Facebook's responsibility

By Ethan Porter

People do not come to believe in misinformation out of nowhere. First, they must be exposed to the misinformation; then, they must *not* be exposed to a correction of the misinformation. Viewed in this light, belief in misinformation is a later step in a long process that implicates not only human psychology, but the architecture of online platforms. These platforms, too, do not come from nothing. Their design reflects human decision-making — decisions that can have profound effects on whether or not people believe misinformation.

Decisions that Facebook has made about fact-checking illustrate how particular decisions about platform design can facilitate belief in misinformation. Some of the company's less defensible choices in this area are well-known. Company leaders, for example, have reportedly prevented misinformation disseminated by political figures from being corrected. The company is aware that its labeling of various posts by Trump as false has proven ineffective, but it has done little to curb misinformation spread by Trump and his allies. In the words of one top executive, it is not the company's role to "intervene when politicians speak."

This vision of non-interference plays out in the everyday approach that Facebook takes to fact-checking. So far as I can tell, the company's present policy works as follows. If a user shares misinformation on the platform, other users may report it, and the misinformation may be fact-checked. If the fact-check finds that the misinformation is indeed false, then Facebook may apply a warning label to the post. But the original poster, and those who saw the false misinformation, are never compelled to see the fact-check. If, say, your uncle posts a story alleging that Venezuelan communists stole votes on behalf of Joe Biden, and a Facebook fact-check judges the story to be false, your uncle will never be directly confronted with the fact-check, nor will anyone who saw the post because of him. The fact-check will live on his newsfeed, but no one — including him — will ever *have* to see it, surely minimizing its impact.

Perhaps these decisions would be defensible were fact-checks themselves ineffective. In general, however, this is not the case. Across several articles and a book, Thomas J. Wood and I have repeatedly found that factual corrections reduce belief in misinformation. After seeing a fact-check, the typical person responds by becoming more factually accurate than they would have been had they not seen the fact-check. The accuracy gains caused by fact-checks are not limited to one side of the aisle or

another. Conservatives who are shown factual corrections of fellow conservatives have become more factually accurate, and liberals have behaved similarly. While there was earlier concern among researchers about the potential for factual corrections to "backfire," and cause people to become more inaccurate, the updated consensus among scholars finds the opposite: Backfire is, at best, vanishingly rare. Insofar as they increase factual accuracy, and lead people away from believing in misinformation, fact-checks work.

Fact-checks can work on Facebook, too, or at least on a simulation thereof. In a recent paper, "Misinformation on the Facebook News Feed," Wood and I administered a fact-checking experiment on a platform meticulously engineered to resemble the real Facebook. (Note: The platform was built in partnership with the group Avaaz. Our partnership with Avaaz did not result in any financial benefits for us, and we were free to report the results of the study as we saw fit.) We recruited large, nationally representative samples via YouGov. Study participants logged on to the platform and saw a news feed which randomly displayed 0-5 fake news stories. Participants then saw a second news feed, which contained a randomly assigned number of fact-checks of any of the fake stories they might have seen. Just like on the real Facebook, participants could choose what to read. They could have chosen not to read the misinformation or the factual corrections. Or they could have read them and been unaffected by them. Either behavior would have rendered the fact-checks useless.

Instead, our evidence indicates that the fact-checks had considerable effects on factual accuracy. We ran the experiment twice, making slight tweaks to the design of the fake Facebook each time, so as to better resemble the real thing. To measure effects on accuracy, we relied on a five-point scale, with questions about the content of each fake item. Across both experiments, our results once again demonstrated the ability of fact-checks to reduce belief in misinformation. On average, when weighting for sample size, the mean "correction effect" — the increase in factual accuracy attributable to corrections — was 0.62 on our five-point scale. Meanwhile, the average "misinformation effect" — the decrease in accuracy attributable to the misinformation alone — was -0.13. On our simulated Facebook, fact-checks decreased false beliefs by far larger amounts than misinformation, sans corrections, increased them.

The version of Facebook used in these experiments was distinct from the real Facebook in several ways. Most importantly for our purposes, it was distinct in the way it showed users fact-checks. While the company's current policy lets fact-checks linger near the bottom of the news feed of users who posted misinformation, and rarely confronts those who read, but did not post, the misinformation with a fact-check, we made our

fact-checks conspicuous for subjects exposed to misinformation. They were also presented with the fact-checks immediately after exposure to misinformation.

If it desired, Facebook could emulate this model. The company could ensure that both posters and consumers of misinformation see fact-checks. It could make such fact-checks impossible to ignore. If you saw misinformation, the next time you logged on, you could see a fact-check at the very top of your news feed. Of course, if Facebook followed this approach, there would often be a longer lag time between exposure to misinformation and exposure to factual corrections than there was in our study. Yet this would still represent a vast improvement over the status quo. By not presenting fact-checks to users who were exposed to misinformation, and by not compelling posters to see fact-checks, Facebook ignores the problem that it has helped create.

By issuing fact-checks to all users exposed to misinformation, and doing so expeditiously and conspicuously, the company could lead many people to greater accuracy, and away from believing in misinformation. But the company has chosen not to do so. It has made this choice in spite of considerable evidence testifying to the effectiveness of fact-checks, some of which is outlined above. Because of these choices, many more people will believe misinformation than would otherwise.

It is easy to blame our friends and relatives for spreading mistruths and believing false claims. To be sure, there is blame to go around. But some of that blame belongs at the feet of Facebook and the particular design decisions the company has made.

Ethan Porter is an assistant professor at George Washington University, where he directs the Misinformation/Disinformation Lab at the Institute for Data, Democracy and Politics.