

Interactions with Potential Mis/Disinformation URLs Among U.S. Users on Facebook, 2017-2019

Aydan Bailey*

Theo Gregersen*

Franziska Roesner

Paul G. Allen School of Computer Science & Engineering, University of Washington

Note (September 10, 2021): After the publication of this paper, we learned that the dataset provided by Facebook contained an error: interactions from the 50% of U.S. users who had no political page affinity (PPA) score were excluded (see <https://www.nytimes.com/live/2020/2020-election-misinformation-distortions#facebook-sent-flawed-data-to-misinformation-researchers>). When the data is updated, we plan to revisit our analyses. As written, the paper reports only on the 50% of U.S. users who had PPA scores and who thus interacted with Facebook in certain ways that may impact our findings.

ABSTRACT

Misinformation and disinformation online — and on social media in particular — have become a topic of widespread concern. Recently, Facebook and Social Science One released a large, unique, privacy-preserving dataset to researchers that contains data on URLs shared on Facebook in 2017-2019, including how users interacted with posts and demographic data from those users. We conduct an exploratory analysis of this data through the lens of mis/disinformation, finding that posts containing potential and known mis/disinformation URLs drew substantial user engagement. We also find that older and more politically conservative U.S. users were more likely to be exposed to (and ultimately re-share) potential mis/disinformation, but that those users who were exposed were roughly equally likely to click regardless of demographics. We discuss the implications of our findings for platform interventions and further study towards understanding and reducing the spread of mis/disinformation on social media.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in collaborative and social computing**; • **Information systems** → **Social networking sites**; • **Security and privacy** → *Social aspects of security and privacy*;

KEYWORDS

misinformation, disinformation, social media

ACM Reference Format:

Aydan Bailey, Theo Gregersen, and Franziska Roesner. 2021. Interactions with Potential Mis/Disinformation URLs Among U.S. Users on Facebook, 2017-2019. In *ACM SIGCOMM 2021 Workshop on Free and Open Communications on the Internet (FOCI'21)*, August 27, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3473604.3474561>

1 INTRODUCTION

The last several years have brought significant concerns about the spread and influence of (unintentionally false) misinformation and (intentionally false) disinformation online [25, 43, 49]. Sometimes referred to as “fake news” or information pollution, mis/disinformation is notorious for spreading through social media platforms like Facebook and Twitter, where large (and increasing) fractions of Americans consume at least some of their news [39].

While social media platforms analyze their own user data internally, it is also important to study mis/disinformation in a public, neutral, academic manner. Such research enables us to compare findings across platforms, contributes to scientific knowledge about human behaviors related to online mis/disinformation more generally, and ultimately informs efforts to reduce its spread and impact. Prior efforts have often focused on Twitter because the openness of the platform lends itself to external research, though prior work has also studied public data on other platforms (see Section 2). For more closed platforms like Facebook, academic work often relies on smaller-scale case studies and/or controlled user studies rather than large-scale organic platform usage data.

In this work, we leverage a released privacy-protected dataset from Facebook [29] to conduct an exploratory study of the spread of potential mis/disinformation and associated user behaviors on the platform. This dataset, made available to researchers through an application process organized by Facebook and Social Science One [40], contains billions of rows of Facebook data from January 2017 to December 2019 about which demographic groups of users saw and interacted with which URLs on the platform. Though this dataset has some privacy-related limitations, it provides us with a unique view into a slice of Facebook user behaviors and the spread of mis/disinformation and other URLs shared on Facebook.

*Co-first authors listed in alphabetical order.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

FOCI'21, August 27, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8640-1/21/08.

<https://doi.org/10.1145/3473604.3474561>

Specifically, we investigate the following research questions, focused on U.S. Facebook users:

- (1) **RQ1: Exposure to and interactions with potential mis/disinformation on Facebook:** How, and how much, did U.S. Facebook users in our dataset see and interact with known or potential mis/disinformation URLs?
- (2) **RQ2: Demographic differences:** What, if any, differences exist in how different demographic groups interacted with U.S.-shared mis/disinformation URLs in our dataset – including people of different genders, age groups, and U.S. political inclinations?

To investigate these questions, we must overcome several challenges. First, we must grapple with the question of “what is mis/disinformation?” in the first place. In the absence of ground truth, we use existing identifications of known or potential mis/disinformation URLs: Facebook’s own third-party fact checking labels and a list of low-quality, potential mis/disinformation domains previously compiled by other researchers [51]. We must also consider several limitations of the dataset, including its privacy-preserving noisy properties and limited scope, as well as more fundamental limitations about the ability of platform data (where behaviors are influenced by a platform’s design) to shed light on human behavior in general [50]. We thus focus on an exploratory analysis that describes what we find in the data; we do not aim to make generalizable claims about people in general.

Among our results, we find that in 2017-2019, known and potential mis/disinformation drew substantial user engagement on Facebook, and that large fractions of users re-shared posts containing (any) URLs without first clicking on them. We also find that older adults with more conservative U.S. political leanings were exposed to, clicked on, and re-shared potential mis/disinformation more than other groups. However, the largest differences are in who was exposed to (i.e., shown) these posts in the first place, which may reflect only partly on explicit user behaviors and instead result from Facebook’s newsfeed algorithm.

Based on our findings, we reflect on the implications for platform-based mis/disinformation defenses (e.g., the importance of preventing initial exposure) as well as future research (e.g., towards defenses that reduce susceptibility and spread after exposure). Overall, given this unique dataset, our results provide a view into U.S. users’ exposure to and interactions with potential mis/disinformation on Facebook in 2017-2019 that could not previously be studied at this scale.

2 BACKGROUND AND RELATED WORK

Fully systematizing the broad and growing research space around online mis/disinformation is beyond our scope, but we point to several summaries [25, 43, 49] and highlight some of the most related examples. We present a more detailed comparison of our findings to prior work in Section 5.

Ecosystem and Platform Studies. Methodologically, our work is most closely related to prior studies of user behaviors based on social media platform data directly. There are several important limitations of this methodology, including limitations of platform APIs [43]. Notably, Facebook’s API is limited and has become more so over time [15, 34], though Twitter’s API also has limitations [46].

Moreover, human behaviors on specific platforms are shaped by the designs of those platforms, and thus findings from these contexts should not be over-generalized [50].

Nevertheless, studying data on platforms themselves gives important insights into how mis/disinformation spreads in practice. Due to its relatively open API, many of these studies have focused on Twitter [1, 6, 7, 18, 38, 44, 48], though significant work has also studied user behaviors on Facebook [12, 13, 16, 20, 27, 45] as well as on platforms like Reddit, WhatsApp, YouTube, and others [4, 14, 21, 23, 28, 35, 36].

Studies of Facebook tend to predate the early 2018 API restrictions and/or focus on specific group or topic case studies, e.g., public Facebook groups on known conspiracy theories or Facebook pages of known news outlets [30, 47]. Our work provides insights from a broader view of Facebook data that was made available through Facebook’s partnership with Social Science One [29]. Research on this dataset is beginning to appear [19]; our investigation complements this recent work, confirming some of its findings with a different view on mis/disinformation as well as providing new results.

User Studies. Other work has studied people directly, conducting survey, controlled lab, or in-depth qualitative studies to test and/or observe interactions with mis/disinformation on social media platforms. As above, these studies have focused primarily on Facebook and/or Twitter [12, 13, 16, 17, 20, 26, 27, 45], or surveyed people about their general social media usage [39]. Findings about the effectiveness of interventions on social media platforms (e.g., “fake news” labels) have been mixed [5, 13, 17, 24, 33].

3 METHODS

3.1 Dataset Overview

The Facebook URL dataset that we analyze was made available to us via an application process in a partnership between the Social Science Research Council, Social Science One, and Facebook [41]. As per our Research Data Agreement with Facebook, we accessed this dataset only via Facebook-provided servers and did not download a copy of the dataset. Between the original public description of the dataset [29] and our analysis in May 2021, Facebook updated the dataset (at least) to cover a longer time period (through March 2020 at the time of our analysis). Since we did not download or save a previous version of the dataset, we cannot verify exactly what changes have been made over time. The results in this paper are based on the version of the dataset that was available to us on Facebook’s servers as of early May 2021, which we describe here.

The version of the dataset we analyzed in May 2021 consisted of 41.7 million URLs from 1.1 million unique parent domains. This includes only URLs that were shared or posted publicly on Facebook at least 100 times (with Laplace(5) noise). The included URLs were posted between January 2011 and March 2020.¹ The dataset also includes 17 trillion datapoints of user interactions with those URLs (e.g., clicks, re-shares) from January 2017 through December 2019, associated with some noisy (privacy-preserving) demographic information. Data about repeated actions (e.g., subsequent clicks

¹Though the bulk of the URL data, and all of the user interaction data, is from the 2017-2019 time range, some URLs shared during that period were first posted as early as January 2011, as reflected by the `first_post_time` attribute in the dataset.

from users who clicked on a URL more than once) or users who later deleted their Facebook accounts is not included.

Our Scope and Baseline. We only consider URL data during the time period for which we have corresponding user interaction data, January 2017 through December 2019. We also consider only interactions from U.S. users (for whom there is additional demographic information, discussed below), and we scope the URL list to those URLs that are labeled as most-shared in the U.S. (as opposed to in other countries). Our final subset thus consists of 9,172,097 URLs from 216,137 unique parent domains (and associated user interaction data). We refer to this set of URLs as our **baseline “U.S. URLs” subset**.

We use this subset of all U.S. URLs as a baseline against which we compare user interactions with known and potential mis/disinformation. We use this full subset rather than a more specific baseline (e.g., mainstream or trustworthy news sources) because this allows us to consider the “average” U.S. URL on Facebook without imposing any value judgements on what is trustworthy or otherwise worth including in a baseline. This full set of U.S. URLs likely includes a wide variety of types of content (which future work may wish to separate out further), but we note that because the dataset includes only URLs that were shared publicly 100 or more times (with Laplace(5) noise), a potentially long tail of unpopular websites is excluded by definition.

Demographic Data. The dataset includes demographic information associated with URL interactions: age bracket, gender, and “political page affinity.” The latter refers to U.S. political leaning and is available only for U.S. users. It is calculated by analyzing the pages that users follow and based on methods from Barberá et al. [3]. The dataset documentation does not detail further how PPA is calculated here, but the end result maps users into five integer buckets from -2 (ideological left, i.e., liberal) to 2 (ideological right, i.e., conservative). For simplicity and to avoid suggesting conclusions where we draw none, we omit demographic categories with substantially less data (i.e., “other” for gender and “NA” for age) from our figures and analysis.

Differential Privacy. In order to release this dataset while protecting the privacy of the users whose data is represented, Facebook applied differential privacy techniques to add noise to the data [9, 10, 29]. Specifically, the dataset has been made differentially private on the level of user interactions, meaning that the dataset would look statistically the same whether or not any particular user interaction is actually included. Theoretically, this gives plausible deniability about any particular action by any particular user, while still accurately reflecting behavior statistically at this large scale.

In practice, differential privacy here is accomplished by adding Gaussian noise to certain columns in the dataset. In our results, we thus display calculated (noisy) values as well as 95% confidence intervals. Signal-to-noise properties apply: though statistics for larger samples include more noise, that noise will generally be smaller relative to the true value.

Because noise was applied to the dataset only for data columns related to user activity and demographics, URL attributes such as post times and fact-check ratings are precise.

3.2 Known and Potential Mis/Disinformation URL Subsets

To investigate our research questions, we need a list of mis/disinformation domains and/or URLs. However, compiling such a list is a major research challenge of its own. Since addressing that challenge is not a goal of this work, we rely on two existing characterizations, one that is likely too narrow and one that is likely too broad:

(1) Third-Party Fact-Checked False List (“TPFC False U.S.”).

Our most straightforward source of “false” URLs are those flagged directly by Facebook, in collaboration with its third-party fact checking partners. This process focuses primarily on “viral” misinformation (i.e., popular posts) as well as clear falsehoods. Posts of flagged URLs are labeled in Facebook’s UIs and/or downranked in its newsfeed algorithm [11]. Specifically, we focus on U.S. URLs marked as “fact checked as false” or “fact checked as mixture or false headline” (as opposed to, e.g., “satire”). There were 6,838 U.S. URLs with such ratings from 2,644 unique parent domains.

The advantage of this list is our high confidence in the “false” labels. The disadvantages include that inclusion in the list depends on decisions made by Facebook and its partners (e.g., a focus on viral content, and limited coverage due to limited resources). Moreover, Facebook’s fact-checking choices are not entirely transparent and have been criticized for being both potentially too harsh and/or too lenient on both U.S. politically left- and right-leaning content [32, 42]. This list thus does not represent a labeling of a random sample of URLs, and exclusion from this list does not imply truth.

(2) Broad Potential Misinformation List (“Low-Quality News U.S.”).

Since Facebook’s fact-checked URL list is a narrow view of potential mis/disinformation, we also consider a much broader view of low-quality news. Specifically, we consider a subset of potential mis/disinformation domains subsampled from a list from Zeng et al. [51]. This list (of which we obtained an updated version in October 2020) includes 1,344 domains compiled from amateur, open source, and professional content checkers. Since this list includes a broad range of low-quality news domains, we subsampled based on labels only from professional fact-checking sources Snopes, FactCheck.org, and Politifact. The resulting subsample contains 120 domains, 103 of which were present and corresponded to 108,408 U.S. URLs in the Facebook dataset.

The disadvantage of this list is that it may be overly general: not all articles on a domain may contain falsehoods, and thus we can consider these URLs only *potential* mis/disinformation, or more generally, low-quality news. However, this list provides us with a much broader perspective that is not influenced by Facebook’s fact-checking decisions.

3.3 Ethics

Our study was reviewed by our university’s human subjects review board (IRB). Because the dataset does not include identifiable information about individuals, our IRB determined that this study did not classify as human subjects research. Nevertheless, we treated the data with caution and adhered to Facebook’s guidelines about its use: accessing it only on Facebook-provided servers, making no efforts to de-anonymize or identify any individuals, and submitting the paper to Facebook to review for potential privacy issues in

advance of publication. (No changes to the paper were requested by Facebook.)

3.4 Limitations

First, as discussed, we lack ground truth for mis/disinformation and instead consider both a narrow (the TPFC False URL list) and a broad (the low-quality news URL list) view for this exploratory study.

Second, the noisy privacy-preserving properties of the dataset limit us to questions about interaction data in aggregate (not individual users or actions). The fact that the dataset (also for privacy reasons) includes only data about URLs that were shared publicly by enough people on Facebook also means there may be a selection bias towards certain URLs (that tend to be shared publicly and widely) and Facebook users (those who make non-private posts).

Third, there is other data that would provide useful context that we simply do not have. For example, we lack demographic information about who *posted* URLs in the first place; the “share” interactions consist only of *re-shares* of posts. We also cannot compare posts containing URLs (reflected in this dataset) with posts that do not contain URLs.

Moreover, the fact that the dataset is not public and available only through an application process and official agreement — though helping to protect the privacy of the Facebook users whose data is represented — makes it challenging for others to reproduce our results. The fact that we only access the data on Facebook’s servers (per our agreement with Facebook), where the dataset may be updated or our access revoked at any time, also makes it potentially challenging for us (and others) to reproduce our own results in the future. As discussed, the May 2021 version of the dataset we analyzed includes more records than described in the June 2020 public description of the dataset [29].

Finally, we cannot separate the role of Facebook’s newsfeed algorithm and other design features from user behaviors [50]. For example, when we consider URL views, we cannot distinguish whether a user saw a post because they chose to follow a relevant public Facebook page, because their friend posted it, or because Facebook pushed a sponsored post to their feed. We must assume that Facebook’s algorithm attempts to push posts to people that they are likely to engage with. Similarly, we lack details about how U.S. political leaning (PPA) was calculated, so we cannot distinguish whether correlations between PPA and certain behaviors (e.g., which URLs a user engages with) are by definition or emergent.

Still, this large Facebook dataset provides a unique opportunity to study user interactions with potential mis/disinformation on the platform itself, providing a complementary perspective to other methodologies with other limitations. In our results, we focus on *describing* what we find in this dataset about how people on Facebook interacted with URLs in 2017–2019; we do not intend to make generalizable claims about human behaviors in general.

4 FINDINGS

We first characterize our URL subsets, and then study how Facebook users interacted with mis/disinformation and low-quality news URLs, comparing to a baseline of all U.S.-shared URLs and across demographic groups.

Parent Domain	Unique URLs	TPFC-F	Total Views
foxnews.com	67.2K	7	[2.22e10, 2.23e10]
cnn.com	76.1K	0	[2.16e10, 2.17e10]
nytimes.com	71.2K	1	[2.11e10, 2.12e10]
buzzfeed.com	27.8K	0	[1.74e10, 1.75e10]
npr.org	25.4K	0	[1.70e10, 1.70e10]
youtube.com	863.6K	163	[1.50e10, 1.54e10]
washingtonpost.com	60.8K	1	[1.26e10, 1.27e10]
huffingtonpost.com	38.6K	0	[1.25e10, 1.26e10]
dailywire.com	27.1K	20	[1.13e10, 1.14e10]
nbcnews.com	33.1K	0	[1.01e10, 1.02e10]

Table 1: U.S. Subset Domains with Most Views. TPFC-F refers to the number of URLs from that domain that Facebook’s third-party fact-checkers labeled as false. Given the differential privacy noise applied to the data, view counts in Tables 1–3 are presented with 95% confidence intervals; overlapping ranges means that we cannot distinguish the underlying non-private data points. Thus, the ordering in these tables (by view count) is approximate.

Parent Domain	Unique URLs	TPFC-F	Total Views
higherperspectives.com	830	2	[1.34e9, 1.35e9]
dailyasnark.com	3.7K	4	[1.08e9, 1.10e9]
madworldnews.com	4.8K	4	[5.69e8, 5.99e8]
awarenessact.com	2.5K	10	[5.19e8, 5.39e8]
disclose.tv	3.7K	6	[4.65e8, 4.91e8]
thegatewaypundit.com	15.0K	42	[4.36e8, 4.82e8]
thepoliticalinsider.com	5.1K	2	[3.51e8, 3.78e8]
infowars.com	8.9K	25	[3.38e8, 3.77e8]
worldnewsdailyreport.com	430	11	[3.43e8, 3.53e8]
conservativepost.com	2.0K	8	[3.12e8, 3.33e8]

Table 2: Low-Quality News Domains with Most Views.

Parent Domain	Unique URLs	TPFC-F	Total TPFC-F Views
youtube.com	863.6K	163	[5.32e7, 5.70e7]
livebr0adcast.com	156	130	[5.56e5, 2.00e6]
yournewswire.com	2.6K	81	[6.71e7, 7.10e7]
overseasdaily.com	132	78	[3.34e6, 4.96e6]
actual-eventstv.com	105	76	[4.10e5, 1.51e6]
channel23news.com	238	61	[1.60e7, 1.99e7]
wfrv9.com	63	59	[1.04e6, 4.94e6]
dailyusaupdate.com	481	55	[6.13e6, 9.97e6]
actual-events.com	139	53	[7.90e4, 9.83e5]
network-channel.com	138	50	[1.14e3, 6.94e5]

Table 3: Domains with Most TPFC False U.S. URLs.

4.1 Characterizing Our URL Subsets

Tables 1 and 2 show the view counts for the top ten most-viewed domains in our baseline U.S. and our low-quality news URL subsets, respectively. Note that “views” here refers to instances when a user saw a Facebook post containing a URL appearing in their newsfeed, *not* instances where a user clicked through to the target of the URL. The tables also show how many unique URLs from each domain appear in that subset, as well as how many URLs were fact-checked by Facebook’s partners as false (the “TPFC-F” column).

We highlight several observations: First, the most popular U.S.-shared domains in our dataset (Table 1) corresponded largely to well-known mainstream news sites, though we note the outsize popularity of youtube.com. Second, we note that the most popular U.S.-shared domains (Table 1) and the most popular domains from our low-quality news list (Table 2) are disjoint. That is, potential mis/disinformation domains in the form of low-quality news sites were not among the *most* popular shared domains.

Table 3 shows the domains in our dataset with the most URLs that were third-party fact checked as false, and the view counts




Figure 1: Average Cumulative View Counts Over Time for URLs in Different Subsets. 95% confidence intervals are shown.




Figure 2: Percentage of URL Views in Each URL Subset by (a) Gender, (b) Age Group, (c) Political Page Affinity. 95% confidence intervals are shown.

for those URLs. This table reveals the comparatively small scope of Facebook’s third-party fact checking labels, which applied only to a small number of URLs, most of which (except youtube.com) came from domains that appeared with relatively few unique URLs on Facebook.

4.2 Views (Exposure): Who Saw What?

Overall Views / Exposure. We begin with a core question: how much were people exposed to (i.e., shown on their newsfeed) URLs in our different mis/disinformation subsets, compared to the baseline dataset of all U.S. URLs? In exploring this question, we emphasize again that Facebook plays a role in who sees which posts. In other words, trends we surface below may reflect Facebook’s newsfeed algorithm rather than (or in addition to) differences in human behaviors (e.g., which Facebook pages they follow) [50].

To investigate views, we return to Tables 1 and 2, which show the total number of views across all URLs from each domain. We find that URLs from the most popular U.S. domains overall received an order of magnitude more views than URLs from popular domains in the low-quality news subset (though the latter still received large numbers of views). We also note that the total TPFC False URL views for the domains in Table 3 were substantial considering the smaller number of TPFC False URLs: just dozens of TPFC False URLs produced several million views (perhaps by definition, given that viral misinformation is prioritized for Facebook’s fact checking).

Spread Over Time. We also ask: how quickly did URLs in the different subsets spread? In Figure 1, we investigate the cumulative views over time, averaged across URLs, for each of our subsets. Since the dataset provides timestamp information only at month-level granularity, we can track only cumulative monthly views. We find that the average URL (in all subsets) spreads most quickly within the

first month after the initial post, suggesting that mis/disinformation interventions must be deployed quickly to be effective.

From Figure 1 we also see that URLs in the TPFC False subset received significantly more views, on average, than those in other subsets (hence likely making them targets for fact-checking). In addition, TPFC False URLs were longer-lived, receiving a greater fraction of views after the first month compared to the others (nearly doubling in cumulative views). Unfortunately, the granularity of the timestamps in the dataset prevent us from investigating whether spread decreased substantially after the time of fact checking.

Demographic Differences in Exposure. We now turn to our second research question, considering different demographic groups (gender, age, and political leaning). While we do not know the overall demographic breakdown of Facebook users in general, we can compare demographics in interactions with our U.S. URLs baseline to interactions with our low-quality news and TPFC False URL subsets.

Figure 2 shows the percentage of views from users in different demographic categories for the average URL in each URL subset. For example, considering the TPFC False URL subset, we find that on average (i.e., averaged across all URLs in this subset), more than 20% of views came from users ages 65 or older. Taking all U.S. URLs as our baseline, we find that known or potential mis/disinformation URLs in both lists were shown more often (on average) than baseline to male users, to users ages 55 and older, and to users with political page affinity of 1 or more (i.e., right-leaning).

However, the breakdown in Figure 2 obscures potential interactions between demographic categories (for example, older users may also tend to be more conservative). To tease apart these potential interactions, Figure 3 shows heatmaps of how often different demographic groups were shown posts containing these URLs.




Figure 3: Percentage of URL Views by Demographic Groups (Age and Political Page Affinity), for Each URL Subset. Percentages are presented in ranges with 95% confidence intervals due to differential privacy noise.

Each heatmap corresponds to a URL subset, and each square of each heatmap represents a different demographic intersection defined by the tuple {age bucket, political page affinity}. (Since gender did not have an impact on our conclusions, we collapse gender for

readability of the heatmap.) The percentages represent the fraction of views that an average URL-containing post received from each demographic group (i.e., for each heatmap, the percentages sum to 100%).

Figure 3 thus compares how frequently, relative to each other, different *overlapping* demographic groups were shown posts containing URLs. For both of our mis/disinformation URL subsets, we find that U.S. users who are older *and* more politically right-leaning were more likely to see misinformation URLs in their feeds. The top baseline heatmap, by contrast, shows a larger percentage of views from younger and farther left-leaning users — in other words, the trends for low-quality news and fact-checked false URLs were not merely reflections of overall trends of everyone’s Facebook usage and/or experience.

4.3 Clicking: Who Clicked What?

When a Facebook user saw a post containing a URL, what did they do? Figure 4 shows a variety of behaviors aggregated across all users in our dataset, for each URL subset. In this subsection, we begin by considering clicking behaviors. Unlike views, which can result from Facebook pushing content to users, clicks are actions taken directly by users (though still predicated on having seen the post in the first place).

Overall Clicking Behaviors. Figure 4’s set of “clicks” bars shows average click-through rates, i.e., average clicks-per-view. We find that compared to the baseline of all U.S. URLs, which users clicked on roughly 6% of the time they see them, URLs in our mis/disinformation sets were clicked on more often, roughly 8-9% of the time. In other words, the average U.S. user who saw a potential mis/disinformation URL was more likely to click on it than the average user who encountered any random link on Facebook. We note that this comparison depends on our choice of baseline, comparing to the “average” U.S.-shared URL; the reasons for increased engagement with potential mis/disinformation here may result not from mis/disinformation tactics but (for example) because news-style content receives more engagement in general.

Demographic Differences in Click-Through Rate. The heatmaps in Figure 5 break down the click-through rates for different demographic groups and URL subsets. Here we show a *baseline click-through rate heatmap* on top: the percentages (with confidence intervals of 95%) represent the average click-through rate per demographic intersection for all U.S.-shared URLs. The next two heatmaps show *differences in click-through rate* from the baseline: for each demographic intersection, we calculated the ratio of the click-through rate to the baseline click-through rate. (For example, a value of 1 would indicate no change, while a value of 2 would indicate that the click-through rate doubled.)

To our surprise, the demographic trends observed when considering views (above) do not fully hold here. In addition to baseline click-through rate behavior being roughly comparable on both sides of the political spectrum, we see that for the TPF False and low-quality news URLs, differences in click-through rates compared to the baseline were *not* skewed by political page affinity. That is, despite the demographic differences in who saw these posts,




Figure 4: Average Actions per View for URLs in each Subset. 95% confidence intervals are shown.

political leaning did not seem to impact on click-through rates for known or potential mis/disinformation URLs.

We do continue to see differences by age, though not unique to mis/disinformation: Older adults were more likely to click on URLs in general (top heatmap), without additional age-related differences for potential mis/disinformation URLs.

We cannot distinguish *why* people clicked on links or the impacts of reading the underlying content (e.g., perhaps politically left-leaning users click on right-leaning links to debunk them). Still, this finding suggests that political differences in engagement with misinformation reported in prior research may be partially due to who is exposed by the platform in the first place rather than more fundamental differences between groups. Moreover, we emphasize that this exposure is not random: Facebook’s newsfeed algorithm may optimize showing these posts to precisely those users more likely to click on them, and our findings may reflect the success (or at least uniformity) of Facebook’s algorithm in doing just that. Still, this result underscores the importance of limiting initial exposure to mis/disinformation, as we discuss further in Section 5.

4.4 Sharing: Who Shared What?

We now consider URL sharing behaviors. Our dataset includes two types of interactions: *shares*, which involve users re-sharing a Facebook post containing a URL after clicking on it, and *shares without clicks*, where users re-share a post containing a URL without first clicking on the URL. These two metrics are mutually exclusive, meaning that “shares” refers only to shares *with* clicks. We also emphasize that our interaction data (with associated demographic information) includes only *re*-shares, not original posts of a URL.

Overall Sharing Behaviors. Figure 4 shows that mis/disinformation URLs garnered more shares — and more shares without clicks — per view than average U.S.-shared URLs. Share rates are particularly high for TPFC False URLs. Since Facebook prioritizes fact-checking viral mis/disinformation, this finding may reflect in part how the TPFC False subset was selected in the first place. However, note that the sharing rates for the low-quality news URLs approached the rates for these noteworthy TPFC False URLs, i.e., they also drew substantial engagement (per view). We emphasize again that the comparison to baseline depends on our choice of a broad baseline: we see that (these) potential mis/disinformation URLs spread more


efficiently than other U.S. URLs, on average, but we leave to future work more detailed comparisons against other potential URL subsets of interest (e.g., mainstream news).

Across all URL subsets, we were surprised at the high rate of shares without clicks. Considering share counts (rather than shares-per-view), nearly half of all shares, for all URL subsets, was done by Facebook users who had not clicked on the URL they were re-sharing. Specifically, the percentage of shares *without* clicks for the average URL was (with 95% confidence intervals): 42.29–42.35% in the overall U.S.-shared dataset, 39.81–40.18% for TPFC False, and 41.66–42.13% for low-quality news. The fact that these ratios are similar across all URL subsets suggests that there were not necessarily URL or content based differences driving this behavior.


Demographic Differences in Sharing URLs. We consider the rates at which different groups shared the URLs they see. Figure 6 is constructed like the click-through rate heatmaps (with a top baseline, and differences from that baseline below), now considering the average shares-per-view ratio per demographic intersection. Here we consider both shares with and without clicks, i.e., all re-shares. Note that as with clicks, we cannot distinguish the goals of a share (e.g., sharing because one agrees with the article or to debunk it).

Unlike for click-throughs, we again see a trend towards higher sharing rates (given that a user is exposed) of mis/disinformation URLs by politically right-leaning users, particularly of TPFC False URLs. Unlike views, here the impact of political leaning seems to dominate compared to age — younger right-leaning users were also more likely to share these URLs compared to baseline. At the same time, we find that older Facebook users were slightly more likely to share URLs in general, regardless of whether they are mis/disinformation.


Finally, we investigate specifically *what* was shared by Facebook users in different demographic buckets. Figure 7 considers the top 50 most viewed domains from our baseline and low-quality news URL subsets, plotting the average age and political leaning of users who shared URLs from those domains. We have labeled some example domains, including a cluster of potential mis/disinformation URLs shared mostly by older, politically conservative users, as well as examples on the political left. (The TPFC False URL subset is omitted from this comparison because there are not enough flagged URLs per most-viewed TPFC False domain to reduce noise via averaging sufficiently to allow for meaningful comparison.)



(a) The top heatmap shows baseline click-through rates (clicks per view) by demographic group, for all U.S.-shared URLs.




(a) The top heatmap shows baseline sharing rates (shares per view) by demographic group, for all U.S.-shared URLs.



(b) Differences in click-through rates from baseline (e.g., TPFC False URLs were clicked through by adults age 65+ and PPA=2 1.37-1.45x more often than average U.S.-shared URLs).

Figure 5: Baseline and Difference in Click-Through Rates by Demographic Groups and URL Subsets. Ranges representing 95% confidence intervals are shown; overlapping ranges do not allow us to distinguish the underlying non-private values.



(b) Differences in sharing rates from baseline (e.g., average TPFC False URLs were shared by adults age 65+ and PPA=2 2.48-2.61x more often than average U.S.-shared URLs).

Figure 6: Baseline and Difference in Shares/View by Demographic Groups and URL Subsets. Ranges representing 95% confidence intervals are shown; overlapping ranges do not allow us to distinguish the underlying non-private values.




Figure 7: Average Demographics for Users Who Shared URLs from the Top 50 Most Viewed Domains. For both the U.S.-shared URLs baseline and the low-quality news URL subset, we consider the top 50 most-views domains, and calculate the average demographics (age and political page affinity) for users who re-shared posts containing URLs from those domains. Several example domains are labeled and 95% confidence intervals are shown.

5 DISCUSSION

Reflecting on Demographic Differences. Our results show demographic differences in who interacted with potential mis/disinformation on Facebook in 2017-2019. Most strikingly, we find that older adults with right-leaning U.S. political affinities were more exposed to (i.e., viewed more posts containing) potential mis/disinformation URLs. As one caveat, recall that identifying mis/disinformation is challenging — while we relied on reputable sources for our URL lists, we cannot exclude the possibility that these lists (and Facebook’s own fact-checking) may disproportionately include politically conservative content. Future work should continue to consider how to identify mis/disinformation and investigate our research questions with those perspectives. Future work should also consider additional baseline URL subsets (e.g., trustworthy news sources by some definition).

Also noteworthy is that political demographic differences flattened out for click-through rates (though they reappeared for sharing rates). In other words, once a user saw a post containing a mis/disinformation URL, political leaning did not seem to play a role in whether a user clicked on it. This finding may reflect in part Facebook’s success targeting users likely to click, and recall that we do not know *why* people clicked or whether they took the false or low-quality content at face value. Still, this result should caution us not to overestimate the impact of political leaning on how and how often people *engage with* mis/disinformation without considering the likelihood of *being exposed* in the first. The fact that we again see demographic differences for shares may result from many reasons that future work might study: one observation we make is that unlike clicks, shares are publicly visible, which may impact behaviors among different demographic groups.

At the same time, we see that older adults were more likely to be exposed to (i.e., view), click on, and re-share potential mis/disinformation than younger adults. This is the case both because

older adults were disproportionately likely to be shown posts containing these URLs, but also because they then exhibited higher click-through and sharing rates on Facebook in general. Future defenses — either on social media platforms or outside of them (e.g., educational efforts) — should thus look towards reducing susceptibility of older adults in particular.

Comparisons to Prior Findings, and Future Research Questions. Our findings contribute to a broader understanding of how people interact with potential mis/disinformation on social media, based on a unique, large-scale dataset. While smaller-scale and controlled studies help explain how or why people interact with “fake news”, larger-scale measurements can shed light on broader trends. Most related to our work, prior results from larger-scale studies of Facebook and Twitter suggested that false news (under various definitions) is shared more by U.S. politically right-leaning users and older adults, though sharing is generally rare [18, 20]; that Facebook users are more likely to see, click, and share content aligned with their political ideology [2]; that users tend to focus on a small number of news sources [37]; and that low-quality news sources may receive more user engagement than mainstream news [8]. Our findings confirm and update these prior (often pre-2017) results with data from 2017-2019, and emphasize an important nuance in how we should interpret findings about demographic differences: overall trends can be heavily influenced by what content users are exposed to by the platform in the first place.

Our findings confirm other recently published results based on this dataset [19], which used a different list of low-credibility news domains (based on NewsGuard [31]) but also found that older, more conservative adults were more likely to see and share those URLs. Since the underlying dataset is not publicly available, we believe that multiple, corroborating peer-reviewed sets of findings are scientifically valuable. In addition to a complementary perspective considering different mis/disinformation URL subsets, our work

also adds an analysis of clicks (in addition to views and shares) and highlights that views should be interpreted as reflecting users' exposure rather than their direct behaviors.

Our quantitative results also raise questions for follow-on, more qualitative investigations. Most importantly, our results can reveal trends in *what* Facebook users do, but not *why* or *under what conditions* they do it. For instance: Why (and which types of content) do people reshare so frequently without clicking on the associated URLs? Why do people click on URLs in posts and what is the result of that action — e.g., how often do they click in order to fact-check? How and why do older adults use Facebook differently? While our data does not allow us to answer these questions, we emphasize that large-scale dataset releases like this one, despite their limitations, are crucial to allowing researchers to see these trends and formulate such follow-on questions.

Implications for Defenses. Platform defenses can aim to prevent mis/disinformation from reaching users (preventing exposure), or warn them when they encounter it (reducing susceptibility) or when they attempt to share it (reducing spread). Our finding that people across the political spectrum were roughly equally likely to click on (but not necessarily re-share) potential mis/disinformation *once exposed* suggests that preventing exposure may be crucial. Put another way, one might question Facebook's role and responsibility in surfacing these posts to susceptible users in the first place.

The fact that people frequently re-share posts *without clicking* on URLs suggests that interventions at sharing time may also be valuable (e.g., a recent Twitter change prompting users to click on URLs before retweeting them [22]).

Regarding interventions at the time of a user's exposure — such as “false news” labels on posts — our dataset unfortunately provides limited insight. Because most shares of URLs happen in the first month or two of their lifetime on Facebook, the month-level granularity of our data prevents us from investigating whether the spread of TPFC-False URLs decreases significantly once they have been fact-checked (and thus labeled in users' feeds). We can confirm from our data that the scope of Facebook's fact-checking efforts is limited (to already viral and clear-cut false cases), meaning that Facebook's fact-checking alone can address only the tip of the iceberg of potential mis/disinformation on the platform. That said, Facebook's methods indeed succeed at fact-checking highly popular mis/disinformation URLs with potentially great impact.

6 CONCLUSION

We conducted an exploratory analysis of a large-scale dataset of URLs shared on Facebook in 2017–2019, investigating how and how much Facebook users were exposed to and interacted with posts containing potential mis/disinformation URLs, and how these interactions differed across demographic groups. We find that potential mis/disinformation URLs received substantial user engagement, particularly from older adults and from U.S. politically-right leaning users (though not uniformly), and add to a rich and growing literature on mis/disinformation on social media. There are many additional questions we could have investigated in this dataset (e.g., considering different URL subsets), and many more questions that this particular dataset cannot answer. We hope that our exploratory analysis provides a foundation for continued investigations.

ACKNOWLEDGEMENTS

We are very grateful to Social Science One and to Facebook for their partnership in making the Facebook URL Shares dataset available to researchers including ourselves, and to the Facebook team for their logistical and technical support working with the data. We also wish to thank Sewoong Oh for his guidance on working with differentially private data. We thank our anonymous reviewers and our shepherd, Drew Springall, for their feedback. We are also grateful to Scott Bailey for reviewing an earlier draft of this paper. We thank Eric Zeng for providing us with the updated low-quality news URL list. Other members of the UW Security and Privacy Research Lab also helped contribute feedback and analysis ideas. This work was supported in part by a Social Media and Democracy Research Grant from the Social Science Research Council, by the National Science Foundation under Awards CNS-1513584 and CNS-1651230, and by the UW Center for an Informed Public and the John S. and James L. Knight Foundation.

REFERENCES

- [1] Ahmer Arif, Leo Graiden Stewart, and Kate Starbird. 2018. Acting the Part: Examining Information Operations Within #BlackLivesMatter Discourse. *Proceedings of the ACM on Human-Computer Interaction (CSCW)* 2, Article 20 (Nov. 2018), 27 pages.
- [2] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [3] Pablo Barberá, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau. 2015. Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science* 26, 10 (2015), 1531–1542.
- [4] Alessandro Bessi, Fabiana Zollo, Michela Del Vicario, Michelangelo Puliga, Antonio Scala, Guido Caldarelli, Brian Uzzi, and Walter Quattrociocchi. 2016. Users Polarization on Facebook and Youtube. *PLOS ONE* 11, 8 (08 2016).
- [5] Leticia Bode and Emily K. Vraga. 2015. In Related News, That Was Wrong: The Correction of Misinformation Through Related Stories Functionality in Social Media. *Journal of Communication* 65, 4 (2015), 619–638.
- [6] Alexandre Bovet and Hernán A. Makse. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature Communications* 10, 7 (2019).
- [7] Ceren Budak. 2019. What Happened? The Spread of Fake News Publisher Content During the 2016 U.S. Presidential Election. In *The World Wide Web Conference (WWW)*.
- [8] Peter Burger, Soeradj Kanhai, Alexander Pleijter, and Suzan Verberne. 2019. The reach of commercially motivated junk news on Facebook. *PLOS ONE* 14, 8 (08 2019).
- [9] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Third Conference on Theory of Cryptography (TCC)*. 20.
- [10] Georgina Evans and Gary King. 2020. Statistically Valid Inferences from Differentially Private Data Releases, with Application to the Facebook URLs Dataset. (2020). <https://j.mp/38NrmRW>.
- [11] Facebook. [n. d.]. Fact-Checking on Facebook. ([n. d.]). <https://www.facebook.com/business/help/2593586717571940>.
- [12] Jessica T. Feezell and Brittany Ortiz. 2019. ‘I saw it on Facebook’: An experimental analysis of political learning through social media. *Information, Communication & Society* (2019), 1–20.
- [13] Martin Flinthenam, Christian Karner, Khaled Bachour, Helen Creswick, Neha Gupta, and Stuart Moran. 2018. Falling for Fake News: Investigating the Consumption of News via Social Media. In *CHI Conference on Human Factors in Computing Systems*.
- [14] Adam Fourney, Miklós Z. Rác, Gireeja Ranade, Markus Mobius, and Eric Horvitz. 2017. Geographic and Temporal Trends in Fake News Consumption During the 2016 US Presidential Election. In *ACM Conference on Information and Knowledge Management*.
- [15] Deen Freelon. 2018. Computational Research in the Post-API Age. *Political Communication* 35, 4 (2018), 665–668.
- [16] R Kelly Garrett and Shannon Poulsen. 2019. Flagging Facebook Falsehoods: Self-Identified Humor Warnings Outperform Fact Checker and Peer Warnings. *Journal of Computer-Mediated Communication* 24, 5 (10 2019), 240–258.
- [17] Christine Geeng, Savanna Yee, and Franziska Roesner. 2020. Fake News on Facebook and Twitter: Investigating How People (Don't) Investigate. In *ACM Conference on Human Factors in Computing Systems (CHI)*.

- [18] Nir Grinberg, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. Fake news on Twitter during the 2016 U.S. presidential election. *Science* 363, 6425 (2019).
- [19] Andy Guess, Kevin Aslett, Joshua Tucker, Richard Bonneau, and Jonathan Nagler. 2021. Cracking Open the News Feed: Exploring What U.S. Facebook Users See and Share with Large-Scale Platform Data. *Journal of Quantitative Description: Digital Media* 1 (April 2021).
- [20] Andrew Guess, Jonathan Nagler, and Joshua Tucker. 2019. Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances* 5, 1 (2019).
- [21] Andrew M. Guess, Brendan Nyhan, and Jason Reifler. 2020. Exposure to untrustworthy websites in the 2016 US election. *Nature Human Behavior* 4 (2020), 472–480.
- [22] Taylor Hatmaker. 2020. Twitter plans to bring prompts to 'read before you retweet' to all users. (Sept. 2020). <https://techcrunch.com/2020/09/24/twitter-read-before-retweet/>.
- [23] Eslam Hussein, Prerna Juneja, and Tanushree Mitra. 2020. Measuring Misinformation in Video Search Platforms: An Audit Study on YouTube. *Proceedings of the ACM on Human-Computer Interaction* 4 (05 2020).
- [24] Daniel Jolley and Karen M. Douglas. 2017. Prevention is better than cure: Addressing anti-vaccine conspiracy theories. *Journal of Applied Social Psychology* 47, 8 (2017), 459–469.
- [25] David Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily Thorson, Duncan J. Watts, and Jonathan Zittrain. 2018. The science of fake news. *Science* 359 (2018), 1094–1096.
- [26] Ji Young Lee and S. Shyam Sundar. 2013. To Tweet or to Retweet? That Is the Question for Health Professionals on Twitter. *Health Communication* 28, 5 (2013), 509–524.
- [27] Yanqin Lu. 2019. Incidental Exposure to Political Disagreement on Facebook and Corrective Participation: Unraveling the Effects of Emotional Responses and Issue Relevance. *International Journal of Communication* 13 (2019).
- [28] Binny Mathew, Ritam Dutt, Pawan Goyal, and Animesh Mukherjee. 2019. Spread of Hate Speech in Online Social Media. In *10th ACM Conference on Web Science (WebSci)*.
- [29] Solomon Messing, Christina DeGregorio, Bennett Hillenbrand, Gary King, Saurav Mahanti, Zagreb Mukerjee, Chaya Nayak, Nate Persily, Bogdan State, and Arjun Wilkins. 2020. Facebook Privacy-Protected Full URLs Data Set. Harvard Dataverse. (2020). <https://doi.org/10.7910/DVN/TDOAPG>.
- [30] Mia Moody-Ramirez and Andrew B Church. 2019. Analysis of Facebook Meme Groups Used During the 2016 US Presidential Election. *Social Media + Society* 5, 1 (2019).
- [31] NewsGuard. [n. d.]. NewsGuard: Restoring Trust & Accountability. ([n. d.]). <https://www.newsguardtech.com/>.
- [32] Donie O'Sullivan. 2020. Facebook fact-checkers to Trump supporters: We are not trying to censor you. (Oct. 2020). <https://www.cnn.com/2020/10/29/tech/fact-checkers-facebook-trump/index.html>.
- [33] Gordon Pennycook, Adam Bear, Evan Collins, and David G. Rand. 2020. The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings. *Management Science* (Feb. 2020).
- [34] Quintly. 2018. Changes to Facebook's API - February 6th, 2018. (Feb. 2018). <https://support.quintly.com/hc/en-us/articles/115004414274-Changes-to-Facebook-s-API-February-6th-2018>.
- [35] Gustavo Resende, Philippe Melo, Hugo Sousa, Johnnatan Messias, Marisa Vasconcelos, Jussara Almeida, and Fabricio Benevenuto. 2019. (Mis)Information Dissemination in WhatsApp: Gathering, Analyzing and Countermeasures. In *The World Wide Web Conference (WWW)*.
- [36] Mattia Samory and Tanushree Mitra. 2018. Conspiracies Online: User discussions in a Conspiracy Community Following Dramatic Events. In *International AAAI Conf. on Web and Social Media*.
- [37] Ana Lucia Schmidt, Fabiana Zollo, Michela Del Vicario, Alessandro Bessi, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. 2017. Anatomy of news consumption on Facebook. *Proceedings of the National Academy of Sciences* 114, 12 (2017), 3035–3039.
- [38] Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Kaicheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. The spread of low-credibility content by social bots. In *Nature Communications*.
- [39] Elisa Shearer and Elizabeth Grieco. 2019. Americans are wary of the role social media sites play in delivering the news. Pew Research Center. (Oct. 2019). <https://www.journalism.org/2019/10/02/americans-are-wary-of-the-role-social-media-sites-play-in-delivering-the-news/>.
- [40] Social Science One. [n. d.]. Building Industry-Academic Partnerships. ([n. d.]). <https://socialscience.one/>.
- [41] Social Science Research Council. [n. d.]. Social Media and Democracy Research Grants. ([n. d.]). <https://www.ssrc.org/fellowships/view/social-media-and-democracy-research-grants/>.
- [42] Olivia Solon. 2020. Sensitive to claims of bias, Facebook relaxed misinformation rules for conservative pages. (Aug. 2020). <https://www.nbcnews.com/tech/news/sensitive-claims-bias-facebook-relaxed-misinformation-rules-conservative-pages-n1236182>.
- [43] Samuel Spies. 2020. How Misinformation Spreads. (July 2020). <https://mediawell.ssrc.org/literature-reviews/how-misinformation-spreads/versions/1-0/>.
- [44] Kate Starbird, Ahmer Arif, Tom Wilson, Katherine Van Koeveering, Katya Yefimova, and Daniel Scarnecchia. 2018. Ecosystem or Echo-System? Exploring Content Sharing across Alternative Media Domains. In *International AAAI Conference on Web and Social Media (ICWSM)*.
- [45] Edson C. Tandoc, Jr. 2019. Tell Me Who Your Sources Are. *Journalism Practice* 13, 2 (2019), 178–190.
- [46] Rebekah Tromble, Andreas Storz, and Daniela Stockmann. 2017. We Don't Know What We Don't Know: When and How the Use of Twitter's Public APIs Biases Scientific Inference. (Nov. 2017). <https://ssrn.com/abstract=3079927>.
- [47] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the National Academy of Sciences* 113, 3 (2016), 554–559.
- [48] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.
- [49] Claire Wardle and Hossein Derakhshan. 2017. *Information Disorder: Toward an interdisciplinary framework for research and policymaking*. Technical Report. Council of Europe.
- [50] Angela Xiao Wu and Harsh Taneja. 2020. Platform enclosure of human behavior and its measurement: Using behavioral trace data against platform episteme. *New Media & Society* (2020).
- [51] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2020. Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites. In *Workshop on Technology and Consumer Protection*.