

Understanding Engagement with U.S. (Mis)Information News Sources on Facebook

Laura Edelson*
New York University
New York, NY, USA
lj992@nyu.edu

Minh-Kha Nguyen*
Université Grenoble Alpes
Grenoble, France

Ian Goldstein
New York University
New York, NY, USA

Oana Goga
Université Grenoble Alpes, CNRS,
Inria, Grenoble INP, LIG
Grenoble, France

Damon McCoy
New York University
New York, NY, USA

Tobias Lauinger
New York University
New York, NY, USA

ABSTRACT

Facebook has become an important platform for news publishers to promote their work and engage with their readers. Some news pages on Facebook have a reputation for consistently low factualness in their reporting, and there is concern that Facebook allows their misinformation to reach large audiences. To date, there is remarkably little empirical data about how often users “like,” comment and share content from news pages on Facebook, how user engagement compares between sources that have a reputation for misinformation and those that do not, and how the political leaning of the source impacts the equation. In this work, we propose a methodology to generate a list of news publishers’ official Facebook pages annotated with their partisanship and (mis)information status based on third-party evaluations, and collect engagement data for the 7.5 M posts that 2,551 U.S. news publishers made on their pages during the 2020 U.S. presidential election. We propose three metrics to study engagement (1) across the Facebook news ecosystem, (2) between (mis)information providers and their audiences, and (3) with individual pieces of content from (mis)information providers. Our results show that misinformation news sources receive widespread engagement on Facebook, accounting for 68.1% of all engagement with far-right news providers, followed by 37.7% on the far left. Individual posts from misinformation news providers receive consistently higher median engagement than non-misinformation in every partisanship group. While most prevalent on the far right, misinformation appears to be an issue across the political spectrum.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; • **Information systems** → **Social networks**.

KEYWORDS

Facebook, news, misinformation, engagement, measurement.

*Laura Edelson and Minh-Kha Nguyen contributed equally to this work.

IMC '21, November 2–4, 2021, Virtual Event

© 2021 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event*, <https://doi.org/10.1145/3487552.3487859>.

ACM Reference Format:

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 *ACM Internet Measurement Conference (IMC '21), November 2–4, 2021, Virtual Event*. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3487552.3487859>

1 INTRODUCTION

After the 2016 U.S. presidential election, there was broad public concern [3, 14] about the impact that online misinformation might have had on public confidence in the fairness of the American electoral system. In response to scrutiny from lawmakers and their users, Facebook announced several initiatives [22, 23] aimed at reducing misinformation on its platforms.

To date, there is little public data about how widespread the misinformation problem on Facebook is. Prior work about misinformation on digital platforms has focused either on mechanisms of spread [13, 18, 21], or on absolute measurements of fake news [15, 19]. With few notable exceptions [34], research has not widely studied the interplay of misinformation and partisanship on digital platforms.

In this work, we aim to shed light on user engagement within the news ecosystem on Facebook. To the best of our knowledge, we are the first to characterize engagement based on the political leaning and factualness of news sources. We explore engagement with (mis)information news from three different perspectives:

- (1) What share of overall engagement with U.S. news sources is taken up by misinformation providers?
- (2) How well do individual sources of misinformation news engage with their audiences, compared to more factual outlets?
- (3) How well do users engage with content from misinformation sources when compared to other news content?

A major challenge to studying these questions is data availability. We need lists of news publishers with good coverage of the ecosystem. We also need partisanship and factualness attributes for each publisher, that is, their political leaning or bias, and whether the source has a reputation for regularly spreading fake news, conspiracy theories, or misinformation. Our proposed methodology derives these attributes from lists of U.S. news publishers acquired from two third-party data providers, NewsGuard and Media Bias/Fact Check. We use Facebook’s CrowdTangle tool to collect 7.5 M public posts and engagement metadata from 2,551 official Facebook pages

associated with those sources during the 2020 U.S. presidential election, from August 10th, 2020 to January 11th, 2021.

To answer our three research questions, we propose three metrics based on the data available through CrowdTangle: (1) total engagement across the ecosystem, (2) per-page engagement normalized by followers, and (3) per-post engagement, where engagement is the number of comments, shares, and reactions such as “like.” Notably absent from CrowdTangle is impression data. We explore video views as an alternative, but show that is not a good substitute.

Our results shed light on several aspects of how misinformation compares to non-misinformation in the Facebook news ecosystem. In all partisan categories, misinformation sources are able to generate higher median per-post engagement, and this difference is statistically significant in all cases. Overall, posts from misinformation sources out-engage those from non-misinformation sources by a factor of six. However, in absolute terms, content from misinformation sources makes up a majority of engagement only for the Far Right, 1.23 B interactions compared to 575 M, or 68.1 %, followed by the Far Left with 37.7 % of engagement coming from misinformation sources. This engagement with misinformation is driven by 109 pages (41.4 %) on the Far Right, and only 16 pages (8.6 %) on the Far Left. (In all other partisan categories, the percentage of misinformation pages is below 6.1 %.) When we compare how much engagement pages generate relative to their number of followers, the picture is mixed. Far Left, Slightly Right, and Far Right misinformation sources out-perform their non-misinformation counterparts, while the reverse is true for Slightly Left and Center pages.

The primary implication of our findings is that it is not only Far Right misinformation that is of concern on Facebook. Misinformation appears to confer a per-post engagement advantage no matter the political leaning of the misinformation source. It is an important question for future research to study why this is the case.

Our work makes the following contributions:

- (1) We propose a methodology for harmonizing multiple news source lists and news quality evaluations.
- (2) We introduce metrics for measuring user engagement with (mis)information from three perspectives, which can serve in the future to measure changes in the news ecosystem and evaluate countermeasures.
- (3) We measure the relative scale and scope of misinformation within the larger context of news on Facebook, and show that it often outperforms non-misinformation.

2 BACKGROUND

Facebook has become an important channel where news publishers promote their content. To do so, news outlets set up a Facebook page where they can publish various types of posts. In the context of news publishers, such posts most commonly link to articles on the web, such as the publishers’ own websites (e.g., Figure 11b in the appendix). Posts can also contain images (such as memes) or prerecorded or live video (Figure 11c).

Even though users could visit these publishers’ Facebook pages manually to read their content, they typically encounter the pages’ posts on their own timeline when Facebook’s algorithms deem them interesting to the user. Among many signals that determine whether a post will be selected are the user “following” a page, how

“engaging” the post has been so far for other users, and Facebook’s prediction of how engaging the post will be to the specific user. Thus, in order to reach a wider audience, Facebook pages need to publish content that is engaging.

Facebook users can engage with content in a range of ways. Users can *react* to a post by clicking the “like” button or one of various alternatives such as “angry” or “sad.” Users can also *share* a post with their friends, and they can write *comments*. We collectively refer to all three types of engagement as *interactions*.

We consider misinformation to be information that is false or misleading [20], regardless of the intent of the author or distributor of the information. We consider misinformation to include disinformation [33], which refers to false or misleading content that is communicated with the intent to deceive. False information communicated in error or by ignorance also falls under the umbrella of misinformation. Misinformation can have a partisan focus, such as right-leaning stories casting doubt on the legitimacy of the 2020 U.S. election, but it also exists in left-leaning or apolitical/center circles, such as misinformation related to environmental causes or the efficacy of vaccines. Figure 10 in the appendix shows misinformation examples across the political spectrum.

3 METHODOLOGY

To study differences in how users engage with (mis)information providers on Facebook, we first obtain manually curated lists of U.S. news publishers. These lists also rate the typical news quality or factualness and political leaning of each publisher. We then discover these news publishers’ official Facebook pages, and extract all public posts along with metadata showing how many Facebook users interacted with these posts.

3.1 News Publishers

We obtained lists of news sources from Media Bias/Fact Check (MB/FC) [6] and NewsGuard (NG) [24] in July 2020. Those two independent data providers have also been used in prior work [29, 34]. NewsGuard create detailed trust ratings for news websites that they refer to as “nutritional labels.” These evaluations are available only to paying customers. Media Bias/Fact Check consider themselves an independent online media outlet “dedicated to educating the public on media bias and deceptive news practices.” While their evaluations of news sources are slightly less comprehensive than NewsGuard’s, they are publicly available. Both lists assign various attributes to each news source; for the purposes of this study, we use the partisanship and factualness labels. *Partisanship* is the political leaning, or bias of the news source. *Factualness* is whether the source has a reputation of regularly spreading misinformation or conspiracy theories. While these evaluations are qualitative and inherently subjective, both data sources follow public methodologies [7, 25] in their labelling and have broadly similar criteria for these two attributes.

We extracted these news source evaluations once at the end of the study period and consider them to be static. In line with prior work [2, 13, 15, 16, 29, 34], we are using classifications at the level of news publishers instead of individually fact-checked news articles. For the purposes of this study, we are most interested in the behavior of news publishers, rather than individual pieces of

content. In the Facebook ecosystem, the reputation of news sources is most relevant since users follow the pages of news publishers and Facebook’s algorithms use these follower relationships as a basis to show articles to users. Furthermore, this approach is more scalable and broader than considering only fact-checked articles. We conservatively apply the misinformation label only to news publishers that have a reputation for repeatedly posting misinformation, and aim to capture all user engagement with this news source. Fact checkers typically focus their attention on high-engagement and verifiable content [10, 27], thus it is likely that not all posts of misinformation providers would be fact checked, and the set of fact-checked articles would not be a representative sample of articles containing misinformation.

In total, NewsGuard had evaluations for 4,660 news sources, and we were able to collect another 2,860 evaluations for news sources from the Media Bias/Fact Check website. Not all of these news sources were relevant for our study, and some list entries were lacking partisanship or factualness attributes, thus we filtered and merged the two lists as follows.

3.1.1 U.S. Publishers. Countries have very different political systems with a different political spectrum or partisan divide. Therefore, we restrict our analysis to a single country and discard 1,047 non-U.S. news sources from NG, and 342 from MB/FC.

3.1.2 Facebook Page. Our analysis is based on user engagement with posts on the publishers’ official Facebook pages. The NewsGuard data set contains the primary Facebook page for some sources, but not all of them, and in some cases, multiple list entries share the same Facebook page. The Media Bias/Fact Check data does not contain any references to Facebook pages at all. We combined duplicate list entries with the same Facebook page in NG, which resulted in the removal of 584 entries. To fill in missing Facebook page information, we queried Facebook for domain-verified Facebook pages having a domain name matching the primary Internet domain name of the news publisher from the list. We had to remove 883 entries from NG and 795 from MB/FC because we were unable to find matching Facebook pages.

3.1.3 Political Leaning. To study engagement with (mis)information providers based on their political leaning, we need partisanship attributes for each news source. While NG considers all news sources without partisanship information as *center*, we discarded 89 entries in MB/FC because they had no partisanship data. (Most of them were labelled as pro-science or conspiracy-pseudoscience.)

Both lists classify partisanship into different categories. We translate them into a common categorization of *Far Left*, *Slightly Left*, *Center*, *Slightly Right*, and *Far Right* as detailed in Table 1. In the 701 cases where we had both a NG and MB/FC evaluation, we gave preference to the latter. In these cases, the two lists only agreed 49.35% of the time. However, most of the differences were slight. Just over 34.24% of cases disagreed between center and either slightly left or slightly right, and 10.41% between slightly left and far left, or slightly right and far right. Overall, neither NG nor MB/FC appear to have a significant left-ward or right-ward skew compared to the other in terms of the distribution of their partisanship evaluations. NewsGuard has a bias toward the center of the distribution, rating 79.7% of news sources as *Center*, compared with 34.11% for MB/FC.

Combined	NewsGuard	Media Bias/Fact Check
Far Left	Far Left	Left, Far Left, Extreme Left
Slightly Left	Slightly Left	Left-Center
Center	N/A	Center
Slightly Right	Slightly Right	Right-Center
Far Right	Far Right	Right, Far Right, Extr. Right

Table 1: Mapping of partisanship labels of the two news publisher lists to our harmonized political leaning attribute.

3.1.4 (Mis)Information. For this study, we require a boolean *misinformation* flag for each news publisher, representing whether the publisher has a history of spreading misinformation, fake news, or conspiracy theories. Both NG and MB/FC use different terms to capture the spectrum of misleading or questionable news practices, but at the far end of the spectrum, the terms “Conspiracy,” “Fake News,” or “Misinformation” were used by both data providers. NewsGuard describes this information in the “Topics” column of their data file, while MB/FC reports questionable news practices in the “Detailed” section of the source evaluation on their website. If any of the three terms above is used to describe a news publisher, we apply the misinformation label to the publisher. NewsGuard and MB/FC were in broad agreement about this measure; we had an evaluation from both lists for 679 publishers, including only 33 disagreements where we broke the tie by applying the misinformation label.

3.1.5 Minimum Page Follower and Interaction Thresholds. To avoid a skew in our results that might come from news publishers with minimal posting activity or reach of their Facebook pages, we remove pages that never reached 100 followers during the study period, a total of 15 pages from NG, and 19 from MB/FC. We also remove pages that average fewer than 100 interactions per week, another 187 pages from NG and 343 from MB/FC.

3.2 Coverage of Publishers

After all filtering steps, our final data set of news publisher Facebook pages consists of 1,944 pages originating from NG, and 1,272 from MB/FC, or 2,551 unique pages in total. The low overlap of only 665 pages is in line with prior work comparing (mis)information publisher lists [5] (although their study did not include NG).

The overall partisan composition and list provenance of our final news publisher data set is shown in Figure 1. The x axis of the upper row shows that most of the publishers are classified as center, followed by slightly left. In terms of which news publisher list contributed a Facebook page to our combined data set, as shown on the y axis, NG has higher coverage than MB/FC, including at least half of publishers for each political leaning except for the far right, where NG contained only 47.1%. NewsGuard contributed most of the unique center publishers, whereas MB/FC contributed a higher share of unique pages in all non-center political leanings, especially among the more extreme pages on the two far ends of the spectrum. In relative terms, overlap between the two lists was lowest in the center, and larger in the two left leaning groups than in the two right leaning groups, whereas in absolute terms, it was highest in the center due to the large number of center pages overall. Weighting pages by total interactions or followers, as shown in the middle and bottom rows of the figure, increases the share of the




Figure 1: Composition of the filtered data set by political leaning (horizontal, colored), and origin publisher list (vertical, hatched). In the middle and bottom rows, news publisher pages are weighted by total interactions and followers, respectively. In relative terms, overlap between the two lists is smallest in the center and larger in the extremes. Overlap increases when pages weighted by interactions or followers.

overlapping pages in the extremes and center, but decreases lightly for slightly left and slightly right pages. Weighting also increases the impact of NG-only pages, especially for slightly right pages, and decreases the impact of MB/FC-only pages. This suggests that NG covers publishers with a larger reach than MB/FC.

Out of the total 2,551 pages, a vast majority (2,315 pages) are classified as non-misinformation pages. The proportions between NG and MB/FC-sourced pages stay roughly similar to the proportions for all pages. One exception is the far right when pages are weighted by interactions or followers; in this case, pages only found in the MB/FC list make up a much larger share, while pages from NG and the list overlap decrease their share. Regarding the 236 misinformation pages, MB/FC contributes no unique misinformation pages with a slightly left or slightly right leaning, whereas more than half of center misinformation pages are unique to MB/FC. However, these MB/FC-only pages make up only a small share of the total interactions or followers of center pages. Compared to the overall page provenance distribution, misinformation pages sourced from NG or the list overlap still account for a majority of interactions or followers in each political leaning group, but MB/FC-only pages in the far left account for a larger share, and in the far right for a smaller share of total interactions or followers. (Separate plots of list provenance for non-misinformation and misinformation pages can be found in Figures 12a and 12b in the appendix.)

3.3 Facebook Post Data

We use Facebook’s CrowdTangle API to extract all public posts of news publishers on their Facebook pages, covering the period between 10 August 2020 and 11 January 2021. Metadata includes the total engagement numbers for the post. Specifically, according to the CrowdTangle API documentation [31], this includes the number of top-level comments below the post on the publisher’s

page (excluding responses to comments), the number of times the original post was shared publicly by users, and the number of reactions, such as likes. The numbers only include interactions with the original post on the publisher’s Facebook page. They do not include interactions when posts are shared beyond the publisher’s page. The metadata also include the page’s number of followers at the time the post was published, where followers are defined as the number of users who have chosen to receive updates from a page.

The API contains engagement metrics at various time steps after a post was published. In order to allow for a fair comparison between posts, we use engagement numbers at a two-week delay since posting. Since social media content tends to be short-lived, we expect that posts will not accrue substantial additional engagement beyond those two weeks. Due to scheduling issues during parts of our data collection, we made API requests too early for nearly 1.4% of all posts and have slightly less than two weeks worth of engagement for those posts, varying from 7 to 13 days. Overall, our data set contains 7,504,050 posts and associated metadata.

3.3.1 Video Posts. Unfortunately, CrowdTangle does not contain any data about post impressions, that is, how often a post was shown to users. CrowdTangle does contain view counts for video posts, defined as users having watched at least 3 seconds of the video [8]. We extracted video view counts from the CrowdTangle web portal because they are not available through the API. This video data collection took place on 8 February 2021. We were able to retrieve view count data for 597,844 video posts, missing data of 46 k video posts (7.1%; see Section 3.3.2 for an explanation). We also excluded 291 posts for scheduled live video because these posts cannot have accumulated any video views yet. We note that 415 pages did not publish any video content, and 1,267 pages published video only intermittently but not regularly; only 869 pages published video every week.

We analyze the video data set separately from the overall posts data set. While the video posts are a subset of the overall posts and were published during the same date range, the web portal shows only the latest view count and engagement metrics. Unlike the overall posts data set where we can analyze engagement after a two-week delay since post publication, the time delay between video post publication and observation of the metrics varies from 3 to 25 weeks due to the separate data collection. For this reason, the video data set is not fully comparable to the overall posts data set.

CrowdTangle reports video views separately for the original post, crossposts, and shares of the same video. In order to allow for a fair comparison to the post-based post engagement metrics, we only consider views from the original post. Similarly, we exclude from consideration posts of external (e.g., YouTube) video since it could have been promoted through third-party channels that do not impact the engagement numbers of the Facebook post.

3.3.2 Impact of CrowdTangle Bugs. After completing our initial analysis of the data for this paper, we became aware of at least two bugs in CrowdTangle that affect our work. First, before September 2021, the CrowdTangle API failed to return a subset of posts even though they were available on Facebook. After Facebook fixed the issue [30], we recollected the posts data (Section 3.3) and merged it with our initial data set. This resulted in an additional 627,946 posts, mostly published in August 2020 and after December 24, 2020.

Second, the API sometimes returned identical posts with different CrowdTangle IDs even though they had the same Facebook post ID. We removed 80,895 accidentally duplicated posts. Overall, the updated posts data set has an additional 7.86 % of posts and 7.08 % of additional engagement. The results in this paper (except for Section 4.4) are derived from the updated posts data set, thus they are not affected by the bugs. We did not notice any materially different trends between the two versions of our data set, and the update does not seem to have introduced any strong bias based on political leaning or (mis)information status of a publisher.

We decided not to update our secondary data set containing video views (Section 3.3.1). Recollecting this data set would only give us the latest view count, which is many months after the videos were first published, in contrast to the two-week delay we use for engagement in the remainder of the paper. Concretely, 46 k videos (7.1 % of video posts from the updated posts data set) are not in the video views data set, and thus excluded from our analysis in Section 4.4. The ratios of missing videos range from 6.1 % (center non-misinformation) to 23.0 % (far-right non-misinformation). The second bug causing duplicated posts does not affect the video view data. Because of the different collection method and time, and the effects of missing videos due to the API bug, results from our two data sets should not be compared quantitatively, only qualitatively based on general trends.

3.3.3 Ethical Considerations. We collect our posts data set using CrowdTangle, an official tool with API access granted to us by Facebook. CrowdTangle contains only public posts of Facebook pages and aggregate engagement data, but no personally identifiable information. We do not have access to information about individual news consumers, and cannot quantify the overlap of followers across pages, for instance.

4 ANALYSIS

In this work, we are interested in news providers with a reputation of spreading misinformation, and how much engagement they generate on Facebook, especially in comparison to content from non-misinformation sources. Using three different metrics, we analyze our data set from three different angles: Overall engagement with (mis)information at the level of the news ecosystem on Facebook in Section 4.1, the range of (mis)information publishers and how well they engage with their primary audiences in Section 4.2, and the performance of individual pieces of (mis)information content across the political spectrum in Section 4.3.

4.1 Ecosystem-Wide Engagement

Our first research question aims to shed light on how much engagement (mis)information from news providers generates on Facebook overall. We compute total engagement by summing the number of interactions of all posts across all publisher pages. During our study period, content from the 236 misinformation sources generated 2 B total interactions, compared to 5.4 B interactions for 2,315 non-misinformation sources. Thus, misinformation providers accumulated a sizeable quantity of engagement overall, but less than non-misinformation news sources. The picture is more nuanced, however, when taking into account the partisanship of the source.




Figure 2: Bar plot of total engagement with (mis)information pages from Far Left to Far Right, with the number of Facebook pages on the x axis. In absolute terms, only Far Right misinformation pages accumulated more engagement than their non-misinformation counterparts.

To do so, we segment publishers into groups based on their political bias and (mis)information status. Figure 2 shows that the 1,527 center publishers as a group accumulate the largest amount of interactions, around 2.6 B, over the course of our study period. We also observe that Slightly Left-leaning news sources generate more engagement than Slightly Right-leaning news sources, whereas the opposite is true for the extreme end of the spectrum, where the Far Right generates by far the most engagement outside of the Center.

The finding that Facebook users engage less with misinformation still holds true across the political spectrum from Far Left to Slightly Right, but not for the Far Right. There, only 109 misinformation publishers account for over 1.2 B interactions, which is more than twice the total engagement of the 154 non-misinformation publishers of the Far Right (and more than half of the engagement of the much larger Center non-misinformation group that includes 1,434 publishers). This suggests that relatively small numbers of misinformation sources can drive disproportionately large engagement, also illustrated by the only 16 misinformation pages of the Far Left that generate more than 60 % of the total engagement of their 171 non-misinformation counterparts, and the 11 Slightly Right misinformation pages that generate almost 38 % of the engagement of their 177 non-misinformation counterparts. A counterexample in our data set are the 7 Slightly Left misinformation pages that generate very little overall engagement, less than 0.3 % of the engagement of Slightly Left non-misinformation sources.

In general, reactions contribute most engagement, as shown in Table 2 (with “likes” being the most common reaction). Exact proportions of interaction types, and the gains or losses when comparing misinformation to non-misinformation, fluctuate across the political spectrum.

When looking at post types in Table 3, posts with links to non-Facebook websites are the most common contributor of engagement with non-misinformation publishers, followed by photo posts and Facebook-hosted video. For misinformation publishers, the largest gains come from photo posts. On the Far Left, for example, photo

Total	Far Left	Left	Center	Right	Far Right
Comments (N)	9.79 %	14.1 %	18.3 %	20.6 %	13.3 %
<i>(misinfo.)</i>	-0.42	-8.51	-11.7	-8.10	+3.36
Shares (N)	11.8 %	8.52 %	12.4 %	12.4 %	14.6 %
<i>(misinfo.)</i>	+6.16	+21.3	-2.69	+5.71	-2.30
Reactions (N)	78.4 %	77.4 %	69.3 %	67.0 %	72.1 %
<i>(misinfo.)</i>	-5.75	-12.8	+14.4	+2.39	-1.06

Table 2: Interaction types: Percentage of total engagement with non-misinformation (N) pages per political leaning, and in alternating rows the misinformation difference in percentage points (misinformation delta relative to non-misinformation pages). Comments, shares and reactions add up to 100 % in each column. Reactions are most common.

Total	Far Left	Left	Center	Right	Far Right
Status (N)	0.46 %	0.34 %	0.21 %	0.36 %	0.64 %
<i>(misinfo.)</i>	-0.08	-0.31	-0.17	-0.00	+2.10
Photo (N)	17.6 %	23.2 %	18.6 %	11.0 %	13.7 %
<i>(misinfo.)</i>	+55.9	+11.4	+16.8	+1.28	+12.3
Link (N)	47.6 %	64.1 %	62.7 %	75.3 %	62.9 %
<i>(misinfo.)</i>	-32.0	-5.50	-13.1	-17.6	-11.6
FB video (N)	33.9 %	6.80 %	13.1 %	7.90 %	20.7 %
<i>(misinfo.)</i>	-25.0	-0.86	-1.20	+13.3	-8.48
Live video (N)	0.38 %	3.45 %	5.24 %	5.37 %	1.87 %
<i>(misinfo.)</i>	+0.99	-2.83	-2.73	-2.63	+5.40
Ext. video (N)	0.12 %	2.07 %	0.20 %	0.10 %	0.19 %
<i>(misinfo.)</i>	+0.24	-1.92	+0.36	+5.66	+0.23

Table 3: Post types: Percentage of total engagement with non-misinformation (N) pages per political leaning, and in alternating rows the misinformation difference in percentage points (misinformation delta relative to non-misinformation pages). Different post types add up to 100 % in each column. Link posts contribute most engagement.

posts contribute 73.5 % of engagement with misinformation sources, as opposed to only 17.6 % for non-misinformation.

The total engagement metric as we have presented it in this section can provide an overview of the Facebook news ecosystem as a whole. However, if the underlying list of news publishers is incomplete and missing large pages from a particular political leaning or (mis)information group, the picture could be biased and lead to incorrect conclusions.

4.2 Publisher/Audience Engagement

A separate question is how well individual publishers generate engagement within their primary audiences. We utilize a per-page metric that sums the interactions of all posts by the page during the study period. The primary audience of pages are their followers, since posts are primarily shown to a page’s followers. To account for pages’ different audience sizes, we divide the per-page sum of interactions by the largest number of followers observed for that page during the study period. This makes it possible to compare small pages with only a niche following to well-established pages

with large audiences. Because the metric sums posts of a page without accounting for the number of posts, pages do not incur a “penalty” for low-performing posts, but they can gain a higher score with posts that perform well relative to the size of each page’s respective follower base.

From this perspective, the median misinformation provider generates 1.46 interactions per follower, or a mean of 3.71, whereas the median non-misinformation provider generates a higher 2.06 interactions per follower, but a lower mean of 3.15. The difference in trend between the median and mean is likely due to the mean being more impacted by outliers (exceptionally well-performing misinformation providers).

The box plot in Figure 3 shows how (mis)information providers perform across the political spectrum. On the Far Left and the Far Right, the median misinformation page engages better with their audience than the respective non-misinformation page of the same political leaning, whereas the opposite is true for the less extreme political leanings. Also worth noting is that the median engagement of the four left-leaning groups is significantly lower than the engagement of the corresponding right-leaning groups, suggesting that Slightly Right and Far Right pages better mobilize their followers, irrespective of (mis)information status.

Because of outliers, averages are considerably higher than the medians. The median Slightly Right misinformation page, for instance, generated 1.2 interactions per follower over the study period, whereas the average was 5.8, which even outperformed the average of Slightly Right non-misinformation (whereas the median did not). Overall, the means indicate that misinformation providers from the Far Left, Slightly Right, and Far Right engage better with their audiences than the corresponding non-misinformation providers. For Slightly Left and Center publishers, those in the misinformation group consistently performed worse than those in the non-misinformation group, both in the median and mean. (Median and mean values can also be found at the bottom of Table 9 in the appendix.)

To determine which of these differences in mean engagement were statistically significant, we need to disentangle the effects of partisanship and factualness. For that purpose, we fit a Multivariate ANOVA model with partisanship and factualness as the independent variables, and their interaction on the natural log-transformed distribution of engagement per follower as the dependent variable (Table 4; see Appendix A.1 for details about the appropriateness of this test). In addition to the significant main effects (that are semantically less interesting for the purpose of our analysis), the interaction of factualness and partisanship is also significant at the .05 level, except for Slightly Left publishers. For publishers of all other partisanship groups, ANOVA showed significance of the effect of factualness on mean engagement per follower (i.e., misinformation associated with decreased engagement for Far Left and Center publishers, and increased engagement for Slightly Right and Far Right publishers). We note that while the effect for the Far Left and Slightly Right was significant, it was on the basis of only 16 and 11 misinformation pages, respectively, thus we do not have a high degree of confidence that this result is representative of these categories of pages in general. Post-hoc testing (Appendix A.2) confirmed the significance of factualness in explaining differences in engagement per follower for Center and Far Right partisanship.

Test (Section)	F	Far Left	Slightly Left	Center	Slightly Right	Far Right
Publisher (4.2)	15	$t(186) = 4.30$ $p < 0.01$	$t(385) = 0.50$ $p = 0.59$	$t(1.53k) = -7.30$ $p < 0.01$	$t(187) = 2.10$ $p = 0.04$	$t(262) = 7.10$ $p < 0.01$
Post (4.3)	5,709	$t(369k) = 112$ $p < 0.01$	$t(1.12M) = 23.8$ $p < 0.01$	$t(4.60M) = 37.4$ $p < 0.01$	$t(497k) = 94.9$ $p < 0.01$	$t(374k) = -15.2$ $p < 0.01$
Video views (4.4)	893	$t(32.5k) = 55.9$ $p < 0.01$	$t(57.8k) = -11.0$ $p < 0.01$	$t(440k) = 41.6$ $p < 0.01$	$t(37.9k) = 19.7$ $p < 0.01$	$t(29.2k) = 30.1$ $p < 0.01$
Video engagement (4.4)	144	$t(32.5k) = 22.3$ $p < 0.01$	$t(57.8k) = -6.50$ $p < 0.01$	$t(440k) = 68.0$ $p < 0.01$	$t(37.9k) = 3.00$ $p < 0.01$	$t(29.2k) = 4.60$ $p < 0.01$

Table 4: Multivariate ANOVA test results for the interaction of partisanship and factualness in four metrics (using natural log transform). Except for Slightly Left in the per-page, per-follower metric, the interaction was significant at the 0.05 level.




Figure 3: Box plot of engagement with (mis)information pages from Far Left to Far Right, normalized by each page’s number of followers. White lines represent the medians and + the means. Some outliers up to 82.8 not shown. Mean engagement per follower is higher for misinformation pages than non-misinformation pages in the Far Left, Slightly Right, and Far Right.

The audience engagement distributions for misinformation and non-misinformation publishers overlap. That is, even though one group performs better than another on average or in the median, many individual pages from the group do perform worse. There are Far Right non-misinformation pages that engage their audiences better than the median misinformation page, for instance, and some center misinformation pages do better than the median center non-misinformation page, even though most do not.

In addition to understanding how well publishers of (mis)information engage with their primary audiences, it is also interesting to look at how large these audiences are. To that end, Figure 4 shows the distribution of followers per page. For the Far Right, misinformation and non-misinformation pages have a similar median of around 200 k followers, whereas for all other political leanings, misinformation pages have a considerably higher median number of followers than the respective non-misinformation counterpart, such as 1.1 M vs. 248 k on the Far Left, or 956 k vs. 128 k for Slightly Right pages. These larger per-page audience sizes for misinformation pages outside of the Far Right come from comparatively few




Figure 4: Box plot of followers per page, from Far Left to Far Right. White lines represent the medians and + the means. Some outliers up to 114M not shown. For each partisanship, there are much fewer misinformation pages than non-misinformation pages, but their median number of followers tends to be higher.

misinformation pages. Indeed, in several political leanings, a vast majority of misinformation pages have larger audiences than most of the non-misinformation pages with the same partisanship. This suggests a potential audience for new misinformation pages to enter the market, whereas the Far Right appears to be more saturated.

The audience-normalized per-page metric we have used in this section has a number of drawbacks. First, using a single value for each page can cause occasional outlier posts to propagate to the final distribution, that is, a page’s viral post could dominate its score, even if it happened only once and is not representative of the entire study period. Second, normalizing by the audience size can cause strong effects at the extremes, such as reducing the contribution of high-engagement posts of a page with an even larger follower base, or amplifying the contribution of moderate-engagement posts of a page with a very small follower base. Some of these effects were mitigated by removing pages with fewer than 100 followers and aggregating across all posts of each page, but Figure 5 shows that there still are cases of this kind in the distribution of normalized per-page engagement.

A third limitation of our per-publisher metric is that it sums post engagement without accounting for the quantity of posts. This

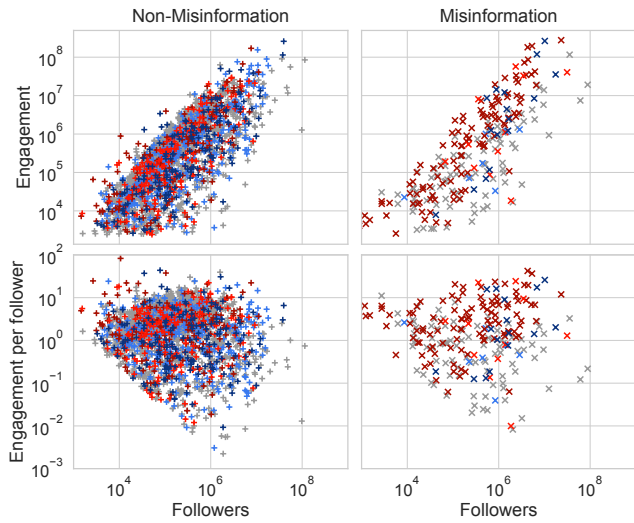


Figure 5: Scatter plot of non-misinformation publishers (left) and misinformation publishers (right), contrasting follower counts with total interactions (above) and with follower-normalized interactions (below). Double log scale. Identical range on the x axes vertically, and on the y axes horizontally. Normalizing by followers can penalize publishers with very high engagement but an even larger follower base, and amplify publishers with moderate engagement from a small follower base.

is done under the theory that the attention of a page’s audience is limited, and that the metric should neither prescribe whether a page can obtain that engagement with a single post or split it across multiple lower-engagement posts, nor should the metric penalize pages for no-engagement posts. However, when pages can generate repeat engagement by publishing more frequently (short of saturating their audience’s limited attention), differences in posting behavior can make it hard to compare the audience engagement of different pages. Figure 6 shows that Far Left and Far Right misinformation pages indeed post more frequently than their non-misinformation counterparts, which could explain why they accumulate higher per-page engagement. Similarly, Slightly Left and Center misinformation pages post less frequently than non-misinformation pages, which could hurt their performance under our per-page audience engagement metric. In other words, higher audience engagement of Far Left and Far Right misinformation publishers, and lower audience engagement of Slightly Left and Center misinformation pages speaks to the effectiveness of publishers, but it does not imply that Far Left and Far Right audiences would be more likely than Slightly Left or Center audiences to engage with a post of misinformation. To explore this aspect, we consider per-post engagement in the next section.

4.3 Post Engagement

To quantify the engagement advantage (or disadvantage) of individual pieces of content from (mis)information pages, we now study posts independently from their pages. This metric is better suited to




Figure 6: Box plot of posts per page, from Far Left to Far Right. White lines represent the medians and + the means. Some outliers up to 62 k not shown. Misinformation pages on the Far Left, Slightly Right and Far Right tend to post more than the respective non-misinformation pages.

study viral content with extremely high engagement because such posts appear as individual data points in the distribution, and their influence on the overall distribution can be reduced or assessed by considering median post engagement or comparing it to the mean. This metric also models low-engagement posts that “disappear” in the summed engagement of the per-page audience engagement or the ecosystem-wide total engagement metrics. As this metric is deliberately independent from individual publishers, a large share of low-engagement posts, for instance, does not necessarily imply that certain news publishers do not engage well with their audiences overall; it only means that many individual posts fail to find an audience (320 k posts, or roughly 4.3 % have no engagement).

We note that all content annotation for factualness is done at the source level. While there is a high likelihood that most posts from a news source follow its general style, tone, and attitude towards factualness and political leaning, there is no guarantee that this holds true in each single instance. This limitation may introduce a level of noise into our results.

We do not normalize post engagement because we have no reliable estimate for the audience size. Most posts are likely shown to a subset of the corresponding page’s followers, but once content becomes viral, Facebook may also show it to non-followers in the “Suggested For You” box [9]; in this case, per-follower normalization would introduce a bias. Additionally, follower-normalized per-post engagement becomes even more unstable than the normalized per-page audience engagement metric from Figure 5 because there is no aggregation at all, thus we avoid normalization here.

At the level of individual posts, those from misinformation sources tend to receive more engagement than those from non-misinformation sources. Of our data set of over 7.5 M posts, 446 k are from misinformation sources. Those posts have a mean engagement of 4,670 interactions, compared to a mean of 765 interactions for non-misinformation sources. Even when breaking down posts by the political leaning of their source, as shown in the box plot in Figure 7, those from a misinformation page have a consistent




Figure 7: Box plot of engagement per post, with partisanship and factualness derived from the publisher (log scale). White lines represent the medians and + the means. Outliers up to 4M not shown. Across the political spectrum, median posts from misinformation pages consistently outperform those of the corresponding non-misinformation pages.

advantage in engagement, both in the median and the mean (except for Slightly Left, where the mean engagement with posts from misinformation sources is slightly lower). This advantage can be very large in absolute terms. In the Far Left, the misinformation median is over an order of magnitude larger than the non-misinformation median and is outperforming even the non-misinformation mean. The misinformation to non-misinformation difference between the medians is similarly large for Slightly Right. The median misinformation advantage is lowest for the Far Right. (Median and mean values can also be found at the bottom of Table 5.)

Using a Multivariate ANOVA with partisanship and factualness as the independent variables and their interaction on the natural log-transformed distribution of per-post engagement as the dependent variable, we found that the impact of factualness on mean engagement was significant across all political leanings at the 0.05 level (Table 4). Post-hoc testing confirmed the significance of factualness in explaining differences in mean engagement for all partisanship.

When looking at the types of interactions in Table 5, reactions are by far the most common type of engagement. In the median, posts from misinformation sources have higher engagement across all interaction types. As the most extreme example, a median post from Far Left non-misinformation pages garners 101 reactions, whereas a median post from a Far Left misinformation page generates 1.5k additional reactions. This may be due in part to different follower bases, as shown in Figure 4. But even for the Far Right, where follower bases are more similar for misinformation and non-misinformation pages, a median post from misinformation pages generates 77% more reactions than a median post from non-misinformation sources, and has almost twice the numbers of comments and shares. In the means, absolute increases in engagement are even higher, but not always in relative terms.

Table 6 shows median and mean engagement numbers for various types of posts. (Table 11 in the appendix further breaks this down into interaction types.) Photo posts and various types of video posts tend to have the highest increases in engagement when

comparing non-misinformation to misinformation sources. Live video posts from misinformation pages perform exceptionally well on the Far Right, with 6.47k additional interactions in the median, or 18.5k in the mean, when compared to a post from a Far Right non-misinformation source.

4.4 Video

When studying misinformation news on Facebook, ideally we would like to assess how often such posts are shown to users. Post impression data, however, is not available on CrowdTangle. Video views are the closest available metric, but different in various regards. Views are not impressions because they require a minimum amount of activity (i.e., watching for at least three seconds). Furthermore, views are available only for video posts, which account for 5.1% to 19.4% of all posts per political leaning, and video posting patterns differ across publishers. We collected a separate data set of video views as described in Section 3.3.1. It covers a subset of the video posts from the remainder of the paper and was collected at a different time, thus we only compare trends.

We begin by looking at total views of videos from (mis)information sources across the Facebook news ecosystem, computed similar to total engagement in Section 4.1. At a high level, view trends are in line with all-post-type engagement. Views of videos from non-misinformation sources significantly outnumber views from misinformation sources across the political spectrum from the Far Left to Slightly Right. Videos from Far Right misinformation sources, however, collect 3.4 times more views in total than those from non-misinformation sources. As Figure 8 shows, the proportions across different political leanings do not match total engagement with all post types (Figure 2); for example, despite their larger number, videos from Slightly Left non-misinformation sources generate only around 54% of views from Far Left non-misinformation sources. (Engagement with all types of content from Slightly Left non-misinformation sources was 1.6 times the total engagement from Far Left non-misinformation sources.) We also note that Slightly Left misinformation providers posted only 337 videos over the course of our study period, which means that the data does not permit us to assess any advantage or disadvantage of misinformation among Slightly Left video.

Figure 9a shows the distribution of views for individual videos across the political and (mis)information spectrum according to their publisher, similar to post engagement in Section 4.3. On a per-video basis, the median number of views from misinformation sources exceeds those from non-misinformation sources for all political leanings (except for slightly left, which may not be significant due to the low number of videos). The means confirm higher views from Center, Slightly Right and Far Right misinformation sources. For the Far Left, the mean indicates fewer views of videos from misinformation sources (whereas the median indicated more). Fitting a Multivariate ANOVA model with partisanship and factualness as the independent variables and the natural log-transformed distribution of views per video as the dependent variable, we found the interaction effect of factualness was significant at the 0.05 level for all political leanings (Table 4), and post-hoc testing confirmed the significance of factualness in explaining differences in mean video views for all partisanship groups.

Median	Far Left	Left	Center	Right	Far Right
Comments (N)	13.0	6.00	5.00	6.00	36.0
<i>(misinfo.)</i>	+351	+3.00	+0.00	+62.0	+37.0
Shares (N)	12.0	4.00	4.00	4.00	28.0
<i>(misinfo.)</i>	+191	+42.0	+5.00	+31.0	+26.0
Reactions (N)	101	37.0	30.0	35.0	212
<i>(misinfo.)</i>	+1.50 k	+100	+65.0	+225	+165
Overall (N)	142	53.0	48.0	53.0	310
<i>(misinfo.)</i>	+2.22 k	+156	+86.0	+362	+251

(a) Median

Mean	Far Left	Left	Center	Right	Far Right
Comments (N)	212	149	91.0	154	388
<i>(misinfo.)</i>	+918	-106	+4.30	+335	+625
Shares (N)	255	90.0	61.9	93.0	425
<i>(misinfo.)</i>	+1.91 k	+139	+79.2	+618	+321
Reactions (N)	1.70 k	818	345	501	2.10 k
<i>(misinfo.)</i>	+7.07 k	-322	+867	+2.22 k	+2.21 k
Overall (N)	2.16 k	1.06 k	498	748	2.91 k
<i>(misinfo.)</i>	+9.90 k	-289	+950	+3.17 k	+3.16 k

(b) Mean

Table 5: Median (a) and mean (b) interactions per post from non-misinformation (N) pages, broken down by interaction type, and in alternating rows the difference for posts from misinformation pages (delta relative to non-misinformation). Values do not add up to the overall aggregate because computations are done independently. A median post from Far Left misinformation pages receives 1.5 k more reactions than the non-misinformation counterpart. Due to outliers, the means can be much higher.

Median	Far Left	Left	Center	Right	Far Right
Status (N)	127	50.0	43.0	48.0	289
<i>(misinfo.)</i>	+728	+67.0	+66.0	+280	+115
Photo (N)	379	299	82.0	47.0	611
<i>(misinfo.)</i>	+21.0 k	+374	+316	+2.07 k	+1.15 k
Link (N)	611	57.0	43.0	17.0	26.0
<i>(misinfo.)</i>	+2.21 k	-29.0	+13.0	+137	+1.38 k
FB video (N)	146	133	45.0	114	1.10 k
<i>(misinfo.)</i>	+2.41 k	+227	+321	+2.75 k	+1.63 k
Live video (N)	183	662	205	285	116
<i>(misinfo.)</i>	+1.11 k	-373	+412	+142	+6.47 k
Ext. video (N)	25.0	23.0	65.0	86.0	54.0
<i>(misinfo.)</i>	+2.55 k	+50.0	-65.0	+902	+199
Overall (N)	142	53.0	48.0	53.0	310
<i>(misinfo.)</i>	+2.22 k	+156	+86.0	+362	+251

(a) Median

Mean	Far Left	Left	Center	Right	Far Right
Status (N)	1.26 k	786	374	661	2.26 k
<i>(misinfo.)</i>	+2.39 k	-109	+801	+2.21 k	+1.72 k
Photo (N)	4.01 k	5.55 k	1.43 k	1.19 k	4.60 k
<i>(misinfo.)</i>	+27.8 k	-4.49 k	+1.23 k	+7.14 k	+9.76 k
Link (N)	1.81 k	2.62 k	404	925	1.57 k
<i>(misinfo.)</i>	+3.95 k	-2.51 k	-213	+3.93 k	+23.0 k
FB video (N)	10.8 k	1.88 k	1.11 k	1.27 k	9.24 k
<i>(misinfo.)</i>	-2.47 k	-1.24 k	+1.57 k	+10.4 k	+1.55 k
Live video (N)	895	2.78 k	707	1.50 k	2.96 k
<i>(misinfo.)</i>	+1.61 k	-1.24 k	+967	+718	+18.5 k
Ext. video (N)	461	539	381	375	650
<i>(misinfo.)</i>	+10.3 k	-403	-306	+6.46 k	+1.47 k
Overall (N)	2.16 k	1.06 k	498	748	2.91 k
<i>(misinfo.)</i>	+9.90 k	-289	+950	+3.17 k	+3.16 k

(b) Mean

Table 6: Median (a) and mean (b) interactions per post of each type from non-misinformation (N) pages, and in alternating rows the difference for posts from misinformation pages (delta relative to non-misinformation). Values do not add up to the overall aggregate because computations are done independently. Photo, Facebook video and live video posts from misinformation pages receive significantly higher median and mean engagement per post than posts from non-misinformation pages.

For comparison, we show engagement with the same set of videos in Figure 9b. Disregarding slightly left video, the trends of the medians and means with regard to the impact of a misinformation source on video engagement are the same as they are for video views. To understand differences in mean engagement per video, we used a Multivariate ANOVA model similar to what we have described previously, and found the difference in mean engagement per video due to factualness was statistically significant at the 0.05 level of confidence for all political leanings of the video publisher (Table 4). Post-hoc testing confirmed the significance of factualness in explaining differences in mean video engagement for all partisanship groups.

To further investigate the similarity of video views and engagement, we plot video views against engagement in Figure 9c. This figure suggests that engagement is generally correlated with views, and that to some extent, engagement-based metrics may serve as a substitute for view-based metrics. However, 283 videos received

more engagement than views, out of which 246 received more reactions than views. Facebook users likely reacted to the videos without watching them, as users can react only once to each post (whereas they could comment or share multiple times). While we could filter out these pathological cases, we cannot account for similar effects of engaging without viewing among videos with more views than engagement, thus views are not a suitable replacement for the lack of impression data. While we would like to study how many users who were shown a misinformation video also engaged with the video, using view data might not be a good approximation.




Figure 8: Bar plot of total views of videos from (mis)information news publishers from Far Left to Far Right, with the number of videos posted in each group on the x axis. In absolute terms, only videos from Far Right misinformation pages accumulated more views than their non-misinformation counterparts. (Different data set than Figure 2, does not compare directly.)

4.5 Summary of Findings

In the ecosystem of U.S. news sources on Facebook, only in the Far Right does misinformation accumulate more overall engagement than non-misinformation. Across all other political leanings, misinformation providers generate less engagement than non-misinformation news sources, but sometimes still a sizeable share of total engagement (e.g., 37.7% in the Far Left). When looking at the performance of individual news providers and how well they engage with their respective follower base, the advantage of misinformation widens to the Far Left when considering median performance, and additionally includes Slightly Right publishers when considering mean audience engagement. In terms of engagement on a per-post basis, median engagement with posts from misinformation providers is consistently higher across the entire political spectrum, and for average engagement, only Slightly Left misinformation sources are at a disadvantage compared to their non-misinformation counterparts. Similar results hold for views of videos posted by news pages. We also find that while there are fewer misinformation than non-misinformation pages, individual misinformation pages tend to have larger average follower bases.

5 DISCUSSION

Given our findings, there is no single answer to the question of how bad is Facebook’s misinformation problem. While we found misinformation news to be highly engaging to audiences across the entire political spectrum, in absolute terms it is driving a significant portion of engagement only on the extreme ends of the spectrum, and particularly on the far right. Our methodology does not give us any insight into *why* the apparent potential for more misinformation is not exploited to the same degree outside of the far right; this is an interesting topic for future research.

Recommendations. Facebook data available to researchers on CrowdTangle are currently very limited in scope. While we were

able to analyze engagement and views, the data are entirely opaque about content recommendations or impressions, and we cannot disentangle the effects of Facebook’s algorithms, content attractiveness, and user behavior. For example, we were able to show that misinformation content is more engaged with, but in order to study whether it is truly more engaging, the *rate* of engagement, we would need impression data. Furthermore, in order to study the sources of engagement, we would need the impression data broken down by categories such as followers, non-followers, shares, and sponsored impressions. These data could help researchers understand whether there is a difference in rate of engagement with misinformation based on partisanship or the type of content promotion. We believe that understanding these factors is crucial to devising effective countermeasures against misinformation.

Limitations. The results of our analysis are as imperfect and limited as the data sources that we rely on. Major threats to the validity of our results are potential biases or omissions in the lists of news publishers that we obtained from NewsGuard and Media Bias/Fact Check. We were unable to discover the Facebook pages of some of these publishers, and our data sets are limited to what is available on CrowdTangle. When publishers delete posts from their Facebook pages, for instance, they also disappear from CrowdTangle, which means that highly controversial posts may be missing in our data set. Our findings reflect the limits of the data available on CrowdTangle; we were able to extract high-level trends in engagement with misinformation news sources, but cannot derive any more detailed data-driven explanations of these phenomena. Lastly, we label news sources, not individual news pieces, based on their reputations for factualness and partisanship as judged by NewsGuard and Media Bias/Fact Check. Our boolean misinformation attribute is based on a threshold instead of modeling the degree to which news publishers vary in how much misinformation, fake news, and conspiracy theories they spread. We do not know which percentage of content promoted by misinformation sources contains misinformation, or how much content from a publisher is in line with its partisanship assessment.

6 RELATED WORK

There is a large body of work on misinformation. We focus on studies of misinformation news in online social media.

Misinformation and the 2016 election. Allcott and Gentzkow [1] explore the economics of misinformation on Facebook during the 2016 election and measure the relative exposure of the U.S. population to fake news. Faris et al. [11] perform a large-scale analysis of content from media sources shared on Facebook via CrowdTangle, and perform a network analysis. They observe that the far-right media ecosystem is distinct from the mainstream media in linking behavior, as well as in journalistic standards and practices. Three of those authors, Faris, Roberts, and Benkler [4] expand this work to additionally perform case studies of left and right partisan disinformation. They note that while individual Facebook posts spreading both left and right disinformation generate high levels of engagement, left-wing partisan media sources are much less likely than their right-wing counterparts to repeat disinformation from other sources. Guess, Nagler, and Tucker [15] study user sharing on Facebook by observing user behavior directly, and matching against a

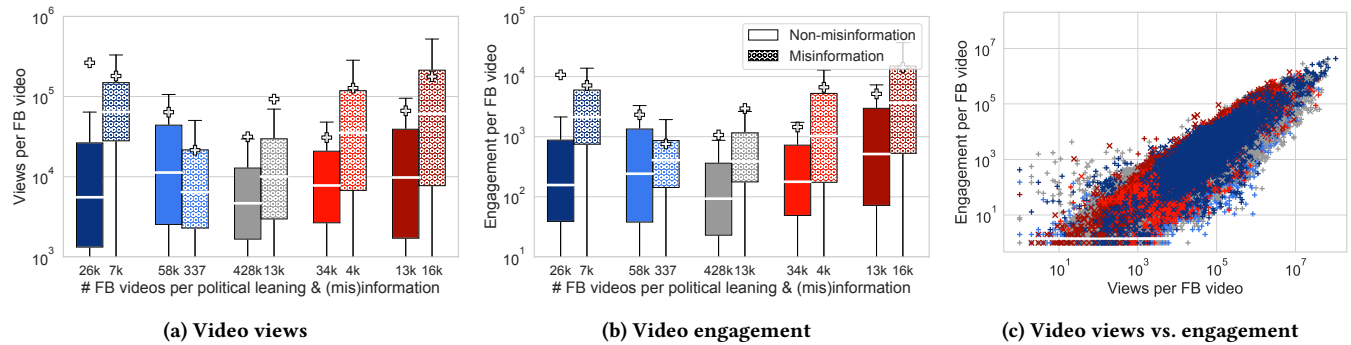


Figure 9: Views and engagement with Facebook-native and live video (separate data set). Medians and means for views and engagement follow similar trends except for content from Slightly Left misinformation pages (they posted only 337 videos, thus likely not reliable). The scatter plot excludes 171 videos with 0 views and an additional 5,511 videos with 0 engagement due to the double log scale. While views and engagement appear largely correlated, outliers suggest users engaging with video posts without viewing.

list of known purveyors of fake news. They find that fake news makes up only a small part of overall user sharing activity. **User engagement with misinformation.** Van Bavel et al. [32] explore psychological motivations and propose a model for understanding the reasons users might have for engaging with misinformation. They point out that users share content even when they do not believe it is true, to signal partisan identity. Several others [18, 21] have also found that users engage with and share news content they do not believe to be true. Geeng, Yee, and Roesner [12] use a browser extension to inject fake news into study participants' Facebook timelines. They observe how participants interact with misinformation when they encounter it, and document their explanations for (not) investigating.

Misinformation ecosystem and audit studies. Closest to our work is a large-scale study of news link sharing in reddit communities by Weld, Glenski, and Althoff [34], who use news source ratings from Media Bias/Fact Check to annotate links with a partisan bias and factualness evaluation. We study a different platform, Facebook, and various different post engagement metrics on the official Facebook pages of news publishers instead of link sharing in unofficial communities. Kornbluh, Goldstein, and Weiner [19] of the German Marshal Fund use news evaluations from NewsGuard to identify trustworthy, fake news, and manipulator news sources, and measure their engagement over time; our work considers additional publisher/audience and per-post engagement metrics, and contrasts publishers based on their partisanship. Glenski et al. [13] study how users share content from deceptive news sources on Twitter, classifying a relatively small number of 282 news providers and monitoring how users share content from those providers. Our work covers 2,551 news sources active on Facebook and focuses on the publisher side instead of user characteristics and sharing behavior. Hussein et al. [17] audit YouTube video recommendations and note factors associated with YouTube's recommendation algorithm promoting misinformation.

Journalistic Uses of CrowdTangle. Journalists have used CrowdTangle data for reporting on Facebook's ecosystem, particularly during election periods. In 2019, the Guardian created a dashboard [26]

showing the election videos with most engagement during the parliamentary election that year. Kevin Roose's Twitter feed *Facebook's Top 10* [28] posts daily the Facebook posts with most engagement over the past 24 hours.

7 CONCLUSION

In this paper, we studied how much Facebook users engage with (mis)information from U.S. news providers on Facebook. We found that in absolute terms, Far-Right misinformation sources accumulate more engagement than non-misinformation sources of the same partisanship (68.1% of overall Far Right engagement, followed by 37.7% on the Far Left). Even though the misinformation ecosystem outside of the Far Right is smaller in absolute engagement terms, misinformation also performs better than non-misinformation on the Far Left and Slightly Right when considering how publishers engage with their respective follower base. In terms of engagement with posts from misinformation sources, the median is higher for all political leanings. Our findings illustrate that the lure of misinformation is not confined to Far-Right audiences, although there are fewer publishers outside of the Far Right to feed the misinformation ecosystem. We hope that future research will be able to investigate *why* misinformation generates more engagement.

ACKNOWLEDGMENTS

We wish to thank the employees of CrowdTangle and Facebook who built the tools to enable our analysis. Particular thanks are owed to Naomi Shiffman for her help working with CrowdTangle and her insightful feedback on drafts of this work. Cybersecurity for Democracy at NYU's Center for Cybersecurity has been supported by Democracy Fund, Luminare, Media Democracy Fund, the National Science Foundation under grant 1814816, Reset, and Well-spring. This research was supported in part by the French National Research Agency (ANR) through the ANR-17-CE23-0014 and the MIAI@Grenoble Alpes ANR-19-P3IA-0003 grants and the European Union's Horizon 2020 research and innovation programme under grant agreement No 101021377. This material is based upon work supported by the Google Cloud Research Credits program.

REFERENCES

- [1] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31, 2 (2017), 211–236.
- [2] Jennifer Allen, Baird Howland, Markus Mobius, David Rothschild, and Duncan J Watts. 2020. Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances* 6, 14 (April 2020).
- [3] Michael Barthel, Amy Mitchell, and Jesse Holcomb. 2016. Many Americans believe fake news is sowing confusion. <https://www.journalism.org/2016/12/15/many-americans-believe-fake-news-is-sowing-confusion/>
- [4] Yochai Benkler, Robert Faris, and Hal Roberts. 2018. *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. Oxford University Press, New York, 472 pages.
- [5] Lia Bozarth, Aparajita Saraf, and Ceren Budak. 2020. Higher ground? How groundtruth labeling impacts our understanding of fake news about the 2016 U.S. presidential nominees. In *ICWSM*.
- [6] Media Bias/Fact Check. 2021. Media Bias/Fact Check. <https://mediabiasfactcheck.com/>
- [7] Media Bias/Fact Check. 2021. Media Bias/Fact Check methodology. <https://mediabiasfactcheck.com/methodology/>
- [8] CrowdTangle. 2021. Introducing Facebook and Instagram video views! <https://www.crowdtangle.com/resources/videoviews>
- [9] Facebook. 2021. *Why do I see suggested posts in my Facebook News Feed?* Facebook. <https://www.facebook.com/help/485502912850153>
- [10] FactCheck. 2021. Our process. <https://www.factcheck.org/our-process/>
- [11] Robert M. Faris, Hal Roberts, Bruce Etling, Nikki Bourassa, Ethan Zuckerman, and Yochai Benkler. 2017. *Partisanship, propaganda, and disinformation: Online media and the 2016 U.S. presidential election*. Research Paper. Berkman Klein Center for Internet & Society. <http://nrs.harvard.edu/urn-3:HUL.InstRepos:33759251>
- [12] Christine Geeng, Savanna Yee, and Franziska Roesner. 2020. Fake news on Facebook and Twitter: Investigating how people (don't) investigate. In *CHI*.
- [13] Maria Glenski, Tim Weninger, and Svitlana Volkova. 2018. Propagation from deceptive news sources: Who shares, how much, how evenly, and how quickly? *IEEE Transactions on Computational Social Systems* 5, 4 (Dec. 2018), 1071–1082.
- [14] Ted Van Green. 2020. Few Americans are confident in tech companies to prevent misuse of their platforms in the 2020 election. <https://www.pewresearch.org/fact-tank/2020/09/09/few-americans-are-confident-in-tech-companies-to-prevent-misuse-of-their-platforms-in-the-2020-election/>
- [15] Andrew Guess, Jonathan Nagler, and Joshua A Tucker. 2019. Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances* 5, 1 (Jan. 2019).
- [16] Benjamin D Horne, Jeppe Nørregaard, and Sibel Adali. 2019. Different spirals of sameness: A study of content sharing in mainstream and alternative media. In *ICWSM*. 257–266.
- [17] Eslem Hussein, Prerna Juneja, and Tanushree Mitra. 2020. Measuring misinformation in video search platforms: An audit study on YouTube. In *ACM CSCW*.
- [18] Shan Jiang, Miriam Metzger, Andrew Flanagin, and Christo Wilson. 2020. Modeling and measuring expressed (dis)belief in (mis)information. In *ICWSM*.
- [19] Karen Kornbluh, Adrienne Goldstein, and Eli Weiner. 2020. *New study by Digital New Deal finds engagement with deceptive outlets higher on Facebook today than run-up to 2016 election*. Technical Report. The German Marshall Fund of the United States. <https://www.gmfus.org/blog/2020/10/12/new-study-digital-new-deal-finds-engagement-deceptive-outlets-higher-facebook-today>
- [20] Merriam-Webster. 2021. Misinformation. <https://www.merriam-webster.com/dictionary/misinformation>.
- [21] Miriam J Metzger, Andrew Flanagin, Paul Mena, Shan Jiang, and Christo Wilson. 2021. From dark to light: The many shades of sharing misinformation online. *Media and Communication* 9, 1 (2021).
- [22] Adam Mosseri. 2021. Taking action against misinformation across our apps. <https://www.facebook.com/formedia/blog/working-to-stop-misinformation-and-false-news>
- [23] Adam Mosseri. 2021. Working to stop misinformation and false news. <https://www.facebook.com/combatting-misinfo>
- [24] NewsGuard. 2021. NewsGuard. <https://www.newsguardtech.com/>
- [25] NewsGuard. 2021. NewsGuard rating process and criteria. <https://www.newsguardtech.com/ratings/rating-process-criteria/>
- [26] David Pegg, Niamh McIntyre, and Pamela Duncan. 2019. Which election videos are getting the most attention online? <https://www.theguardian.com/politics/2019/nov/14/which-election-videos-are-getting-the-most-attention-online>
- [27] Facebook Journalism Project. 2020. How our fact-checking program works. <https://www.facebook.com/journalismproject/programs/third-party-fact-checking/how-it-works>
- [28] Kevin Roose. 2020. Facebook's Top 10. <https://twitter.com/FacebooksTop10>
- [29] Mattia Samory, Vartan Kesiz Abnoui, and Tanushree Mitra. 2020. Characterizing the social media news sphere through user co-sharing practices. In *ICWSM*.
- [30] Mark Scott. 2021. Thousands of posts around January 6 riots go missing from Facebook transparency tool. <https://www.politico.eu/article/facebook-crowdtangle-data-january-6-capitol-hill-riots-transparency/>
- [31] Naomi Shiffman. 2021. CrowdTangle codebook. <https://help.crowdtangle.com/en/articles/3213537-crowdtangle-codebook>
- [32] Jay J Van Bavel, Elizabeth A Harris, Philip Parnamets, Steve Rathje, Kimberly C Doell, and Joshua A Tucker. 2021. Political psychology in the digital (mis)information age: A model of news belief and sharing. *Social Issues and Policy Review* 15, 1 (Jan. 2021), 84–113.
- [33] Clare Wardle and Hossein Derakhshan. 2021. Information disorder: Toward an interdisciplinary framework for research and policy making. <https://rm.coe.int/information-disorder-toward-an-interdisciplinary-framework-for-research/168076277c>.
- [34] Galen Weld, Maria Glenski, and Tim Althoff. 2021. Political bias and factualness in news sharing across more than 100,000 online communities. *arXiv* (Feb. 2021).

A APPENDIX

A.1 Appropriateness of ANOVA the Model

The overall goal of our analysis was to understand whether factualness (i.e., misinformation or non-misinformation) was associated with significant differences in engagement within the five partisanship groups of news publishers. We first tested more generally whether engagement was distributed differently between the five partisanship groups and the binary factualness variable using the Kolmogorov-Smirnov (KS) test. This test addresses that question directly though comparison of the empirical CDF (rather than a measure of centrality alone), and does not rely on the specific distributional assumptions or strict sample size requirements that many exact tests require. Using the KS test, we made pairwise comparisons of the ten partisanship/factualness combinations and adjusted our p-value threshold to account for multiple comparisons. The KS test results indicated that the distributions of the ten groups differ.

We further explored these differences by fitting a Multivariate ANOVA model, given that our data satisfied the general assumptions of that model. The ANOVA model has the advantage of being interpretable with regard to our research question and allowing us to reason about the interaction of factualness and partisanship.

A.2 Post-Hoc Testing

In cases where the ANOVA F-statistic was significant, we could conclude that there were in fact significant differences in mean engagement, and performed post-hoc comparisons to identify which levels of the test were significant. We applied the Tukey HSD test in a pairwise manner across our combination groups of partisanship and factualness, and adjusted our p-values using Bonferroni correction. Table 7 shows as an example the results for the per-page, per-follower metric from Section 4.2. In that case, the results of the post-hoc testing confirm the finding of significance of factualness for explaining differences in engagement per follower in the Center and Far Right partisanship groups from our ANOVA model. For the three other metrics (per post, video views, and video engagement), similar post-hoc testing with the Tukey HSD test confirmed significance for all political leanings, likely due to the larger sample sizes of the misinformation groups.

Pairwise comparison between groups		Meandiff	p-adj	Lower	Upper	Reject
Center (N)	Center (M)	-1.14	0.01	-1.60	-0.69	True
Center (N)	Far Left (N)	-0.58	0.01	-0.93	-0.24	True
Center (N)	Far Left (M)	0.04	0.90	-1.03	1.11	False
Center (N)	Far Right (N)	-0.14	0.90	-0.51	0.22	False
Center (N)	Far Right (M)	0.42	0.10	0.00	0.85	True
Center (N)	Slightly Left (N)	-0.38	0.01	-0.63	-0.13	True
Center (N)	Slightly Left (M)	-1.21	0.46	-2.82	0.40	False
Center (N)	Slightly Right (N)	-0.21	0.72	-0.55	0.13	False
Center (N)	Slightly Right (M)	-0.34	0.90	-1.63	0.94	False
Center (M)	Far Left (N)	0.56	0.08	0.01	1.11	True
Center (M)	Far Left (M)	1.18	0.08	0.03	2.34	True
Center (M)	Far Right (N)	1.00	0.01	0.44	1.56	True
Center (M)	Far Right (M)	1.57	0.01	0.97	2.17	True
Center (M)	Slightly Left (N)	0.76	0.01	0.27	1.25	True
Center (M)	Slightly Left (M)	-0.07	0.90	-1.74	1.60	False
Center (M)	Slightly Right (N)	0.94	0.01	0.39	1.48	True
Center (M)	Slightly Right (M)	0.80	0.76	-0.56	2.15	False
Far Left (N)	Far Left (M)	0.62	0.81	-0.49	1.74	False
Far Left (N)	Far Right (N)	0.44	0.17	-0.04	0.91	False
Far Left (N)	Far Right (M)	1.01	0.01	0.49	1.53	True
Far Left (N)	Slightly Left (N)	0.20	0.89	-0.19	0.59	False
Far Left (N)	Slightly Left (M)	-0.63	0.90	-2.27	1.01	False
Far Left (N)	Slightly Right (N)	0.37	0.33	-0.08	0.83	False
Far Left (N)	Slightly Right (M)	0.24	0.90	-1.08	1.56	False
Far Left (M)	Far Right (N)	-0.19	0.90	-1.30	0.93	False
Far Left (M)	Far Right (M)	0.38	0.90	-0.76	1.52	False
Far Left (M)	Slightly Left (N)	-0.42	0.90	-1.51	0.66	False
Far Left (M)	Slightly Left (M)	-1.26	0.64	-3.18	0.67	False
Far Left (M)	Slightly Right (N)	-0.25	0.90	-1.36	0.86	False
Far Left (M)	Slightly Right (M)	-0.39	0.90	-2.05	1.28	False
Far Right (N)	Far Right (M)	0.57	0.05	0.04	1.10	True
Far Right (N)	Slightly Left (N)	-0.24	0.77	-0.64	0.17	False
Far Right (N)	Slightly Left (M)	-1.07	0.65	-2.71	0.57	False
Far Right (N)	Slightly Right (N)	-0.06	0.90	-0.53	0.41	False
Far Right (N)	Slightly Right (M)	-0.20	0.90	-1.53	1.13	False
Far Right (M)	Slightly Left (N)	-0.80	0.01	-1.27	-0.34	True
Far Right (M)	Slightly Left (M)	-1.64	0.11	-3.30	0.02	False
Far Right (M)	Slightly Right (N)	-0.63	0.01	-1.15	-0.11	True
Far Right (M)	Slightly Right (M)	-0.77	0.79	-2.11	0.58	False
Slightly Left (N)	Slightly Left (M)	-0.83	0.89	-2.46	0.79	False
Slightly Left (N)	Slightly Right (N)	0.17	0.90	-0.21	0.56	False
Slightly Left (N)	Slightly Right (M)	0.04	0.90	-1.26	1.34	False
Slightly Left (M)	Slightly Right (N)	1.01	0.71	-0.63	2.64	False
Slightly Left (M)	Slightly Right (M)	0.87	0.90	-1.19	2.93	False
Slightly Right (N)	Slightly Right (M)	-0.14	0.90	-1.46	1.18	False

Table 7: Result of the Tukey HSD post-hoc test for the ANOVA finding from Section 4.2 (log per-page, per-follower engagement). Factualness: (M) misinformation, (N) non-misinformation.

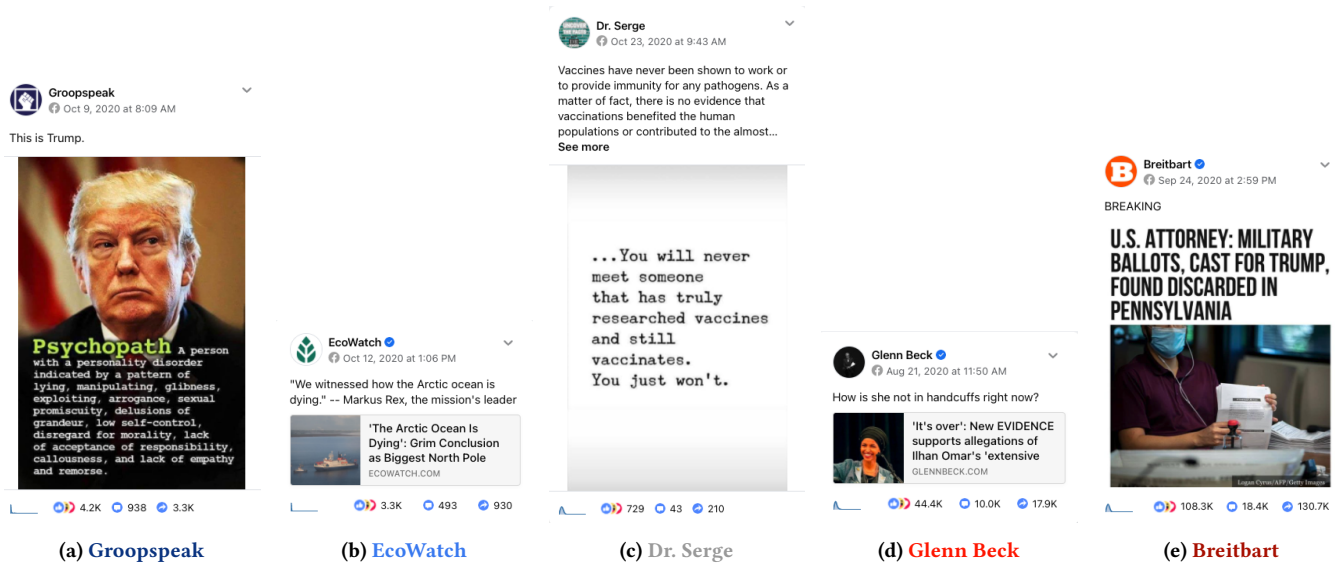


Figure 10: Misinformation: Sample Facebook posts ordered by the publisher’s political leaning from far left (a) to far right (e). Publishers classified as misinformation based on NewsGuard and Media Bias/Fact Check data.

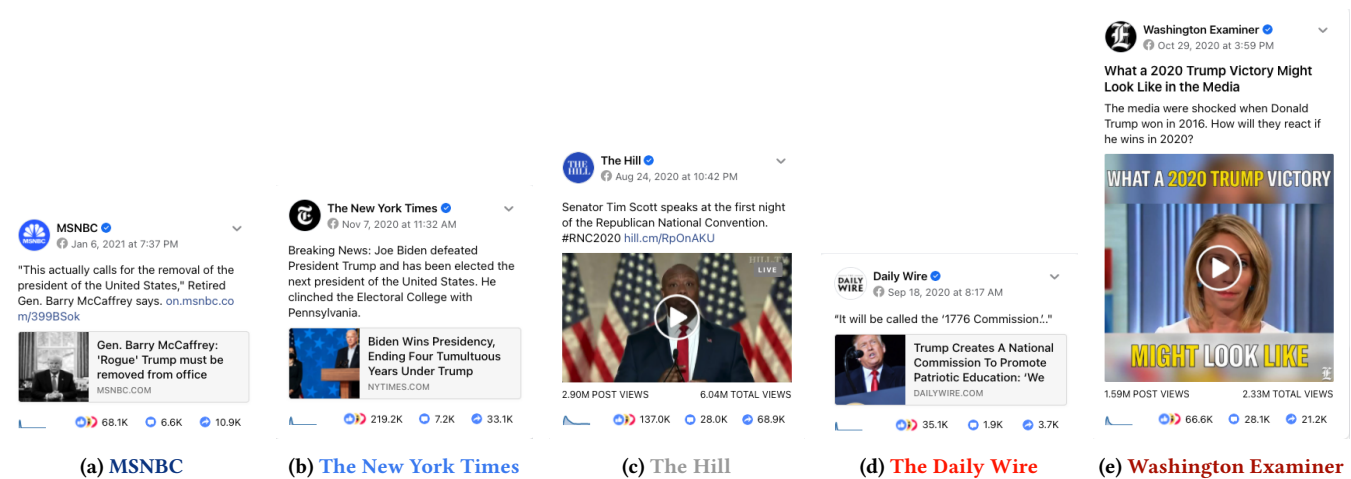


Figure 11: Non-Misinformation: Sample Facebook posts ordered by the publisher’s political leaning, far left (a) to far right (e). Publishers classified as non-misinformation based on NewsGuard and Media Bias/Fact Check data.

Partisanship	#	Non-Misinformation	Misinformation
Far Left	1	The Dodo	Occupy Democrats
	2	CNN	The Other 98%
	3	Washington Press	NowThis
	4	Rappler	Trump Sucks
	5	MSNBC	Bipartisan Report
Left	1	Bleacher Report Football	Dr. Josh Axe
	2	ABC News	True Activist
	3	Rudaw	EcoWatch
	4	NBC News	Mint Press News
	5	The New York Times	National Vaccine Information Center
Center	1	World Health Organization (WHO)	Jesus Daily
	2	CGTN	China Xinhua News
	3	The Hill	RT
	4	BBC News	The Epoch Times
	5	ESPN	Higher Perspective
Right	1	Fox Business	David J Harris Jr.
	2	Daily Wire	NTD Television
	3	Forbes	Glenn Beck
	4	IJR	Todd Starnes
	5	The Babylon Bee	Sputnik
Far Right	1	Ben Shapiro	Fox News
	2	Trending World by The Epoch Times	Breitbart
	3	The White House	Dan Bongino
	4	PragerU	Donald Trump For President
	5	Blue Lives Matter	NewsMax

Table 8: Top 5 pages with the highest total engagement within each partisanship and factualness group.




Figure 12: Composition of the filtered data set by political leaning (horizontal, colored), and origin publisher list (vertical, hatched). In the middle and bottom rows of each plot, news publisher pages are weighted by total interactions and followers, respectively. Same data as Figure 1, but plotted separately for (a) only non-misinformation pages, and (b) only misinformation pages. In relative terms, overlap between the two lists is smallest in the center and larger in the extremes. Overlap increases when pages are weighted by interactions or followers.

The biggest difference between non-misinformation and misinformation pages lies in the distribution of pages, interactions, and followers across the political spectrum, as visible on the x axes. Based on our two source lists NG and MB/FC, slightly left misinformation pages, for instance, account for a disproportionately low 7 out of 386 slightly left pages overall, whereas the 109 far-right misinformation pages are disproportionately numerous compared to the 154 far-right non-misinformation pages. (Absolute numbers of pages per political leaning and (mis)information status can be found on the x axis of Figure 2.) We discuss these effects in Section 4.

Median	Far Left	Left	Center	Right	Far Right
Comments (N)	0.11	0.24	0.49	0.41	0.24
<i>(misinfo.)</i>	+0.19	-0.22	-0.44	-0.25	+0.08
Shares (N)	0.15	0.19	0.34	0.26	0.28
<i>(misinfo.)</i>	+0.11	+0.00	-0.18	+0.00	+0.33
Reactions (N)	0.63	1.04	1.49	1.22	1.36
<i>(misinfo.)</i>	+0.57	-0.80	-1.03	-0.33	+0.60
angry (N)	0.07	0.08	0.09	0.10	0.16
<i>(misinfo.)</i>	+0.07	-0.05	-0.08	-0.07	+0.10
care (N)	0.01	0.01	0.02	0.01	0.01
<i>(misinfo.)</i>	+0.01	-0.01	-0.02	+0.00	+0.00
haha (N)	0.03	0.06	0.09	0.08	0.06
<i>(misinfo.)</i>	+0.08	-0.05	-0.08	-0.03	+0.08
like (N)	0.38	0.63	0.86	0.73	0.76
<i>(misinfo.)</i>	+0.33	-0.42	-0.53	-0.14	+0.44
love (N)	0.05	0.09	0.14	0.08	0.06
<i>(misinfo.)</i>	+0.04	-0.07	-0.11	+0.05	+0.07
sad (N)	0.03	0.07	0.14	0.06	0.03
<i>(misinfo.)</i>	+0.02	-0.05	-0.13	-0.04	+0.01
wow (N)	0.01	0.03	0.06	0.05	0.03
<i>(misinfo.)</i>	+0.01	-0.02	-0.05	-0.02	+0.02
Overall (N)	0.99	1.50	2.44	2.00	2.00
<i>(misinfo.)</i>	+0.67	-1.04	-1.67	-0.71	+1.12

Median	Far Left	Left	Center	Right	Far Right
Comments (N)	0.34	0.43	0.77	0.63	0.50
<i>(misinfo.)</i>	+0.37	-0.36	-0.62	+0.17	+0.26
Shares (N)	0.45	0.31	0.46	0.42	0.76
<i>(misinfo.)</i>	+0.54	-0.06	-0.15	+0.58	+0.20
Reactions (N)	1.94	1.75	2.05	1.96	2.88
<i>(misinfo.)</i>	+2.38	-1.13	-1.23	+2.10	+0.81
angry (N)	0.27	0.16	0.15	0.20	0.51
<i>(misinfo.)</i>	+0.18	-0.08	-0.10	+0.69	+0.01
care (N)	0.02	0.02	0.04	0.03	0.02
<i>(misinfo.)</i>	+0.00	-0.02	-0.03	+0.00	+0.01
haha (N)	0.22	0.11	0.16	0.24	0.24
<i>(misinfo.)</i>	+0.49	-0.10	-0.11	+0.08	+0.13
like (N)	1.11	1.09	1.15	1.12	1.74
<i>(misinfo.)</i>	+1.50	-0.68	-0.58	+0.97	+0.53
love (N)	0.20	0.17	0.24	0.17	0.19
<i>(misinfo.)</i>	+0.15	-0.12	-0.16	+0.23	+0.14
sad (N)	0.07	0.13	0.21	0.14	0.10
<i>(misinfo.)</i>	+0.05	-0.09	-0.18	+0.02	-0.01
wow (N)	0.05	0.06	0.09	0.07	0.08
<i>(misinfo.)</i>	+0.02	-0.03	-0.06	+0.12	+0.01
Overall (N)	2.73	2.48	3.29	3.02	4.14
<i>(misinfo.)</i>	+3.30	-1.55	-2.00	+2.85	+1.27

(a) Median

(b) Mean

Table 9: Median (a) and mean (b) engagement per non-misinformation (N) page, normalized by each page’s number of followers, and in alternating rows the difference for misinformation pages (delta relative to non-misinformation pages). Broken down by interaction and reaction type. Values do not add up to the overall aggregate because computations are done independently.

Median	Far Left	Left	Center	Right	Far Right
Status (N)	0.00	0.00	0.00	0.00	0.00
<i>(misinfo.)</i>	+0.00	-0.00	-0.00	-0.00	+0.00
Photo (N)	0.02	0.02	0.09	0.04	0.07
<i>(misinfo.)</i>	+0.45	+0.04	+0.08	+0.12	+0.17
Link (N)	0.68	1.27	1.98	1.53	1.06
<i>(misinfo.)</i>	-0.14	-1.02	-1.81	-1.09	+0.37
FB video (N)	0.00	0.01	0.02	0.01	0.01
<i>(misinfo.)</i>	+0.01	-0.00	-0.02	+0.07	+0.01
Live video (N)	0.00	0.00	0.00	0.00	0.00
<i>(misinfo.)</i>	+0.00	-0.00	-0.00	+0.04	+0.00
Ext. video (N)	0.00	0.00	0.00	0.00	0.00
<i>(misinfo.)</i>	+0.00	+0.00	+0.00	+0.00	-0.00
Overall (N)	0.99	1.50	2.44	2.00	2.00
<i>(misinfo.)</i>	+0.67	-1.04	-1.67	-0.71	+1.12

(a) Median

Mean	Far Left	Left	Center	Right	Far Right
Status (N)	0.19	0.01	0.01	0.01	0.04
<i>(misinfo.)</i>	-0.14	-0.01	-0.01	+0.06	+0.08
Photo (N)	0.64	0.31	0.30	0.30	0.68
<i>(misinfo.)</i>	+3.00	-0.16	+0.06	+0.01	+0.86
Link (N)	1.70	2.00	2.55	2.42	3.14
<i>(misinfo.)</i>	+0.36	-1.27	-1.89	+2.08	-0.14
FB video (N)	0.17	0.11	0.25	0.19	0.23
<i>(misinfo.)</i>	+0.05	-0.08	-0.14	+0.34	+0.27
Live video (N)	0.01	0.05	0.17	0.09	0.02
<i>(misinfo.)</i>	+0.02	-0.04	-0.05	+0.06	+0.18
Ext. video (N)	0.01	0.01	0.01	0.00	0.02
<i>(misinfo.)</i>	+0.01	+0.00	+0.03	+0.30	+0.02
Overall (N)	2.73	2.48	3.29	3.02	4.14
<i>(misinfo.)</i>	+3.30	-1.55	-2.00	+2.85	+1.27

(b) Mean

Table 10: Median (a) and mean (b) engagement per non-misinformation (N) page, normalized by each page’s number of followers, and in alternating rows the difference for misinformation pages (misinformation delta relative to non-misinformation pages). Broken down by post type. Values do not add up to the overall aggregate because computations are done independently.

Status	Far Left	Left	Center	Right	Far Right
Comments (N)	12.0	5.00	5.00	6.00	36.0
(<i>misinfo.</i>)	+116	+3.00	-1.00	+48.5	+18.0
Shares (N)	11.0	4.00	4.00	4.00	25.0
(<i>misinfo.</i>)	+47.0	+22.0	+5.00	+21.0	+13.0
Reactions (N)	91.0	34.0	27.0	32.0	198
(<i>misinfo.</i>)	+486	+45.0	+50.0	+172	+75.0
Photo	Far Left	Left	Center	Right	Far Right
Comments (N)	23.0	25.0	8.00	3.00	33.0
(<i>misinfo.</i>)	+1.26 k	-13.0	+2.00	+114	+124
Shares (N)	43.0	16.0	8.00	3.00	86.0
(<i>misinfo.</i>)	+3.47 k	+153	+10.0	+254	+122
Reactions (N)	278	206	47.0	33.0	421
(<i>misinfo.</i>)	+15.0 k	+276	+29.1	+1.59 k	+748
Link	Far Left	Left	Center	Right	Far Right
Comments (N)	96.0	9.00	8.00	1.00	2.00
(<i>misinfo.</i>)	+661	-7.50	+9.00	+42.0	+177
Shares (N)	32.0	2.00	2.00	1.00	3.00
(<i>misinfo.</i>)	+54.0	+4.50	+1.00	+13.0	+95.5
Reactions (N)	412	25.0	17.0	11.0	20.0
(<i>misinfo.</i>)	+1.48 k	-4.00	+2.00	+78.0	+1.00 k
FB video	Far Left	Left	Center	Right	Far Right
Comments (N)	14.0	12.0	3.00	11.0	87.0
(<i>misinfo.</i>)	+359	+3.00	+22.0	+128	+223
Shares (N)	22.0	13.0	6.00	15.0	196
(<i>misinfo.</i>)	+368	+107	+33.0	+490	+195
Reactions (N)	101	91.0	31.0	78.0	733
(<i>misinfo.</i>)	+1.55 k	+138	+248	+2.07 k	+1.07 k
Live video	Far Left	Left	Center	Right	Far Right
Comments (N)	27.0	136	83.0	106	17.0
(<i>misinfo.</i>)	+398	-72.0	+87.0	-22.0	+2.18 k
Shares (N)	19.0	18.0	12.0	22.0	13.0
(<i>misinfo.</i>)	+68.0	+55.0	+38.0	+28.0	+464
Reactions (N)	122	315	85.0	122	74.0
(<i>misinfo.</i>)	+595	-142	+238	+158	+3.32 k
Ext. video	Far Left	Left	Center	Right	Far Right
Comments (N)	4.00	1.00	9.00	10.0	4.00
(<i>misinfo.</i>)	+453	+2.00	-9.00	+148	+34.0
Shares (N)	2.00	1.00	5.00	6.00	5.00
(<i>misinfo.</i>)	+265	+17.5	-5.00	+94.0	+15.0
Reactions (N)	18.0	18.0	39.0	56.0	38.0
(<i>misinfo.</i>)	+1.87 k	+30.5	-39.0	+585	+145

(a) Median

Status	Far Left	Left	Center	Right	Far Right
Comments (N)	183	129	76.4	142	355
(<i>misinfo.</i>)	+417	-83.6	+5.92	+294	+365
Shares (N)	108	71.1	43.5	73.1	282
(<i>misinfo.</i>)	+317	+139	+21.4	+350	+121
Reactions (N)	970	585	254	446	1.63 k
(<i>misinfo.</i>)	+1.66 k	-164	+774	+1.57 k	+1.23 k
Photo	Far Left	Left	Center	Right	Far Right
Comments (N)	251	335	157	157	415
(<i>misinfo.</i>)	+1.97 k	-299	-29.6	+245	+1.27 k
Shares (N)	743	342	203	207	809
(<i>misinfo.</i>)	+5.24 k	-60.8	+125	+1.34 k	+1.07 k
Reactions (N)	3.02 k	4.87 k	1.07 k	823	3.38 k
(<i>misinfo.</i>)	+20.6 k	-4.13 k	+1.13 k	+5.55 k	+7.42 k
Link	Far Left	Left	Center	Right	Far Right
Comments (N)	439	313	182	158	155
(<i>misinfo.</i>)	+929	-308	-138	+446	+1.65 k
Shares (N)	154	253	40.4	104	253
(<i>misinfo.</i>)	+295	-229	-19.5	+528	+1.54 k
Reactions (N)	1.21 k	2.05 k	181	664	1.16 k
(<i>misinfo.</i>)	+2.72 k	-1.97 k	-55.7	+2.95 k	+19.8 k
FB video	Far Left	Left	Center	Right	Far Right
Comments (N)	495	164	72.0	129	685
(<i>misinfo.</i>)	+471	-136	+71.1	+694	+581
Shares (N)	1.39 k	250	167	271	1.77 k
(<i>misinfo.</i>)	+597	+0.09	+325	+3.08 k	+723
Reactions (N)	8.90 k	1.47 k	872	871	6.78 k
(<i>misinfo.</i>)	-3.54 k	-1.10 k	+1.17 k	+6.65 k	+247
Live video	Far Left	Left	Center	Right	Far Right
Comments (N)	194	850	328	580	955
(<i>misinfo.</i>)	+730	-357	+159	+117	+6.98 k
Shares (N)	79.9	147	68.9	126	363
(<i>misinfo.</i>)	+130	+123	+200	+199	+1.40 k
Reactions (N)	621	1.78 k	311	797	1.64 k
(<i>misinfo.</i>)	+747	-1.01 k	+608	+402	+10.1 k
Ext. video	Far Left	Left	Center	Right	Far Right
Comments (N)	65.9	118	123	124	84.2
(<i>misinfo.</i>)	+1.44 k	-110	-118	+749	+157
Shares (N)	47.7	42.2	28.7	34.0	85.2
(<i>misinfo.</i>)	+1.65 k	-4.83	-20.1	+966	+167
Reactions (N)	348	378	229	218	481
(<i>misinfo.</i>)	+7.25 k	-288	-168	+4.74 k	+1.14 k

(b) Mean

Table 11: Median (a) and mean (b) interactions per post from non-misinformation (N) pages, separately for each type of post and broken down by interaction type, and in alternating rows the difference for posts from misinformation pages (misinformation delta relative to non-misinformation pages). Values do not add up to the overall aggregate because computations are done independently. This is a version of Table 5 computed separately for each type of post.