

Crowdsourced Fact-Checking: A Scalable Way to Fight Misinformation on Social Media

JENNIFER ALLEN, ANTONIO A. ARECHAR, DAVID G. RAND, MIT Sloan School of Management
GORDON PENNYCOOK, Hill/Levene Schools of Business, University of Regina

1. INTRODUCTION

With concerns about fake news growing in the leadup to the 2020 U.S. presidential election, many people have questioned social media platforms' ability to combat disinformation. In response, Facebook, Twitter, and Google have invested in fact-checking as a way to combat misconception. However, although many studies have shown that fact-checking can be effective in correcting misconceptions, the strategy has problems with both scalability and trust.¹ Fact-checking is a laborious process that cannot keep pace with the enormous amount of content on social media. For example, according to a recent article published by The Hill, Facebook's six fact-checking partners have a combined 26 full-time staff that fact-check roughly 200 pieces of content per month -- a tiny fraction of potentially inaccurate content on Facebook.² Furthermore, according to a Poynter study, 70% of Republicans and 50% of Americans overall think that fact-checkers are biased and distrust fact-checking corrections.³ Thus, despite fact-checking's potential, these drawbacks have hindered its widespread adoption.

Our study explores a solution to both of these problems: applying the "wisdom of crowds" to fact-checking. Previous work has shown that laypeople are surprisingly good at distinguishing high from low quality sources.⁴ Here we ask how well laypeople can tackle the substantially harder problem of ratings the veracity of individual articles. As part of a collaboration with Facebook, our team was granted access to a set of articles that were flagged as potentially problematic by Facebook's algorithms and ran two studies in which laypeople rated the articles for accuracy. In the first study, we asked how well layperson judgments correlated the judgments of professional fact-checkers and how large of a "crowd" of laypeople was required to achieve reasonable results. Then, in a follow up study, we asked how displaying source information about the article to laypeople affects results.

2. MATERIALS AND METHODS

Dataset: We examine a set of 209 URLs identified as potentially inaccurate by Facebook's internal algorithm.

Fact-checking: We recruited three separate fact-checkers from the freelancing site Upwork who had prior experience in fact-checking. We asked them to rate all articles using seven questions related to the accuracy and bias of the articles (7-point Likert scales). The fact-checkers did extensive research for each article and provided justifications for their judgments.

Crowdsourcing: For study 1, we recruited participants from Amazon Mechanical Turk ($n = 791$) and asked each of them to rate a subset of 50 randomly selected articles from the 209 URL dataset. Each article was rated by an average of 96 participants. Participants were shown only the headline and lede of the article, and were asked to judge each to the best of their ability using the same seven questions about accuracy and bias that were given to the fact-checkers.

For study 2, we recruited participants ($n = 1204$) to rate 20 articles each, with each article rated by on average of 104 participants. Additionally, half of participants were randomly assigned to a condition that displayed the source domain of the article (e.g. *breitbart.com*) in addition to the headline and lede of the article.

For both studies, participants were not asked to do any research or provide a source for their claims, but rather to rely on their own judgment. Participants were also asked a series of demographic questions, including their political preference, following the study.

3. RESULTS

3.1 STUDY 1

For each article, we took the average of the seven questions among Democrats and Republicans, respectively, and then averaged those two ratings to create a politically balanced layperson rating. These layperson ratings were then correlated with the average rating from the three fact-checkers, and the resulting correlation was compared to the average correlation amongst the three fact-checkers ($r = .62$).

Within this analytic framework, we examined the impact of the the size of the crowd on the correlation between layperson and fact-checker ratings. We did so using bootstrap-style simulations in which the specified number of Democrat and Republican ratings were drawn with replacement from the full set of ratings for each article.

The results are shown in Figure 1. As can be seen, after roughly 20 responses (evenly balanced between Democrats and Republicans), the performance of the crowd matched the performance of fact-checkers; and above that actually out-performed the fact-checkers.

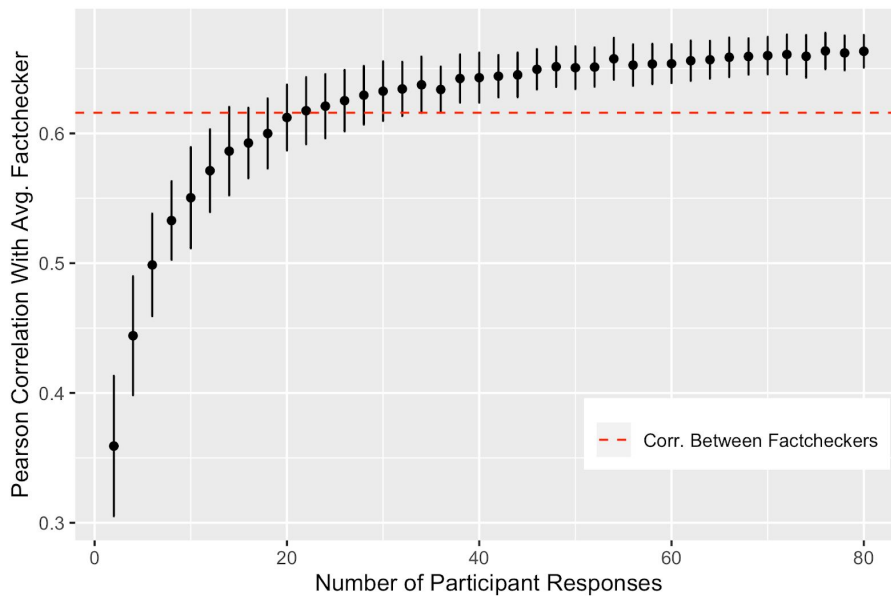


Fig 1: Correlation across articles between politically-balanced layperson headline ratings and average fact-checker research-based ratings, as a function of the number of layperson ratings per headline.

3.2 STUDY 2

In the second study, we examined whether adding source information increases the correlation between laypeople and fact-checkers. Previous work suggests adding source information could improve laypeople’s ability to detect false news when there is a mismatch between headline plausibility and source trustworthiness⁵. Thus, in our set of headlines that are for the most part relatively implausible (since they are headlines that were flagged by Facebook’s algorithm as potentially false), learning that the headline is from a trusted source could improve discernment.

We calculated separate politically-balanced, all Democrat, and all Republican layperson ratings for both the source and no-source condition. We then ran the same bootstrap procedure as in study 1, the results of which are shown in Figure 2. As can be seen, we found that adding source information increases the correlation between laypeople and fact-checkers. Interestingly, we found that this increase is higher for Republicans than for Democrats. These results imply that this increase in agreement is not driven by common bias against conservative outlets by fact-checkers and Democratic raters, as might have been predicted. Rather, these results are consistent with previous results that found that Democrats and Republicans largely agree on the relative trustworthiness of mainstream sources.⁴




Fig 2: Correlation across articles between layperson headline ratings and average fact-checker research-based ratings for the source and no source conditions, shown for a balanced crowd, a Democratic crowd, and a Republican crowd

4. CONCLUSION

The experiments we have presented here provide evidence in support of crowdsourcing's ability to detect misinformation. We find that, given only the headline and lede of an article, a crowd of 20 laypeople can match the performance of fact-checkers researching the full article. We also provide some practical guidance for those wishing to employ such an approach: provide information about the headline's source. Together, our results suggest that crowdsourcing could be a powerful tool for scaling fact-checking on social media, especially in combination with existing fact-checking efforts.

REFERENCES

1. Pennycook, G. & Rand, D. G. The implied truth effect: Attaching warnings to a subset of fake news stories increases perceived accuracy of stories without warnings. *Available at SSRN* 3035384 (2017).
2. Mali, M. Critics fear Facebook fact-checkers losing misinformation fight. *TheHill* <https://thehill.com/policy/technology/478896-critics-fear-facebook-fact-checkers-losing-misinformation-fight> (2020).
3. Flamini, D. Most Republicans don't trust fact-checkers, and most Americans don't trust the media - Poynter. *Poynter* <https://www.poynter.org/ifcn/2019/most-republicans-dont-trust-fact-checkers-and-most-americans-dont-trust-the-media/> (2019).
4. Pennycook, G. & Rand, D. G. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 2521–2526 (2019).
5. Dias, N., Pennycook, G. & Rand, D. G. Emphasizing publishers does not effectively reduce susceptibility to misinformation on social media. (2020) doi:10.37016/mr-2020-001.