

Testimony of Laura Edelson, NYU Cybersecurity for Democracy

Before the
Subcommittee on Investigations and Oversight
U.S. House Science, Space, and Technology Committee

Hearing on “The Disinformation Black Box: Researching Social Media Data”

September 28, 2021

Good afternoon Chairman Foster, Ranking Member Obernolte, and the members of the subcommittee. I am extremely grateful for the committee’s attention to the harms being caused by misinformation on social media, and the difficulties researchers face in trying to study this threat to public health, safety, and our democracy.

My name is Laura Edelson, and I am a Ph.D. candidate in Computer Science at New York University. I also co-lead the Cybersecurity for Democracy project at NYU’s Center for Cybersecurity, and I am an ADL Belfer Fellow. As cybersecurity and privacy researchers, my colleagues and I at NYU study systemic vulnerabilities in online platforms that expose people to misleading and outright false claims – from fake Covid-19 cures, to voting disinformation, to investment scams. I believe we will look back and see this moment in history as a turning point, when we stepped up as a society and took action to address what is now a pervasive problem. Previous generations of Americans have taken on similar social ills such as smoking, drunk driving, and industrial pollution. Revelations from investigations by the Wall Street Journal and The New York Times over the past week show that Facebook has internal research demonstrating specific harms and differential treatment and on its platforms, and has considered certain mitigation strategies, only to discard them. These findings which have now been made public echo the conclusions of both my work and the work of other independent researchers: Facebook amplifies misinformation and extreme content. Facebook’s most influential and powerful users are often the ones spreading the most far-reaching misinformation because Facebook doesn’t apply its own policies and rules evenly. Facebook can be harmful to the mental health of its users.

Clearly, we can’t rely on the platforms alone to reduce these harms on their own. Members of this committee will understand that we need to base solutions on independent, rigorous, scientific inquiry based on concrete data. Unfortunately, researchers have been severely hampered not just by lack of such data, but also outright hostility from platforms toward our research. Indeed, this summer, Facebook cut off my team’s access to their data. We used that very data to support [the finding](#) in our recent study that posts from misinformation sources on Facebook got six times more engagement than factual news during the 2020 elections, to identify multiple security and privacy vulnerabilities that we have reported to Facebook, and to audit Facebook’s own, public-facing Ad Library for political ads.

Every day that my team cannot access the data needed to do their research puts us further behind in the race to find answers. Meanwhile, mis- and disinformation continues to contribute to real world harms. Facebook needs to reinstate our accounts immediately so that we can resume our vital work. While we hope that Facebook will do this soon, we must also acknowledge that the platform's voluntary transparency efforts have failed. It's time for Congress to act to ensure that researchers, journalists, and the public have access to the data we need to both study online misinformation and build real solutions.

The Social Media Black Box

Currently, researchers have only limited access to social media data, even when that data is technically public on the platform. Though there are some avenues for study, they fall far short of what is necessary. We really are faced with a black box when it comes to understanding how information flows on these platforms.

Twitter sells access to a small percentage of public tweet data via an interface called the Firehose API. This tool was designed for business use, but researchers who can afford to pay for access to it have used it to study the spread of disinformation and hate on Twitter. Because of its high cost, however, it is out of reach for many researchers based at universities and non-profit organizations. Nonetheless, it represents the most comprehensive dataset that a large platform has allowed researchers to access.

Facebook makes some public Facebook and Instagram content available via its [CrowdTangle](#) platform. CrowdTangle was originally developed, and is still primarily marketed, as a business analytics tool that major Facebook accounts can use to understand how well their content performs on the platform. More recently, Facebook has offered CrowdTangle to researchers and journalists, and the tool does contain vital information for those tracking and studying misinformation on the platform. However, Facebook severely limits journalists' and researchers' access to CrowdTangle. In the past month, I have heard from researchers across a wide range of institutions who have been unable to secure access to this source of data; in fact, the tool seems to be out of reach for the majority of the disinformation research community. Additionally, Facebook maintains an [online ad library](#). However, researchers and developers may access the Ad Library API only if they sign an agreement that limits how they use and share the data, which significantly hampers meaningful publication of any research findings, as the dataset that would be necessary for other researchers to reproduce any findings cannot be publicly shared. Meanwhile, Facebook has shown itself willing (as in my case) to cut off access to the API when it threatens to produce research that is inconvenient or embarrassing to the platform. And the data it does provide is often unreliable. In recent weeks, both the [New York Times](#) and [Politico](#) have reported on separate major data errors in Facebook's transparency tools for researchers that were discovered by academics attempting to use those tools to audit the company's claims.

Other platforms offer even less information. For example, Google has no comparable program to Facebook's for researcher access to public YouTube data. It does [publish](#) transparency

reports with extremely limited data about electioneering ads in the U.S. I consider Google to be the least transparent of the major social media platforms. Finally, most smaller social media platforms, such as TikTok, have no mechanism of any kind for researcher access to either public user generated content or ads.

Many of the transparency efforts that do exist, even for public data, are so access-limited that other researchers cannot reproduce studies that rely on that data. This cuts off a vital part of the scientific process and hamstringing scientific research into what happens online. Lastly, the data that is available via existing transparency efforts are so incomplete that their usefulness for research is needlessly limited. For example, Facebook doesn't make information about how many users have seen a piece of content available through CrowdTangle, even though they could easily do so, and having access to this user impression data would help us better understand why misinformation sometimes spreads so quickly.

As is obvious by now, these avenues for study are grossly insufficient to inform and meet the needs of the public and researchers. In the face of these serious deficiencies, journalists and academics have developed some independent data collection efforts. For example, because the initial focus of my research was focused on the role that advertising plays in disinformation and discrimination online, my team developed Ad Observer, a browser extension that allows users to voluntarily – and anonymously – share information about the ads that Facebook and YouTube show them. Crucially, the tool collects information about the way each ad was targeted by the advertiser. This information – which might reveal, for example, that an advertiser targeted “married men” or “African-American culture – helps us understand how advertisers exploit Facebook's micro-targeting tools, often to spread disinformation and to single out groups they believe will be particularly receptive to it. The Ad Observer tool does not collect personally identifying information, both because that would not meet our ethical standards and because we don't need to do so to answer our scientific questions.

Ad Observer has yielded information we could not get if we relied on Facebook alone. For example, thanks to our 16,000 volunteers, we've been able to collect data from Ad Observer that demonstrates that the archive of political ads Facebook makes available to researchers is missing more than 100,000 ads. Using Ad Observer data, we have also discovered that there are several categories of political ads that Facebook has deliberately and categorically excluded from its political ad archive — such as ads purchased by organizations that Facebook determines to be “news publishers,” and posts that advertisers pay social media influencers to promote.

In addition to my team's project, there are a small number of other research efforts. One notable example is the Markup's [Citizen Browser](#) project. This project, in a little over a year, has uncovered how Facebook continues to allow advertisers to [target users by race](#) via proxy interests even after it said it would end the practice, how it [continues to recommend](#) anti-vaccine groups, and how it [continues to allow](#) financial products to be targeted by age.

One of the reasons there are so few of these projects is because of the legal threats that platforms – and Facebook in particular – have made against independent researchers. Last October, for example, Facebook sent me and my colleague, Professor Damon McCoy, a letter claiming that Ad Observer violated Facebook’s terms of service. Facebook demanded that we deactivate the tool and delete the data that our volunteers had donated for study. Shortly after, your colleagues on the Energy & Commerce committee sent Facebook a letter highlighting both the safety and importance of our work, and encouraging the company to work with us to allow the continued operation of our tool. Regardless, on August 3, after months of negotiations over Facebook’s demand, the company suspended our Facebook accounts. Although the suspension of our accounts has not affected the Ad Observer tool, it has effectively terminated our access to Facebook’s Ad Library API and to CrowdTangle. We relied heavily on both the Ad Library and CrowdTangle in our research. Other researchers and journalists have received similar [threats](#) from [Facebook](#).

While these independent data collection projects are needed, they are not enough on their own to provide the information that we need about the platforms. There are vital holes in our understanding that we simply can’t fill without additional data. The most serious limitation that researchers face today is that there are many platforms, such as YouTube and TikTok, that do not provide even limited researcher access to data. The Mozilla Foundation has published crowdsourced reports illustrating problems such as [political influencers evading a ban on political ads](#) on TikTok, and users’ [stories](#) of YouTube algorithmic recommendations promoting misinformation. We need more data from these platforms to dig more deeply into these concerns.

What Researchers Need, and Why – a Better Black Box

After a plane crash, NTSB investigators seek access to the flight data recorders – the black box. This allows them to analyze what happened so future disasters can be avoided and culpability can be determined. Social media platforms control the equivalent of the black box of data about online misinformation that could help researchers analyze how it contributes to violent events such as the January 6 attack on the Capitol.

There is a great deal of public data on social media platforms, most notably ads, that can and should be made publicly available in a machine-readable format for the use of researchers and journalists from all types of institutions. Last year, I led a public call for [Universal Ad Digital Transparency](#). I believe this proposal should be implemented by all major digital ad platforms immediately. Real harm is still being done by misinformation in digital ads:

- In the summer of 2020, an advertiser called “Protect My Vote” [ran ads discrediting mail-in balloting](#) that appeared to be aimed at African-American voters in the upper Midwest.
- Conservative retirees were targeted with misleading and out-of-context claims in Facebook ads, then guided to sites to convince them to trade in their retirement funds for

precious metals with a company called Metals.com, according to [reporting](#) by journalist Jeremy Merrill for Quartz in 2019.

- While Facebook announced it would remove racial categories as a way advertisers could target ads, The Markup's Citizen Browser project [reported](#) that proxy categories were still in use—examples include “African- American history,” “Hip hop music,” “Latino music,” and “National Hispanic Heritage Month.”

We also need broader access to public, non-paid content on social media platforms. Right now, this could be done by broadening access to platforms like CrowdTangle, and by Google allowing researchers more access to public YouTube data. We know there is harm being done by the rampant spread of misinformation, hate, and misrepresentation online, even if we don't yet know the full extent of the problem.

- Our [forthcoming study](#) shows that, across the political spectrum, posts from news sources that regularly traffic in misinformation have a statistically significant and large engagement advantage—by a factor of six—over posts from news sources that have a record of factualness.
- People who rely on Facebook for information have substantially lower vaccination rates than those who rely on other sources, according to a [survey](#) conducted by the COVID States project in June/July 2021. Of those who rely exclusively on Facebook for news, 25 percent say they do not intend to get vaccinated. This is not an issue of partisanship: both the vaccination odds and the vaccination rates for people who get their news exclusively from Facebook are lower even than for those who get their information exclusively from Fox News.
- A [report](#) from the Anti-Defamation League found that exposure to videos from extremist or white supremacist channels on YouTube remains common, with one in ten study participants being exposed to such content.
- In March 2020, Facebook and Twitter [announced](#) that they removed a network of Russian-backed accounts, originating in Ghana and Nigeria, that targeted Black communities in the U.S. Similar to voter suppression campaigns in 2016, the accounts appeared to be operated by people in the U.S. and attempted to build an audience by posting about Black history, Black excellence and police brutality.
- Nearly 40 percent of Latinx respondents said they've seen material or information that makes them think the COVID-19 vaccine is not safe or effective, according to a [survey](#) earlier this year by Change Research on behalf of the Latino Anti-Disinformation Lab. Another 20 percent said they have directly received wrong or harmful information about the vaccine, primarily on Facebook (53%).

Ad Observer: Why and How we Protect User Privacy

Protecting user privacy is crucial, and Congress and the public are rightly concerned about it. But the good news is we don't need – and definitely don't want– to expose people's private information in order to study misinformation and share our findings with the public. Lung cancer researchers don't publish the names of individual smokers, and we don't need to reveal

identities and people's online browsing habits to expose the systemic vulnerabilities leading to the spread of misinformation.

We believe Facebook exploited concerns about user privacy as a pretext to squelch our research and use us as an example to chill other researchers in our field when they cut off my team's access to data. Our Ad Observer tool does not collect data about our volunteers or their friends. Ad Observer does collect the names and Facebook pages of paying advertisers – examples might be ExxonMobil, Biden for President, The Daily Wire, or the Democratic Underground. What we've learned is that when Facebook said we were violating user privacy, they were talking about advertisers -- not users like you and me. And as Facebook itself makes clear, all Facebook ads are public information.

In their [blog post](#) defending cutting off our access to data, Facebook attempted to lay the blame for their own actions at the feet of the FTC by citing the 2011 consent order they signed. We were extremely grateful when, two days later, Sam Levine, Acting Director of the Consumer Protection Bureau issued a letter clarifying that the FTC consent order does not, in fact, bar Facebook from creating exceptions for good-faith research (like ours) that is in the public interest. Despite this, Facebook has still not restored our accounts.

We go to great lengths to ensure that our tool does not collect personal information about the users who install it, or their friends. Ad Observer has been vetted by Mozilla's security engineers, who [found](#) no privacy concerns. And of course, all of our research protocols are regularly reviewed by NYU's Institutional Review Board, which oversees all our work. Ad Observer meets the highest standards of user privacy and ethical research. Facebook hasn't acted against us because our project is a threat to its users, but because they perceive our research as a threat to themselves and their bottom line. They are attempting to silence science when our findings are inconvenient, and their message to other academics is clear: criticize us at your peril.

Congress must take action to ensure data access needed to study misinformation

After having now spent several years researching this field, I believe that it is time to acknowledge that voluntary transparency has failed. It has failed to protect consumers from dangerous disinformation, failed to provide scientific researchers with sufficient data to make constructive recommendations to the platforms and the public, and failed to be a trusted source to inform users about their practices and consequences. These transparency schemes that platforms have put in place as a stop-gap measure until Congress is able to act are falling prey to many of the common pitfalls of voluntary regulation. Facebook in particular has changed its rules for what it makes "transparent" multiple times, and now it wields its API agreements like a weapon, threatening academics and journalists whose research it doesn't like with the threat it will cut them off from data necessary for their work.

First, Congress should pass a law requiring Universal Digital Ad Transparency now. The biggest digital ad platforms should be required to make all the ads they run publicly available in a machine-readable format. Along with nearly a dozen researchers, I [called](#) for universal digital ad transparency last year. We will soon be publishing a draft proposal that spells out the technical specifications needed in detail.

Second, I believe that a researcher safe harbor law would help protect the many researchers who engage in direct collection of data from platforms. The passage of this law would not directly give researchers access to data, but it would clarify the legality of a great deal of work that currently exists in limbo.

Third, platforms should be required to make public data available to the public: that is, public content with meaningful reach or content from public figures with meaningful audiences should be made available to researchers via tools or searchable interfaces that are accessible to researchers and journalists for analysis. Posting on public pages is analogous to slapping up a notice on a town bulletin board, or writing a letter to the editor. The intended audience is: everybody. Researchers should be able to collect this information for analysis.

Conclusion

In closing, I want to thank the committee for the opportunity to share my experience and perspective. When I began studying for my Ph.D. I did not expect the road I was on would lead here. But I believe that more data and the scientific process is exactly what we need to meet this moment in history, as we grapple with unforeseen consequences of new technologies. Science has helped us meet tough challenges before, helped us save lives and make the lives we save more enjoyable and fulfilling. Science can help us now, but only if we provide researchers the data they need to study and describe the problems we face. Your attention to this topic is vital if we are to make progress, and I know I'm not alone among my colleagues in offering whatever help I can provide. Thank you.

###

Laura Edelson is a Ph.D. Candidate in Computer Science at NYU's Tandon School of Engineering. Laura studies online political communication and develops methods to identify inauthentic content and activity. Her research has powered reporting on social media ad spending in the New York Times, the Wall Street Journal and the Atlantic. Prior to her current time in academia, Laura was a software engineer for Palantir and Factset. During her time in industry, her work focused on applied machine learning and big data.