



Addendum

Addendum to: Research note: Examining potential bias in large-scale censored data

Addendum to HKS Misinformation Review “Research note: Examining potential bias in large-scale censored data” (<https://doi.org/10.37016/mr-2020-74>), published on July 26, 2021.

Authors: Jennifer Allen (1), Markus Mobius, (2) David M. Rothschild, (2) Duncan J. Watts (3,4,5)
Affiliations: (1) Sloan School of Management, Massachusetts Institute of Technology, USA, (2) Microsoft Research, USA, (3) Department of Computer and Information Science, University of Pennsylvania, USA, (4) Annenberg School of Communication, University of Pennsylvania, USA (5) Operations, Information, and Decisions Department, University of Pennsylvania, USA
How to cite: Allen, J., Mobius, M., Rothschild D. M., & Watts, D. J. (2022). Addendum to: Research note: Examining potential bias in large-scale censored data. *Harvard Kennedy School (HKS) Misinformation Review*, 3(1).
Received: January 20th, 2022. Published: February 24th, 2022.

In Allen (2021), we examined bias in Facebook’s 10-trillion cell URLs dataset hosted with research consortium Social Science One caused by censoring any URL with less than 100 public shares from the data. The data includes the number of views, clicks, shares, comments, likes, loves, angers, hates, sorries, and wows on Facebook, as well political identity, country, age, gender (and a few additional attributes) for every URL shared publicly 100 or more times since January 1, 2017. Just a few weeks after the publication of Allen (2021), on September 10, 2021, Facebook announced there was a major error in the U.S. data the company provided to Social Science One. Instead of including all site users, Facebook had accidentally excluded data from Facebook users who had no detectable political leanings—a group that amounts to roughly half of Facebook’s users in the U.S. (Timberg, 2021). Eager to determine how this error may have impacted our analyses, we accepted an invitation from the Academic Partnership team at Facebook to review the work we had just published, and get access to the corrected data as soon as possible. Both our team and this team at Facebook, sought to replicate our initial findings (our updated findings, published in this Addendum, match what Facebook found in their own analysis).

The main finding of Allen (2021) was that the descriptive statistics from the Facebook URLs dataset overstate the percentage of clicks on Facebook that are fake news, and news overall, by as much as four times (4x). This happens because the data is censored to URLs with 100 or more public shares, and conditional on being viewed, fake news URLs are more likely be shared publicly than not-fake news URLs, inflating the percentage of fake news in all manners of engagement in the Social Science One data versus all of Facebook. Yet, despite the magnitude of the data omission (roughly 50% of the population, not missing at random), our initial result that fake news consumption on Facebook is inflated in the Social Science One URLs data relative to a representative sample remains unchanged. However, the level of

¹ A publication of the Shorenstein Center on Media, Politics and Public Policy at Harvard University, John F. Kennedy School of Government.

overstatement has decreased slightly (from 4x to 3x) due to a lower level of fake-news consumption in the Social Science One dataset. Indeed, Figure 1 shows that there is a slightly lower percentage of fake news and not-fake news clicks in the updated Social Science One data compared to the original data.

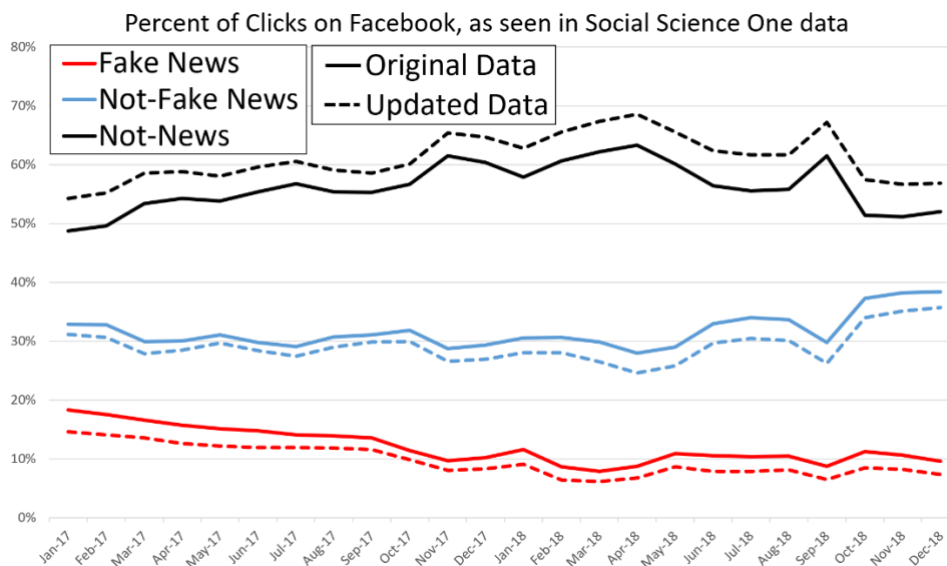


Figure 1. Estimated percentage of clicks on Facebook for fake news, not-fake news, not-news using the raw Facebook URLs data from January 2017 through December 2018. Solid lines are from the original paper and dashed lines are the same output updated to reflect Facebook's correction to Social Science One data.

The main goal of the original paper was to show how the censoring of the Social Science One Facebook data inflated the percentage of fake news clicks. We did this by reconciling the Social Science One data with other representative datasets on which we had performed a censoring process similar to that done on the updated Social Science One dataset (i.e., the excluded users clicked on fake news at a much lower rate than the users in the original data). Table 1 recreates the key table in Allen (2021) (also Table 1) where we show how similar the Social Science One results (column b1) are to either censoring our representative sample to the URLs that made it into the Social Science One data (column c) or censoring the representative sample using CrowdTangle's censoring (column d).

In the original paper, we got an estimate of 10% fake news clicks for the Social Science One dataset, which was somewhat above the 7% for the other two methods designed to mimic the censoring process. However, with the updated data, we got an updated estimate of 7% fake news clicks for the Social Science One data (column b2), which matches the other two methods exactly. The closing of the gap between the estimates actually helps to address questions from our initial paper about whether other differences between the datasets, like mobile vs. desktop use, could be responsible for the discrepancies between the prevalence estimates. Now that the estimates match more closely, we are even more confident that it is the censoring process that is responsible for the difference in fake news prevalence estimates between the Social Science One dataset and the other representative datasets.

Table 1. Estimated percentage of clicks on Facebook for fake news, not-fake news, and not-news under various data and assumptions in December 2018. Red column is added to reflect Facebook’s correction to Social Science One data.

	(a) Nielsen: from Facebook	(b1) Facebook URLs dataset, Old Version	(b2) Facebook URLs Dataset, New Version	(c) Nielsen: in Facebook URLs Dataset	(d) Nielsen: >100 CrowdTangle Public Shares
Fake-News	2.50%	10%	7%	7%	7%
Not-Fake News	24%	38%	36%	39%	42%
Not-News	74%	52%	57%	55%	51%

When Facebook announced this error in their Social Science One data there was concern among researchers that the error could further erode trust in the data that Facebook provides about the use of its platform; Social Science One itself was announced in an effort to foster trust in the wake of the Cambridge Analytica scandal (Alba, 2021). Yet, we can report that the error actually made Facebook overstate the percentage of clicks on fake news, casting Facebook in a worse light than the reality. And luckily, this error did not appear to have altered our results substantially.

Bibliography

- Alba, Davey. (2021, September 10). Facebook sent flawed data to misinformation researchers. *The New York Times*. <https://www.nytimes.com/live/2020/2020-election-misinformation-distortions#facebook-sent-flawed-data-to-misinformation-researchers>
- Allen, J., Mobius, M., Rothschild D. M., & Watts, D. J. (2021, July 26). Research note: Examining potential bias in large-scale censored data. *Harvard Kennedy School (HKS) Misinformation Review*, 2(4). <https://doi.org/10.37016/mr-2020-74>
- Timberg, Craig. (2021, September 10). Facebook made big mistake in data it provided to researchers, undermining academic work. *The Washington Post*. <https://www.washingtonpost.com/technology/2021/09/10/facebook-error-data-social-scientists/>