

Over the past decade, the digital advertising industry has increasingly adopted automated tools to streamline the targeting and delivery of ads. The goal of this targeting is to reach customers who are most likely to be interested in particular products and services. However, in many cases this targeting has instead exacerbated discriminatory and harmful outcomes, particularly related to protected classes and other sensitive categories. For example, ad targeting and delivery systems can be calibrated—either intentionally or by perpetuating historical biases hidden in data—so that users of a certain gender or race do not receive ads for employment opportunities. As a result, these users are deprived of the opportunity to apply for these jobs. The case studies in this report demonstrate that because the digital advertising ecosystem features an array of ads for employment, financial services, housing, and more, these online discriminatory outcomes often have very real offline impacts.

In the context of political ads, the targeted advertising ecosystem can be—and has been—used to spread propaganda and disinformation, enable media manipulation, and even incite genocide. These destructive outcomes heighten concerns regarding fairness, accountability, and transparency around algorithmic decision-making.

This report provides an overview of the advertising platforms and operations of Google, Facebook, and LinkedIn. It also describes the range of advertising-related controversies these platforms have faced, and the subsequent platform and policy changes they have enacted. Google and Facebook are important case studies, as the two companies together account for approximately 60% of the U.S. digital advertising market, with the next closest competitor, Amazon, only accounting for 8.8%. The ad targeting and delivery practices of Google and Facebook are therefore pervasive, and the harmful outcomes that can result from the use of algorithmic decision-making by these platforms are extensive. LinkedIn offers an interesting case study, as the platform’s advertising operations have transformed radically since its acquisition by Microsoft in December 2016. Initially the company’s advertising business model was not a significant component of its revenue stream, and as a result, its targeting and delivery practices were less invasive. However, under Microsoft, the company’s ad targeting and delivery practices have been amended to resemble many of the practices that Google and Facebook implement.

Key Findings

1. Although internet platforms collect vast amounts of user data, they do not sell this data. Rather, they sell the attention of their users to brands and companies who run advertisements on their platforms, and use the data they collect to power their targeted advertising operations. These platforms are therefore still able to monetize and financially benefit from these rampant data collection practices. The financial success of the targeted advertising business model also incentivizes companies to continue and expand these practices.
2. The targeted advertising business model seeks to deliver “relevant” content to users in order to increase user engagement and how long a user spends on the platform. The longer a user spends on a platform, the more ads can be targeted and delivered to them. Therefore, platforms’ desire to deliver relevant content to users are financially motivated.
3. Advertising revenue is based on more than just the price an advertiser will bid to run an ad on a platform. It also relies on an algorithmic determination of which ads will best “engage” users and increase the time they spend on the platform. As a result, when platforms conduct ad auctions, they rely on factors beyond the price of an advertiser’s bid. These factors include algorithmic assessments of how relevant an ad is to a user, as well as the quality of an ad, both of which play into determining which ads are the most likely to keep users on the platform.
4. Both ad targeting and delivery algorithms hold the potential to generate discriminatory and harmful results which often reflect societal biases present in historic training data. Even if an advertiser does not choose to target users based on protected or sensitive characteristics, ad delivery algorithms can deliver ads along these lines based on inferences and subsequent optimization decisions it makes on which categories of users an ad is most relevant to. For example, an employment ad for a job in a field that has been historically dominated by men may end up being delivered mostly to male users due to an algorithmic inference based on historic data that men are more interested in such jobs.
5. Internet platforms offer little transparency and accountability around the automated tools and policies they use to carry out ad targeting and delivery. In addition, few (if any) internet platforms offer their users a comprehensive set of mechanisms with which they can control how advertisers can use their personal information to target them. Many users also lack awareness of how internet platforms target and deliver advertisements.
6. Following the controversy around the misuse of political advertising tools during the 2016 U.S. presidential election, some internet platforms have created political ad databases that provide limited information on what political ads are run on a platform. Generally, these databases do not include granular information such as the reach and engagement of an ad and how an ad was targeted and subsequently delivered. This information is vital for identifying instances of discrimination or other harm. Further, these ad transparency efforts are currently limited to certain geographic regions and they therefore do not provide a comprehensively transparent overview of global political advertising operations on these services.

Recommendations

Going forward, internet platforms, civil society, researchers, and U.S. policymakers should consider the following recommendations in order to promote greater fairness, accountability, and transparency around algorithmic decision-making in this space. Because the First Amendment limits the extent to which the U.S. government can direct how internet platforms decide what content to permit on their sites, we offer limited recommendations for U.S. policymakers.

Internet platforms that offer digital advertising services should:

1. Publish comprehensive and comprehensible descriptions of advertising content policies.
2. Publish comprehensive and comprehensible descriptions of advertising targeting policies.
3. Prohibit targeting based on protected classes and sensitive characteristics that could result in discriminatory outcomes, including characteristics that have been shown to be proxies for protected characteristics.
4. Establish and disclose a comprehensive process to review ads for categories of ads that could have significant real-life consequences such as political, housing, education, employment, and financial services-related ads before they are permitted to run on a platform.
5. Hire independent auditors to conduct regular periodic audits of ad targeting algorithms in order to identify potentially harmful outcomes related to privacy, freedom of expression, freedom of information, and discrimination, and take steps to eliminate or mitigate any harms identified through the audits.
6. Hire independent auditors to conduct regular periodic audits of ad delivery and optimization algorithms in order to identify potentially harmful outcomes related to privacy, freedom of expression, freedom of information, and discrimination, and take steps to eliminate or mitigate any harms identified through the audits.
7. Empower users with comprehensive tools that help them understand how and why ads are targeted and delivered to them.
8. Explain to users why the platform collects, infers, and shares user data.
9. Improve user controls so that users can easily manage whether and how data is collected, inferred, and shared, how this data is used, and how it influences the content that they see. This should include the option to delete this data entirely.
10. Provide clear labels for sponsored and paid content across all of the platform's products, services, and ad networks.

11. Create a publicly available online database of all of the ads that a company has run on its platform.
12. Publish a transparency report that provides a granular overview of the platform's advertising operations across all regions that it operates in.
13. Provide meaningful notice to advertisers who have had their ads or accounts flagged or removed, as well as to users who have flagged ads or accounts.
14. Offer advertisers who have had their ads or accounts flagged or removed, as well as users who have flagged ads or accounts, a robust appeals process. This appeals process should be offered for categories of ads that could have significant real-life consequences, such as political ads, housing ads, employment ads, and credit ads.
15. Fund further research and investigations regarding how the digital advertising ecosystem can be used to reinforce societal biases and discriminatory outcomes, and how to redress these problems.

Civil society organizations and researchers should:

1. Conduct further research on how the digital advertising ecosystem can be used to reinforce societal biases and discriminatory outcomes through advertising, and how to redress these problems.
2. Collaborate to develop a set of industry-wide best practices for transparency and accountability around algorithmic ad targeting and delivery practices. These best practices should explicitly prioritize the public interest above corporate business models and concerns about trade secrets.

U.S. policymakers should:

1. Clarify that all offline anti-discriminatory statutes apply in the digital environment.
2. Enact rules to require greater transparency from online platforms regarding their ad targeting and delivery practices.

Acknowledgments

In addition to the many stakeholders across civil society and industry who have taken the time to talk to us about ad targeting and delivery over the past few months, we would particularly like to thank Dr. Nathalie Maréchal from Ranking Digital Rights and Joe Westby from Amnesty International for their help in closely reviewing and providing feedback on this report. We would also like to thank Craig Newmark Philanthropies for its generous support of our work in this area.