

Title: Internal Facebook investigation appendix for “Research note: Examining potential bias in large-scale censored data”

Authors: Jennifer Allen (1), Markus Mobius (2), David M. Rothschild (2), Duncan J. Watts (3,4,5)

Date: July 26th, 2021

Note: The material contained herein is supplementary to the article named in the title and published in the Harvard Kennedy School (HKS) Misinformation Review.

Appendix 1: Internal Facebook investigation

An internal analysis performed by Facebook, in response to the above findings, provides further evidence that fake news was over-represented in the URLs dataset due to the 100-public-share threshold. After we shared our findings, researchers at Facebook examined the URLs labeled as false by their third-party fact-checker program from 1/1/2017-12/31/2018. Note that the Facebook team is examining individual URLs marked false, rather than publisher-level designations of fake used in the above analysis.

The Facebook team found that the number of unique false news URLs as a proportion of all unique URLs in the dataset increased by two orders of magnitude when a 100 public share threshold was implemented, going from 0.00023% with no threshold to 0.022% with one (a two-magnitude increase). Additionally, when the researchers focused on the total public URL shares (a closer analog to the click analysis that we did previously, but clicks are more common than shares) rather than the number of unique URLs, they found that the proportion of false news URLs shares doubled with the 100-share threshold, going from 0.091% with no threshold to 0.185% with the 100-share threshold. These data reveal how skewed the distribution of false news sharing really is. There are 30,500,000 public shares of the 37,860 unique false URLs, yet 30,300,000 of those public shares come from the 8,660 unique false URLs that have over 100-shares. That is, just 200,000 of the 30,500,000 million shares are from URLs shared less than 100 times publicly, even though a significant majority of false news URLs are under that threshold. A summary of these findings can be found in Table A1 below.

Table A1. Results of Facebook’s internal analysis of how the number and composition of URLs in the URL dataset changes when a 100 public share threshold is applied.

	All URLs		False URLs		False to All URL Ratio	
	No Threshold	>100-shares	No Threshold	>100-shares	No Threshold	>100-shares
Unique URL Count	16.2 bil.	38.2 mil.	37,860	8,660	0.00023%	0.022%
Public Share Count	33.7 bil.	16.3 bil.	30.5 mil.	30.3 mil.	0.091%	0.185%

The Facebook team also explored how changing the threshold changes the proportion of fake news URLs in the dataset, the results of which are shown in Figure A1. While they did not find the sharp discontinuity from 0 to 1 that we saw in our CrowdTangle analysis (perhaps due to different definitions of “public share” for CrowdTangle vs. the Facebook URLs dataset), their analysis does show that the number of unique URLs drops off much more sharply at low thresholds (around 10 public shares) for non-false news URLs than for false news URLs.

This analysis differed from the one presented in the main text largely due to privacy and dataset limitations. Due to privacy limitations, the Facebook analysis was conducted at the share rather than click level, since Facebook could not share any data on clicks below the 100-share threshold. Additionally, they defined URLs as false on the URL level rather than domain level, since they as a rule did not want to perform analysis on the de-anonymized data based on an external research team’s domain list (which could have violated privacy in unknown ways and would have to be subject to an intensive privacy review). Unfortunately, this prevents apples-to-apples comparison with the Nielsen dataset, since we do not have

share data for Nielsen, nor do we have URL-level false news labels for below the 100-share hold level. However, substantively, both analyses—ours and Facebook’s—ask the same question about whether fake news is overrepresented in the Facebook URLs dataset. Despite differences in methodologies, we believe that the internal results from Facebook largely replicate the findings in our analysis using external data sources. In particular, they give credence to the claim that the 100-share threshold can introduce unexpected bias in the URLs dataset and that bias causes the share of fake news to be overestimated in the Facebook URLs dataset.

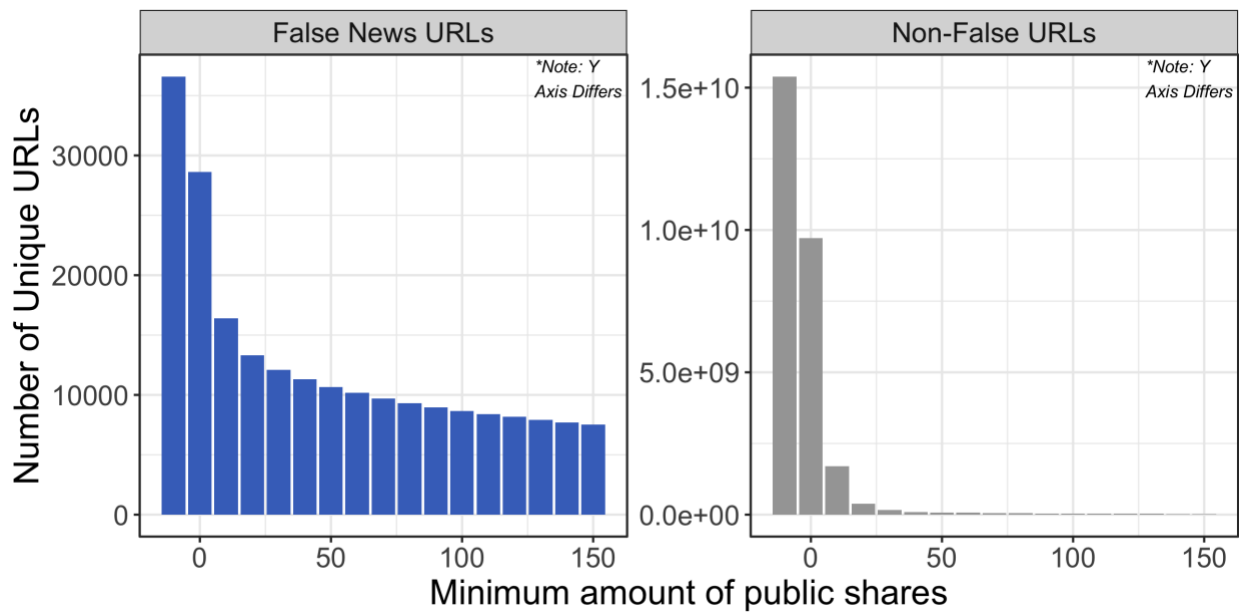


Figure A1. Number of unique URLs that are eligible for inclusion in the Facebook URLs dataset at different public share thresholds. The threshold starts below 0 due to the Laplacian noise applied to the number of public shares. Note that the scales of the Y axes differ for false and non-false URLs.