UNIT = PERCENT/2WEEKS)

Class	Left	Left-Center	Least Biased	Right-Center	Right
Pre	0.28	0.35	0.46	0.69	0.16
Covid	0.41	0.61	0.97	1.44	0.38
Post	0.06	0.16	0.38	0.61	0.13

line using the OLS model in the *pre-covid* period is 0.69 percent/unit (p/u) (each unit being two weeks) but increases to 1.44 for the *covid* period. This class hence experienced more than a doubling (0.75 p/u increase) in the followers' growth rate. Table III summarizes the slopes of the regression lines for all classes and time periods.

When assessing the trend change from the *covid* to *post-covid* time, the absolute drop is greatest for the "least biased" and "right-center" classes, but the relative drops are also significant for the 'other classes. Furthermore, while the growth slope for the "right", "right-center", and "least-biased" classes in the *post-covid* time is reverting to around *pre-covid* time (although the differences are still statistically significant), the *post-covid* growth rates are much lower than the *pre-covid* rates for both left-party classes. Finally, we should mention that in each of the periods, the slopes of the regression lines are statistically different between different classes (except for the "right" and "left" classes during the *covid* period)

Key observation: There is a lower growth rate of followers during *post-covid* than during the *pre-covid* period for all bias classes. Further, compared to the other classes, both left-party groups have experienced significantly higher rate drops when comparing the *post-covid* to the *pre-covid* time period.

Reliability-based follower analysis: We next consider the temporal patterns in the rate that new users followed the publishers associated with different reliability classes. Fig. 3 and Table IV summarize these results.

First, and perhaps most interesting, the two extreme ("low" and "very high") classes experienced the greatest follower growth (∼55%) over the full study period, while the "mostly factual" class experienced the smallest overall growth. Both

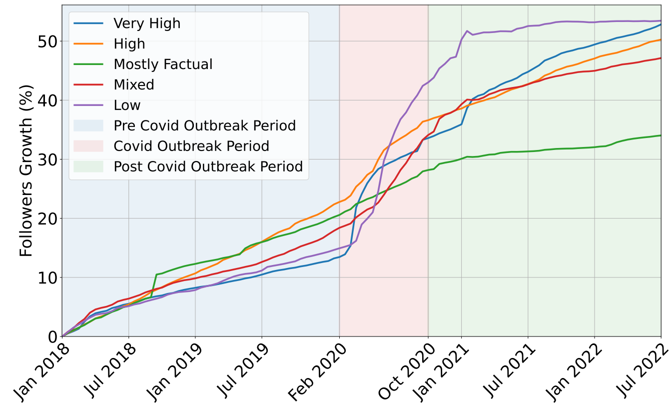

Fig. 3. Followers growth dynamics for different reliability classes

TABLE IV
SLOPE OF THE FOLLOWERS GROWTH REGRESSION LINE FOR EACH RELIABILITY CLASS AND TIME PERIOD (UNIT = PERCENT/2WEEKS)

Class	Very High	High	Mostly Factual	Mixed	Low
Pre	0.23	0.44	0.41	0.31	0.27
Covid	1.25	0.96	0.46	1.02	2.12
Post	0.45	0.34	0.11	0.26	0.17

of the extreme classes received much of their gain during the *covid* period. Second, we observe significant growth differences between different reliability classes when we look at each time period. For example, except for two cases, there is a significant pairwise difference between the slopes of the regression lines of different classes in all time periods (The exceptions are "very high" vs. "mixed" and "high" vs. "mixed" classes both during the *covid* period). Table IV summarizes the slopes of the corresponding regression lines.

Finally, there are significant structural breaks for all the classes when looking at both *pre-covid* vs. *covid* (except for the "mostly factual" class) and *covid* vs. *post-covid*. Here, we again observe some classes with substantially bigger growth rates during the *covid* period. Interestingly, the two extreme cases of "low" and "very high" reliability experienced the largest increases during the *covid* period. Another interesting observation is that the "very high" reliability class experienced its sharpest growth at the beginning of the *covid* period, while the increase for the "low" reliability class lags by roughly two months. This suggests that people initially may have turned to reliable sources and a large number of users later subscribing to the least reliable publishers. In contrast to the above extremes, the "mostly factual" class experienced the smallest increase during the *covid* period and saw the biggest drop in follower growth after the *covid* period (i.e., comparing *pre-covid* vs. *post-covid*). Perhaps the most positive observation here is that the two most reliable classes (i.e., "very high" and "high") are experiencing the greatest growth in the *post-covid* period and the "low" reliability class saw the biggest absolute drop (from 2.12 p/u to 0.17 p/u).

Key observation: In terms of the number of followers, the two most reliable classes experienced lower growth rates during the *covid* period than the "low" reliability class but are experiencing the greatest followers growth during the *post-covid* period.

Bias-based analysis of the temporal trends in the interaction rate: We next examine the publisher based temporal patterns of the interaction rate. Here, we first consider the impact of bias. Fig. 4 summarizes how the moving averages of interaction rates of each bias class (as well as when all sources belonging to any of the five bias classes taken into account) have changed over time.

A noteworthy finding is that none of the classes have a statistically higher mean interaction rate at the end of the *covid* period than they did at the beginning. Looking closer at the *post-covid* time period, it can be clearly observed that, except for certain seasonal components (e.g., the rise in Q4 of each year) and certain events (e.g., the Ukraine conflict), the interaction rates of all classes are decreasing (Table V summarizes the slopes of the regression lines for different bias classes and time periods). There are two potential explanations for this behavior. First, the engagement rate of *pre-covid* users has dropped during the *covid* and *post-covid* time periods. Second, the new population of users associated with each class (during the *covid* and *post-covid* periods) have a lower interaction rate than the general population of followers of the publishers that they have started to follow. In this case, these newly added users would contribute less to the numerator of the interaction rate than to the denominator. We refer to this effect as a deflation effect. Because of Facebook's privacy policies, we cannot isolate the two groups of users and can therefore not answer which of these factors (if any) contributed the most to the observed results.

We also note that, except for the "left" class , the descending trend in interaction rate started already during the *covid* period. While the changes between the time periods are less visible for the interaction-rate analysis, the structural break is statistically significant for all classes and shifts between time periods.

TABLE V
SLOPE OF THE INTERACTION RATE REGRESSION LINE FOR EACH BIAS CLASS AND TIME PERIOD (UNIT = PERCENT/2WEEKS)

Class	Left	Left-Center	Least Biased	Right-Center	Right
Pre	0.04	0.05	0.12	0.24	0.21
Covid	0.22	-0.02	-0.11	0.0	-0.12
Post	-0.20	-0.11	-0.14	-0.18	-0.22

TABLE VI
SLOPE OF THE INTERACTION RATE REGRESSION LINE FOR EACH RELIABILITY CLASS AND TIME PERIOD (UNIT = PERCENT/2WEEKS)

Class	Very High	High	Mostly Factual	Mixed	Low
Pre	0.04	0.11	0.09	0.11	0.02
Covid	-0.01	-0.07	-0.11	0.32	0.60
Post	-0.13	-0.13	-0.19	-0.27	-0.36

Here, the largest changes in slope are observed for "right" and "right-center" classes when comparing the (positive) slopes *pre-covid* and with the (negative) slopes *post-covid*. For example, for these two classes, the rate changes are around -0.43. In contrast, the "left center" publishers experienced the least change in their interaction rate. The final point to note is that when focusing on each period, the differences between classes are less significant here (with exception of *pre-covid* period).

Key observation: For all the five bias classes and accordingly their general population, the interaction rate at the end of our study (July 2022) is significantly lower than at the beginning of our study period (Jan. 2018), even though the trend was increasing during the *pre-covid* time period.

Reliability-based analysis of the temporal trends in the interaction rate: We now again turn to the impact of reliability. Fig. 5 presents a longitudinal view of the results moving averages, and Table VI summarizes the slopes of the regression lines.

Perhaps the most concerning observation is that the two lowest reliability classes (and only these classes!) exhibit significant positive interaction rates during the *covid* period. (For the "mixed" class the slope is positive due to the ascending

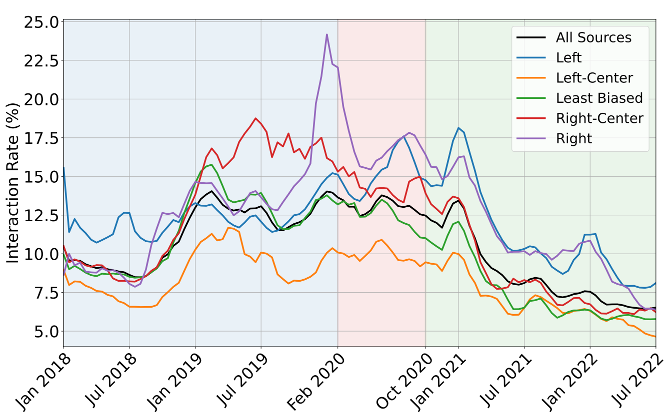

Fig. 4. Interaction rate dynamics for different bias classes

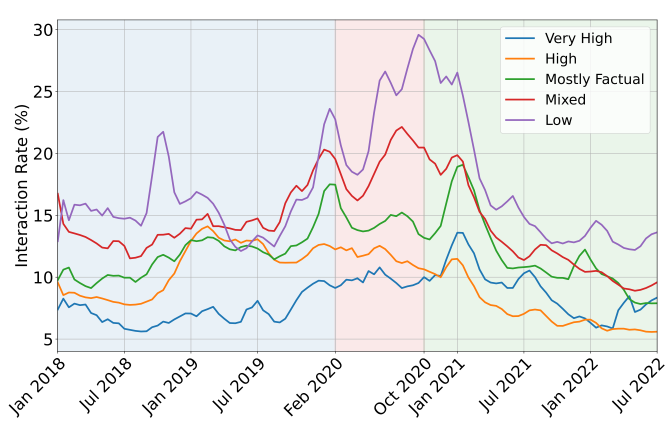

Fig. 5. Interaction rate dynamics for different reliability classes

middle segment of the time series during the *covid* period.) The decreasing interaction rates seen during *post-covid* appear to be triggered after a temporary increase in 2020-Q4 (mainly linked to the U.S. elections and what followed) and then keeps falling to below *pre-covid* values.

Also for this analysis, the structural breaks for all the classes are statistically significant (except for the “high” class during the *covid* to *post-covid* transition and the “mostly factual” class during *pre-covid* to *covid* transition). Furthermore, while the least reliable classes have observed the highest interaction rates (almost throughout the whole studied time period), we should also consider comparing the differences of rates between the *pre-covid* and *post-covid* time periods. For example, when comparing the slopes for the *pre-covid* and *post-covid* time periods, we observe a negative correlation between the slope change and level of reliability.

Key observation: In terms of interaction rate, there is negatively correlation between the reliability level and the covid impact. In other words, high reliability classes are much less affected.

B. Aggregated Analysis

Thus far, we have treated all publishers within a class equally. In this section, we present the aggregated results in which all posts published by any publisher labeled as belonging to a class c are treated equally. For this analysis, we did not place any restrictions on the minimum number of followers of a publisher but simply added them all to the set of publishers associated with the class. The number of followers and the increase in the number of followers are hence calculated only on a per-class basis. For example, $N^{followers}_{\mathcal{P}_c,j} = \sum_{p\in\mathcal{P}_c} N^{followers}_{p,j}$ corresponds to the number of followers time series of class c and $N^{interactions}_{\mathcal{P}_c,j} = \sum_{p\in\mathcal{P}_c} N^{interactions}_{p,j}$ is used to compute the number of interactions timeseries for that class. These values are calculated over all interactions related to any publisher in the set $\mathcal{P}_c$ in time bucket j. Otherwise, equations (1) and (2) can be used seamlessly to calculate these per-class statistics. Again, these results form a single time series for each parameter and class. Due to space constraints, we limit this analysis to political bias.

Interaction changes: Fig. 6 presents the aggregated results of the total interactions for different classes of bias. To better understand the discrepancy in the patterns highlighted, we plot the growth as a percentage compared to that observed in Jan. 2018 (i.e., the start of the analysis period). As a reference point, we include the number of interactions in Jan. 2018 (within brackets) in the figure legend.

Following the same approach outlined in Section II-D, we first needed to decide on the periods. Here, we deciced on four time periods. We call these *pre-covid*, *covid*, *post-covid*, and *after-post-covid*. With the exception of the period surrounding the U.S. elections (and the events that followed), there is a clear downward trend in the total number of interactions during

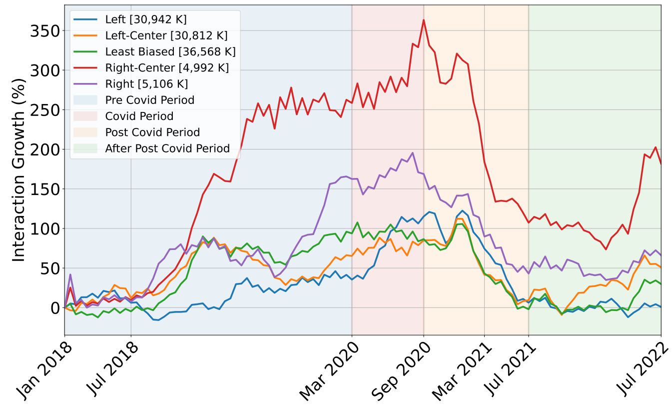

Fig. 6. Interaction growth dynamics for different bias classes

the *post-covid* period. After this period, during the *after-post-covid* period, we have seen a rate convergence followed by an upwards trend. This uptrend, starting in Mar. 2022, can be attributed to the Ukraine conflict discussions.

Follower changes: Fig.7 illustrates the temporal dynamics of the total number of followers. In accordance with Jan. 2018 results (see numbers in the figure legend), the “left-center” class had 280M followers and formed the largest community at that time. Let us now examine the temporal changes that occurred between Mar. 2020 and July 2021 in both Fig. 6 and Fig. 7. Interestingly, the total number of interactions among all classes has decreased at the end of the *post-covid* period as compared to the beginning of the *covid* period. The “right-center” class experienced the most significant decrease and has returned to its Nov. 2018 levels. As we can see from the *after-post-covid* period levels, the deflation effect has converged across all classes. In Mar. 2022, the interactions for all classes began to increase again. It is once again the “right-center” class that holds the largest increase gain. Moreover, except for the “left” class, all other classes in our panel experienced an increase in the total number of interactions between Jan. 2018 to July 2022. However, these values for all classes are significantly smaller than that at beginning of the *covid* period. In addition, by analyzing the total number of interactions and followers for all classes, we observe that the group ranking changes between the start and end of our study timeline (Jan. 2018 to July 2022). Both rankings indicate that the “right” and “right center” classes have changed positions. Accordingly, the “right center” class has a greater number of followers and interactions at the end of our timeline. Both the “left” and “left center”display the same pattern in terms of the number of interactions.

Key observation: For all bias classes, total interactions at the end of the *post-covid* period is significantly lower than that at the beginning of the *covid* period.

IV. Related Work

This study is related to a few recently published works that try to quantify social media users’ engagement with various

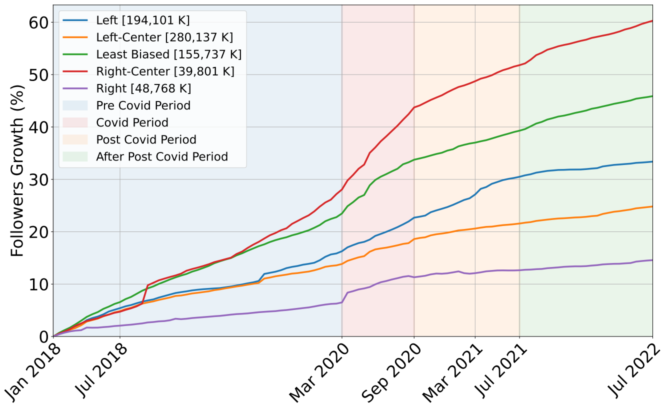

Fig. 7. Followers growth dynamics for different bias classes

types of news. For example, Edelson et al. [2] investigate how Facebook users interact with different bias and reliability classes of news. In addition to aggregated analysis, they provide per-publisher results with the goal of quantifying how individual publishers engage with their audiences. Their per-publisher analysis, for example, shows that in general the median engagement of the left party is significantly lower than that of the right party. Galen et al. [13] carried out a similar investigation as Edelson et al. on Reddit rather than Facebook. Their results show that low-factual content receives 20% fewer upvotes and 30% fewer cross-posting exposures than more factual information. Allen et al. [14] revealed that 15% of the total interactions on Facebook are for misinformation or hyperpartisan material and Guess et al. [15] report that articles from trustworthy news sources are shared over 5.5 times more frequently than articles from low-credibility news sources.

Switching the study scope to Twitter users, Osmundsen et al. [16] show that republicans were more likely than democrats to share fake news sources. Grinberg et al. [17] also do a study on Twitter users during the 2016 U.S elections. Their result shows that 0.1% of users were responsible for sharing 80% of the misinformation.

In another line of research, Allcott et al. [18] examine how users engage with fake news information and websites. Their findings indicate that through the end of 2016, user interactions with fraudulent information increased consistently on both Facebook and Twitter.

None of the above works studied the temporal dynamics of user interactions. Recently, Altay et al. [3] quantified the pandemic's effect on people's online news consumption. Using temporal data from 2017 to 2021 they studied the effects of the covid outbreak on the temporal dynamics of user reaction in different classes of reliability. However, they did not consider bias and no per-publisher based analysis was provided.

V. Conclusions

In this paper, we studied the temporal dynamics of user engagement with various news types. Using a longitudinal dataset capturing Facebook users' engagement with news publishers of diverse types of bias and reliability, we studied both the effects on follower counts and the user interactions with the published news associated with different classes of publishers. Our analysis was performed both on a per-publisher basis (**RQ1**) and at an aggregate level (**RQ2**).

In the reported results, we quantified the temporal dynamics of each class, and using statistical tests, highlighted differences in how different classes were impacted by the *covid* period and how they have kept up during the *post-covid* period. For example, after identifying statistically significant structural breaks in the temporal dynamics resulting from the covid outbreak, we quantified these effects for the different classes of publishers. Here, we have found that the least reliable news saw significant temporal gains during the *covid* period and that the reliability level of the news negatively correlates with the *covid* period's impact on the interaction rate. However, looking at the *post-covid* period, the results are more encouraging. For example, considering the per-publisher-based study, it has been demonstrated that the two most reliable classes are experiencing the greatest increase in followers growth during the *post-covid* period. More broadly, the interaction rate level at the end of the *post-covid* period is lower than the level at the beginning of the *covid* period for all bias classes. Combined, the observations and insights highlighted throughout the paper provide a unique perspective into the pandemic effects on how Facebook users have engaged with publishers of different political bias and reliability classes.

References

[1] N. Newman, R. Fletcher, A. Schulz, S. Andi, C. T. Robertson, and R. K. Nielsen, "Reuters institute digital news report 2021," *Reuters*, 2021.
[2] L. Edelson, M.-K. Nguyen, I. Goldstein, O. Goga, D. McCoy, and T. Lauinger, "Understanding engagement with us (mis) information news sources on facebook," in *21st ACM IMC conference*, 2021.
[3] S. Altay, R. K. Nielsen, and R. Fletcher, "Quantifying the "infodemic": People turned to trustworthy news outlets during the 2020 coronavirus pandemic," *Quantitative Description*, vol. 2, 2022.
[4] "Media bias/fact check," https://mediabiasfactcheck.com.
[5] "Md fontes media," https://adfontesmedia.com.
[6] "All sides," https://www.allsides.com.
[7] "Newsguard," https://www.newsguardtech.com.
[8] "Similar web," https://www.similarweb.com.
[9] "Crowdtangle," https://crowdtangle.com/.
[10] J. Andrade and M. Estévez-Pérez, "Statistical comparison of the slopes of two regression lines: A tutorial," *Analytica Chimica Acta*, vol. 838, pp. 1–12, 2014.
[11] S. Seabold and J. Perktold, "Statsmodels: Econometric and statistical modeling with python," in *Proc. Python in Science*, vol. 57, 2010.
[12] G. C. Chow, "Tests of equality between sets of coefficients in two linear regressions," *Econometrica*, pp. 591–605, 1960.
[13] G. Weld, M. Glenski, and T. Althoff, "Political Bias and Factualness in News Sharing across more than 100,000 Communities," *ICWSM*, 2021.
[14] J. Allen, M. Mobius, D. M. Rothschild, and D. J. Watts, "Research note: Examining potential bias in large-scale censored data," *Harvard Kennedy School Misinformation Review*, 2021.
[15] A. Guess, K. Aslett, J. Tucker, R. Bonneau, and J. Nagler, "Cracking open the news feed: Exploring what us facebook users see and share with large-scale platform data," *Quantitative Description*, vol. 1, 2021.
[16] M. Osmundsen, A. Bor, P. B. Vahlstrup, A. Bechmann, and M. B. Petersen, "Partisan polarization is the primary psychological motivation behind political fake news sharing on twitter," *American Political Science Review*, vol. 115, no. 3, pp. 999–1015, 2021.
[17] N. Grinberg, K. Joseph, L. Friedland, B. Swire-Thompson, and D. Lazer, "Fake news on twitter during the 2016 us presidential election," *Science*, vol. 363, no. 6425, pp. 374–378, 2019.
[18] H. Allcott, M. Gentzkow, and C. Yu, "Trends in the diffusion of misinformation on social media," *Research and Politics*, vol. 6, 2019.