



February 2020

Special Delivery

How Internet Platforms Use Artificial Intelligence to Target and Deliver Ads

Spandana Singh

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.

About the Author(s)

Spandana Singh is a policy analyst in New America's Open Technology Institute.

About New America

We are dedicated to renewing America by continuing the quest to realize our nation's highest ideals, honestly confronting the challenges caused by rapid technological and social change, and seizing the opportunities those changes create.

About Open Technology Institute

OTI works at the intersection of technology and policy to ensure that every community has equitable access to digital technology and its benefits. We promote universal access to communications technologies that are both open and secure, using a multidisciplinary approach that brings together advocates, researchers, organizers, and innovators.

Contents

Introduction	6
The Growth of Today’s Digital Advertising Ecosystem	9
The Role of Data in the Targeted Advertising Industry	11
The Role of Automated Tools in Digital Advertising	15
Concerns Regarding Digital Advertising Policies and Practices	18
Case Study: Google	21
General Overview of Google’s Advertising Platform	21
A Technical Overview of Google’s Advertising Platform	22
Controversies Related to Google’s Advertising Platform	25
Steps Google Has Taken to Adjust Its Advertising Platform Since the 2016 U.S. Presidential Election	28
User Controls on Google’s Advertising Platforms	30
Case Study: Facebook	33
A General Overview of Facebook’s Advertising Platform	33
A Technical Overview of Facebook’s Advertising Platform	34
Controversies Related to Facebook’s Advertising Platform	38
Steps Facebook Has Taken to Adjust Its Advertising Platform Since the 2016 U.S. Presidential Elections	43
User Controls on Facebook’s Advertising Platform	46

Contents Cont'd

Case Study: LinkedIn	48
A General Overview of LinkedIn's Advertising Platform	48
A Technical Overview of LinkedIn's Advertising Platform	49
Controversies Related to LinkedIn's Advertising Platform	51
Steps LinkedIn Has Taken to Adjust Its Advertising Platform Since the 2016 U.S. Presidential Elections	52
User Controls on LinkedIn's Advertising Platform	53
Promoting Fairness, Accountability, and Transparency Around Ad Targeting and Delivery Practices	54
Recommendations for Internet Platforms	54
Recommendations for Civil Society and Researchers	62
Recommendations for Policymakers	62

Introduction

Over the past 20 years, the collection and monetization of internet users' personal and behavioral data for the purpose of delivering targeted advertising has emerged as a ubiquitous business model for the internet economy. According to projections made in 2019 by eMarketer—a market research company that focuses on digital marketing, media, and commerce—U.S. advertisers would spend more than \$129 billion on digital advertising during the year, whereas they would spend only \$109 billion on more traditional advertising methods such as television, radio, and newspapers.¹

The digital advertising industry has grown from an online version of the traditional advertising system, featuring advertising agencies and publishers, to a data-driven environment in which internet platforms such as Google and Facebook take center stage. Today, the two companies account for approximately 60 percent of the U.S. digital advertising market, with the next closest competitor, Amazon, only accounting for 8.8 percent of the industry market share.² On a global level, Google accounts for approximately \$103.73 billion in net digital ad revenues, followed by Facebook with \$67.37 billion, Alibaba with \$29.20 billion, and Amazon with \$14.03 billion.⁴

One of the core components of the success of the online advertising industry has been the introduction of new targeting tools that have enabled advertisers to segment and select their audiences along very specific lines. These targeting tools categorize consumers using a range of data points, which can include demographic characteristics, behavioral information, and personally identifiable information (PII).

Many companies assert that advertising can enhance the lives of their consumers, as it can connect users with products or services that they may find useful and may not otherwise be aware of. Although a user may find a certain advertised product or service relevant, it is important to recognize that the main goal of an advertisement is to drive purchases and revenue, or to elicit changes in behavior or beliefs. In addition, by delivering useful content to users, platforms aim to maximize the amount of time users spend on their services, which in turn increases the number of ads users will view and the amount of revenue the platforms will earn.⁵

Over the past decade, the digital advertising industry has increasingly adopted automated tools to streamline the targeting and delivery of advertisements. The goal of this targeting is to reach customers who are most likely to be interested in particular products and services.⁶ However, in many cases this targeting has instead exacerbated discriminatory and harmful outcomes. It is important to recognize that ad targeting and delivery practices inherently involve “discriminating” among users as they require advertisers to delineate which

categories of users they would like to target; therefore their reach “discriminates” in favor of reaching certain users and against reaching others.

Not all of these delineations are based on sensitive or protected categories, nor do all of them result in harmful outcomes. However, some distinctions can generate particularly damaging outcomes,⁷ especially where they involve discrimination against protected classes or other sensitive categories. This can harm protected classes and exacerbate hidden societal biases. For example, ad targeting and delivery systems could be calibrated and optimized so that users of a certain gender or race do not receive ads for certain employment opportunities. As a result, these users are deprived of the opportunity to apply for these jobs. Given that the digital advertising ecosystem features an array of ads for services such as employment, financial services, and housing, these online discriminatory outcomes often have very real offline impacts.

In addition, the targeted advertising ecosystem can be (and has been) used to spread propaganda and disinformation, enable media manipulation, and even incite genocide.⁸ These destructive outcomes have heightened concerns regarding fairness, accountability, and transparency around algorithmic decision-making. In addition, some researchers have discussed the harms posed by the advertising business model in terms of human rights risks.⁹ Because of these concerns, some researchers, scholars, and members of civil society believe that the targeted advertising business model as a whole must be overhauled in order to enact meaningful change. This is an ongoing debate that will develop as internet platforms further expand their ad targeting and delivery practices.¹⁰ While eliminating this business model would solve many problems described in this report, the advertising business model is likely to last for the immediate future. This report explores actions that could help to mitigate some of these harmful effects.

This report is the third in a series of four reports that explore how automated tools are used by major technology companies to shape the content we see and engage with online. It focuses on the use of automated tools to target and deliver ads to internet users and relies on case studies of three internet platforms—Google, Facebook, and LinkedIn—to highlight the different ways algorithmic tools can be deployed by technology companies to enable ad targeting and delivery. These case studies will also highlight the challenges associated with these practices. This report also offers recommendations on how internet platforms, civil society, and researchers can promote greater fairness, accountability, and transparency around these algorithmic decision-making practices. This report also provides recommendations for policymakers in this regard. However, 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, this report provides only limited recommendations for action by policymakers.

Editorial disclosure: This report discusses policies by Google, Facebook, and LinkedIn. Google, Facebook, and Microsoft (owner of LinkedIn) are funders of work at New America, and LinkedIn's co-founder is on New America's board of directors, but none of these companies or individuals contributed funds directly to the research or writing of this report. New America is guided by the principles of full transparency, independence, and accessibility in all its activities and partnerships. New America does not engage in research or educational activities directed or influenced in any way by financial supporters. View our full list of donors at www.newamerica.org/our-funding.

The Growth of Today's Digital Advertising Ecosystem

Over the past two decades, the digital advertising industry has grown significantly thanks to the introduction of digital tracking tools and the adoption of expansive data collection practices by online platforms. These practices have enabled advertisers to more precisely segment their audiences based on a range of characteristics such as age, education, gender, and ethnicity. In addition, these tools have permitted segmentation based on more nuanced inferences such as personality type, user likes and shares, and predicted behaviors.¹¹ Based on these insights, advertisers are able to deliver tailored ads to users online,¹² often at a much lower cost than through traditional advertising methods. As a result, many advertisers have begun moving away from more traditional forms of advertising, such as print and television.

The targeted advertising business model that is widely used today was pioneered by Google in the late 1990's.¹³ In 2008, Sheryl Sandberg left Google and joined Facebook as its Chief Operating Officer. One of her major first steps was to introduce these digital advertising practices at the company, ushering in a new era for the online advertising ecosystem. Subsequently, the targeted advertising business model rapidly spread across the internet industry. As a result, more and more companies are placing less emphasis on generating revenue in their early days. Rather, they focus on rapidly growing their user bases, collecting vast amounts of data on their users, and then using this data to enable and support ad targeting and delivery. As a result of this model, these internet platforms have been able to significantly monetize and profit off of their users' data.¹⁴

Today, the online advertising ecosystem consists of a number of actors. These include ad publishers, ad networks, ad exchanges, and numerous outlets that enable advertisers to gain reach and traction.

Internet platforms such as Facebook and Twitter are particularly important within the online advertising ecosystem, as they operate "single-site" advertising platforms; in other words, these platforms serve only one specific website. These single-site platforms have revolutionized the ways that ads are targeted and delivered to users, as they have created several centralized locations where advertisers can engage with and target a plethora of individuals.¹⁵ These platforms also operate ad networks that place ads on third-party sites. Other companies, such as Amazon, are also increasingly gaining traction and prominence in the digital advertising industry. Although the company still has a small percentage of the U.S. and global digital advertising industry market share, it is considered a potential future contender with the larger platforms.¹⁶

Telecommunications companies also play a role in the online advertising industry. These companies are responsible for providing users with access to the internet, and seek to provide advertisers with an alternative to the “walled garden” advertising options provided by companies such as Google and Facebook.¹⁷ Over the past few years, telecommunications companies around the world have made significant investments in the digital advertising space. For example, in 2018, AT&T acquired AppNexus, a programmatic advertising marketplace¹⁸ billed as the world’s largest independent digital ad exchange, and a major competitor to Google and Facebook’s ad operations.¹⁹ In addition, in 2015 and 2017 respectively, Verizon acquired the ad-tech stacks of AOL and Yahoo!. The company subsequently established the Oath Ad Platform. However, in 2019 the Oath ad platform was shuttered as part of a strategic business review.²⁰

The Role of Data in the Targeted Advertising Industry

User data has been integral to the rise of digital advertising. The growth in the digital advertising industry would not have been possible without expanding digital data collection practices. As Ben Scott and Dipayan Ghosh outline in their report *Digital Deceit: The Technologies Behind Precision Propaganda on the Internet*, there are a variety of mechanisms that companies in the digital ad ecosystem can use to collect user's data. These can be broken down into four broad categories:²¹

1. **Web Tracking:** When a user visits a website, the website loads built-in web tracking technology. The most common of these technologies is the web cookie, which can be “first-party” or “third-party.” First-party cookies are developed and placed on a website by the website owner. This enables the website to track a user's movements and activity between web pages on the website. Third-party cookies are developed and placed on a website by a third-party entity, in partnership with the website. This enables the third party to monitor and track a user's activity on and across the website, as well as across every website a user visits that has the third-party cookie code embedded on it. Cookies can enable the tracking party to accrue data that can help them infer the interests, preferences, behaviors, and routines of a user based on their behavior across a network of third-party sites. This information can then be used to target these users with specific and relevant ads.

Although cookies are a highly pervasive form of web tracking, users can manage how cookies monitor and collect information on them by clearing cookies in their browser settings. Users can also deploy tools such as the “Do Not Track” feature available on some browsers, which sends a request to websites a user visits to disable its cross-site user tracking, which includes cookies. However, respecting Do Not Track is a decision made by websites, and some advertisers can actually use the fact that users have opted to use Do Not Track use as a signal in browser “fingerprinting,” discussed later in this section.²²

2. **Location Tracking:** Granular location data is an integral component of the digital advertising ecosystem, as it provides a significant amount of information about the interests, preferences, behaviors, and routines of a user. For example, a user's location information can provide insight into where a user lives, where they work, what stores or businesses they regularly visit, and where they spend their free time. This information can be used to determine which ads a user should be targeted with. Smartphone application makers routinely use GPS signals, cellular

network triangulation, Wi-Fi SSIDS, and Bluetooth connectivity to collect such location information.

Users have limited avenues to prevent or manage collection of their location information. These include turning off location settings on their phone, managing the location information access of individual smartphone apps, and opting out of location-based ads on single-site ad platforms. However, it is important to note that many smartphone applications, such as rideshare applications, require location information in order to operate, and they often mandate that users provide constant access to this data in order to use the app.

3. **Cross-Device Tracking:** Consumers, especially in the United States, often access the internet from multiple devices. However, advertisers generally want to avoid delivering duplicate advertisements to consumers across multiple devices. In order to control when and where the ads are delivered, the digital advertising industry has developed cross-device tracking technologies that monitor user activity across devices. Once cross-device inferences are made with high enough confidence, many companies will associate a unique identifier with a user. This identifier becomes a central anchor of user data that is collected across multiple applications, internet platforms, and devices.
4. **Browser Fingerprinting:** A browser fingerprint is a compilation of data about the setup of a user's browser or operating system. This information can include a user's browser provider, browser version, operating system, preferred language, browser plugins (software that is typically third-party that adds functionality to a browser when installed), tracking settings, ad blocker settings, and time zone. Because users are able to customize their browser settings and preferences, their browser fingerprint can be used to identify them across the internet. A 2015 study conducted by the Electronic Frontier Foundation concluded that 84 percent of participating internet users had unique browser fingerprints. Another study has found that websites were able to track a user's browser fingerprint even if they were using a different browser, as a user retains similar browsing characteristics such as similar operating systems, IP addresses, etc.²³ This raises significant privacy concerns. To evade browser fingerprinting some individuals are using Tor, an open-source software that enables anonymous communication. Tor manipulates tracking signals such as IP address and operating system or normalizes them across all Tor Browser users in order to prevent fingerprinting.

Data brokers such as Acxiom, Experian, and Oracle also play an important role in data collection and user targeting. These companies combine user records from a

range of sources, including retail purchases and census data, in order to provide advertising platforms such as Facebook with hundreds of unique data points that they can use to enhance their profile database.²⁴ Internet platforms, advertisers, and third-party data brokers can also use data modeling techniques to make further inferences and predictions about consumer traits and behaviors. Data modeling enables these parties to use information on observed actions and self-reported preferences and interests to supplement and fill in profile information. For example, a data broker could use a user's zip code and name to infer their ethnicity based on census or other data showing ethnic breakdowns by zip code. Data modeling can also be used to categorize users based on factors such as creditworthiness or interest in certain topics.²⁵ In 2018, Facebook patented a system that combines a range of data points to predict "socioeconomic group classification."²⁶

According to Scott and Ghosh, the pervasive use of such methods and the rampant data collection of internet platforms foster a vicious cycle when it comes to the collection of highly personal data. The more highly personal data companies are able to collect, the more relevant advertisements they are able to deliver to users. The more relevant these advertisements are, the longer these companies can keep users on their platforms, thus maximizing the potential advertising space for each user, and driving revenue for the platform.²⁷ However, there is disagreement around what relevant really means in this context. Companies assert that relevant content is content that a user is likely to be interested in based on their online profile and history. But this means that an advertiser could target their advertisements to focus on specific personality types and perceived emotional states. They could also use profiles to target and exploit an already vulnerable category of users.²⁸ This has already been documented. For example, cigarette companies have targeted low-income communities with their advertisements, and similarly vulnerable communities have been the primary audience for products such as diet pills.²⁹ Similarly, in 2013, media planning agency PHD released a study highlighting the dates, times, and occasions when U.S. women felt the least attractive. Based on this, they laid the groundwork for advertisers to target women during "prime vulnerability moments" such as on Monday mornings when the research indicated women "feel least attractive."³⁰ This raises a number of ethical concerns regarding how such relevant advertising can be used to exploit marginalized groups or other individuals when they are most vulnerable for the purpose of driving further revenue.

Although internet platforms collect a significant amount of user data, the largest companies typically do not share or sell these vast datasets. They assert that this is due to concerns regarding user privacy, competitiveness, and the protection of intellectual property. However, the collection of user data is still central to these platform's business models, and although these companies do not sell this data, they are still able to monetize it and gain significant financial benefits. For example, platforms such as Google and Facebook have been able to successfully

profit off of the vertical integration of behavior tracking and ad targeting.³¹ Similarly, by offering advertisers tools to target users based on the platforms' analysis of their users' data, the platforms are effectively selling their users' attention.³²

Today, there are few legal safeguards in the United States that limit what companies can do with their users' data. Although the California Consumer Privacy Act went into effect in January 2020, the United States still lacks a comprehensive federal law to protect consumer privacy. Further, in 2017, President Trump repealed broadband privacy regulations that required internet service providers to “obtain consumer consent before using precise geolocation, financial information, health information, children’s information and web-browsing history for advertising and marketing.”³³ In the United States, internet platforms and civil society organizations have called for comprehensive privacy legislation. However, it is unlikely that any legislation will be enacted soon, as there is little consensus over what specific provisions and safeguards such legislation should cover.

The Role of Automated Tools in Digital Advertising

As internet platforms have explored ways to optimize the ad targeting and delivery process, they have increasingly adopted artificial intelligence and machine-learning tools. Given that advertising relies on making situation-based decisions to determine what content to show a user during a certain scenario, there is significant room for automation and algorithmic decision-making to streamline this process. For example, machine learning systems can be used to optimize the ad delivery process by evaluating and extracting insights from datasets to predict future user behavior and target ads to these users based on these predictions. In addition, automated tools can be used to deliver ad campaigns to users quickly and precisely, based on a range of granular factors such as content format and timing.³⁴

Generally, the advertising operations of internet platforms can be categorized into two components: ad creation and ad delivery. Ad creation is the process by which an advertiser submits an ad to a platform. For most internet platforms, there are three key stages to the ad creation process:³⁵

1. **Producing the ad creative:** Advertisers develop content such as the headline, text, media (e.g. images or video), and the website that an ad should link to. This content is collectively known as the ad creative.³⁶
2. **Audience selection and targeting:** Advertisers select which segments of a platform's users they want to target with their ad.³⁷
3. **Bidding strategy:** Advertisers specify how much they are willing to pay to have their ads shown to users. There are a number of metrics that can be used when making and assessing a bid, including per-impression or per-click bids. Advertisers can also place a maximum cap on their bid and allow the platform to bid for them.³⁸

Once an advertiser has submitted all of this information to a platform, the platform may review the ad, either using automated review, human review, or a combination of the two, to ensure it is in compliance with policies such as its ad content policies, ad targeting policies, and its general platform content policies. It is important to note, however, that not all internet platforms have instituted such a review process, raising concerns that platforms may be allowed to run discriminatory and harmful ads. Once an ad is ready to be run, it will be moved to the ad delivery phase.³⁹ Ad delivery is the process by which an advertising platform shows ads to users. Most ad delivery systems are automated in part or in full. For every ad slot that is available, an ad platform will host an automated auction to determine which ads should be shown to a specific user. A platform

will deliver an ad to a specific set of users based on a range of factors, including the advertiser's budget and how an ad campaign is performing⁴⁰ (as measured by metrics such as engagement or impressions).

Although this seems relatively straightforward, there are some caveats to this process.⁴¹ First, platforms generally avoid showing ads from the same advertiser to a user in quick succession. As a result, a platform's ad delivery system may disregard bids for advertisers who recently won an auction for the same user. Second, platforms assert that one of the benefits of targeted advertising is that it provides users with relevant content that enhances their user experience. Although relevance is a rather subjective notion, some platforms try to quantify it by calculating a relevance score when considering which ads to show a user. However, since the ad delivery decision-making process typically involves algorithmic decision-making, ad platforms do not solely rely on a bid to determine the winner of an auction. This process permits ads with cheaper bids, but higher relevance scores, to win against ads with a higher bid in an auction. This is because platforms benefit financially from delivering relevant content, including ads, to users, as it increases the likelihood users will continue using a platform.⁴² Third, platforms and advertisers may prefer to evenly spread an advertiser's budget over a certain period of time, rather than spending it all at once. This introduces further caveats related to which ads should be considered by platforms for which auctions. It also demonstrates that internet platforms consider more than just an advertiser's bid when determining which ads to deliver.⁴³

Once an advertiser enters the ad delivery phase, they will typically receive data from the internet platform detailing how their ad is performing. This data can include demographic and geographic information on the users who are viewing the ads, as well as some information on the users who actually click or engage with the ad.⁴⁴

Examples of how automated tools can be used during the ad targeting and delivery phases include:

- 1. Dynamic Creative Optimization:** Internet platforms are using automated tools to help advertisers identify which users are most likely to react to a certain ad. During the Dynamic Creative Optimization process, advertisers can tailor ad content for particular users. They can do this by, for example, using A/B testing⁴⁵ to evaluate and collect data on how a user reacts to an advertised price, an ad delivered at a certain time, or an ad delivered in a certain format. During these trials, the advertising system captures engagement statistics. These can be used to inform the future efforts of advertisers and the internet platform. For example, this information can inform the internet platform's algorithms on what the most successful targeting strategies are. The more resources that an

advertiser spends on such tests, the more refined and effective their campaigns are likely to become.⁴⁶

2. **Digital Ad Mediation:** Digital ad mediation is the process of connecting mobile publishers or consumer engagement platforms (such as the Twitter mobile application) with a brand in real time in order to deliver the most relevant ad to a user. Given the volume of content and the number of users on mobile applications, the digital ad mediation process is most effective when internet platforms deploy it using automated tools.⁴⁷

3. **Social Media Management Software:** Social media management companies are increasingly using automated tools to produce, target, and deliver social media campaigns. Leading social media management providers such as Hootsuite are harnessing the power of automated tools to gather detailed information about users' reactions and sentiments related to certain brands. When an event that is relevant to a target user occurs, these tools use these sentiment insights to immediately dispatch an ad that may be relevant to them.⁴⁸ For example, luxury fashion house Givenchy could establish a contingency to initiate promoted ads for their products anytime their spokesperson Ariana Grande appears on television or in a Givenchy-related hashtag. Further, social media management companies are also increasingly adopting machine learning algorithms into their client workflows in order to fuel recommendations regarding audience segmentation, content delivery, and more. These insights aim to significantly expand the reach and influence of brands.⁴⁹

Concerns Regarding Digital Advertising Policies and Practices

The rise of the digital advertising industry has created new roles for brands, publishers, and internet platforms. It has also raised a number of concerns regarding privacy, surveillance, and fairness, accountability, and transparency around algorithmic decision-making practices.

In the early 2000s, a Minnesota father found Target ads for maternity clothing and nursery furniture addressed to his teenage daughter in the mail. According to Target statistician Andrew Pole, the company was able to use historical buying data on all of the women who had signed up for Target baby registries to identify purchasing patterns. This information was used to create an algorithm that could identify women who were likely pregnant. Target then delivered ads to women who had been identified as pregnant. For example, the company's statisticians found that women on the baby registry were buying more quantities of unscented lotion around the beginning of their second trimester. Additionally, many pregnant women purchased supplements such as calcium, magnesium, and zinc during their first 20 weeks. By bundling pattern-based data points such as these together, Target was able to calculate and assign each shopper a "pregnancy prediction score" and estimate each pregnant woman's due date within a narrow window.⁵⁰ In the case of the Minnesota teen, this meant that the company knew she was pregnant and acted on this knowledge before she decided to tell her own family. This was over 10 years ago.⁵¹ Since then, the digital advertising ecosystem has significantly changed, and such practices have become even more refined, more pervasive, more automated, and less visible. Today, the industry relies on and monetizes user data at an unprecedented scale. In this way, data has become the lifeblood of the digital advertising industry. Simultaneously, these rampant data collection and monetization practices have raised a number of concerns.

As Harvard Business School scholar Shoshana Zuboff outlined, the digital advertising industry can be situated within the framework of "surveillance capitalism." In her book *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, Zuboff describes how in the twentieth century, companies such as General Motors and Ford sparked the rise of mass production and managerial capitalism. In the twenty-first century, companies such as Google and Facebook have initiated the rise of surveillance capitalism. These platforms commodify "reality" by tracking the behaviors of individuals online and offline, making predictions about how they may act in the future, and constructing mechanisms to influence these future behaviors, whether such behaviors are voting or making purchases.⁵²

In this new digital advertising model, internet platforms might not sell access to user data. But they do sell the attention of these consumers to brands and

companies that are willing to pay for it.⁵³ Additionally, these platforms monetize data by using it to facilitate ad targeting. These practices are extensive and invasive. Such pervasive online and offline surveillance, and the subsequent monetization of users' behaviors and ideas treats the "private human experience as raw material for product development and market exchange."⁵⁴ This model incentivizes internet platforms to collect as much data on users as possible, as this will enable them to offer precise targeting and delivery tools to advertisers.⁵⁵ In addition, the digital advertising ecosystem incentivizes rampant data collection as these systems require vast datasets in order to operate and improve.

Despite the fact that data-driven ad targeting and delivery practices have become an integral component of the business models of many internet platforms, there is still little transparency around the mechanisms and policies that these internet platforms use to carry out these practices.⁵⁶ Many users also lack awareness of how internet platforms target and deliver advertisements to them. Additionally, few (if any) internet platforms offer their users a comprehensive set of mechanisms with which individuals can control how they fit into the online ad ecosystem.

Nonetheless, a growing body of research has emerged documenting the concerning and often discriminatory results of algorithmic decision-making in the digital advertising ecosystem. The precise nature of modern ad targeting and delivery systems has enabled advertisers to specify which categories of users they would like to include in their ad campaigns. While these determinations might result in some users receiving ads that are relevant to them, this can also result in the discriminatory exclusion of certain categories of users. This can happen even if an advertiser sets apparently non-discriminatory parameters regarding which audiences they would like to target with an ad. This is because automated tools used to target and deliver ads at scale make inferences based on engagement metrics and data on what categories of users are more likely to engage with an ad. This can, and has been found to, reinforce certain societal biases regarding race, gender, and socioeconomic status⁵⁷ through instances such as price discrimination in ads on online retail sites and gender and race discrimination in ads on job and employment sites.⁵⁸ For example, an ad delivery algorithm might only deliver ads for traditionally-male dominated jobs, such as doctors or engineers, to male job-seekers—as it bases its optimization strategy on current and historical data on job occupants that reflect societal gender discrimination in these career fields. Automated tools can therefore prevent some categories of users from receiving critical and relevant information based on false inferences about who will be interested in an ad, which largely reflects societal prejudices and biases.

In addition, political advertisements on internet platforms have come under increased scrutiny. This was in part triggered by revelations that Russian operatives had used political advertising services on a number of internet platforms to influence and suppress voting in the 2016 U.S. presidential election.

⁵⁹ Since then, these platforms have come under increased scrutiny and pressure to provide greater transparency and accountability around their advertising operations, and to develop clearer policies and processes governing who can purchase and run political advertising campaigns.⁶⁰ Some platforms, such as Pinterest and most recently Twitter, have opted to ban political advertising on their platforms altogether.⁶¹ However, because political advertisements cannot be easily categorized or defined, and given that some of these decisions were made fairly recently, the impact of these new bans is yet to be seen.

Policymakers are also considering what approaches they can take to prevent election interference in the future, and to promote greater transparency around political advertising practices online. One example of such an effort is the introduction of the Honest Ads Act, which aims to regulate political advertisements online. However, some commenters have questioned whether such legislative approaches are consistent with the First Amendment, and whether this is a space in which the government can legally take action.⁶²

Case Study: Google



General Overview of Google's Advertising Platform

Google is the world's largest search engine with approximately 92.19 percent of the global search engine market.⁶³ The platform currently ranks first for global internet engagement on Alexa rankings,⁶⁴ meaning that it is the single most visited website in the world.⁶⁵ Google is also the largest player in the digital advertising market, with a 37.2 percent market share.⁶⁶ Advertising funds the majority of Google's, and its parent company Alphabet's, operations, with about 85 percent of Alphabet's annual revenue coming from ads.⁶⁷

Although Google was not the first digital advertising platform, it is responsible for establishing targeted advertising as a key industry practice.⁶⁸ In doing so, the platform reaped massive financial rewards and laid the groundwork for surveillance capitalism to become an established practice throughout the technology industry.⁶⁹ In the late 1990's, Google engineers recognized the vast potential of the continuous flows of behavioral data that could be inferred from the company's expansive data logs from search queries. Google initially used this data to refine and improve its product offerings. In late 2000, Google also harnessed this data to begin matching digital ads with queries, with the aim of delivering relevant ads to users and ensuring value for advertisers.⁷⁰ The insights provided by this behavioral data were paramount, and Google eventually also used this data to predict user behavior, which further spawned the growth of the company's ad business.⁷¹

Like many other online platforms and digital advertising platforms, Google collects and monetizes a staggering amount of user data, including personal information on users such as their age and location, which is used to target and deliver ads to them across the internet. It collects this information based on factors such as what a user searches for using its Search product, the videos they watch on YouTube, and the places they look for on Google Maps.⁷² The company also collects sensitive information about individuals even when they are logged out of their Google accounts, browsing in a private browsing mode, or not even Google account owners in the first place. All of the information that Google collects on individuals online is used to target and deliver ads to them across the internet.

In July 2018, Google rebranded and reshuffled its ad products so that they focused on three core offerings: Google Ads, Google Marketing Platform, and Google Ad Manager. Google Ads is a rebranded version of AdWords, a search-based paid advertising system. The Google Marketing Platform merges Google Analytics 360 Suite and DoubleClick Digital Marketing. Google Ad Manager is a programmatic advertising platform that merges DoubleClick for Publishers and DoubleClick Ad Exchange. These will be discussed in more depth in the next section.⁷³

A Technical Overview of Google's Advertising Platform

The primary mechanism that Google uses for targeting ads is DoubleClick, an advertising delivery and tracking company that Google acquired in 2007. As previously mentioned, portions of DoubleClick were merged into the new Google Marketing Platform and Google Ad Manager in June 2018. Therefore, mentions of DoubleClick refer to its technology, as the actual brand has been distributed throughout Google's platform.

DoubleClick relies on web cookies to track and collect information on users' browsing behaviors. It primarily does this by monitoring user IP addresses. This information can be used to infer a great deal about users' behaviors, habits, and interests, which is then used by DoubleClick to target and deliver ads to users.⁷⁴ However, DoubleClick does not have access to a user's real identity. Information such as users' names, email addresses, and search histories lies with Google proper.⁷⁵ When a user visits a web page that hosts advertising content through Google's Display Network or its DoubleClick Ad Exchange (now part of Google Ads and Google Ad Manager respectively), DoubleClick is able to identify the users' digital identity and determine which ads the user should see.⁷⁶ This is possible even if the user does not arrive at this web page through a Google product or service such as Google Chrome or Google Search.

The primary mechanism through which businesses can deliver and place ads on Google is through the Google Ads platform. Google AdWords was launched in

2000 and rebranded as Google Ads in 2018. Google Ads enables advertisers to create and deliver mobile and desktop ads. These ads can appear on Google's search engine results page (SERP), on other Google products and platforms (e.g. YouTube and Blogger), and on websites that opt in to having Google-placed ads on their web pages.

Using Google Ads, an advertiser can run three types of campaigns—search ad campaigns, display ad campaigns, and video ad campaigns. Search ads are text-based ads that appear alongside search results on the SERP. Display ads are conventionally image-based ads and they appear on websites that opt in to having Google place text- and image-based ads on their web pages. These websites opt in to this process by becoming part of an ecosystem of such websites, known as the Google Display Network (GDN). A website that is interested in hosting online ads must install software onto their website that relays information about a current user, such as their search criteria, to Google. Google will then place ads it deems relevant on their website. Through this process, these host websites receive a cut of the money from the ads. When an ad placed by Google appears on a website other than Google.com, it will be labeled with the banner “Ads by Google.”⁷⁷ The third category of ads, video ads, typically last between six and 15 seconds and appear on YouTube before, during, or after a video plays.⁷⁸

There is no upfront cost for advertisers who want to run ads using the Google Ads platform.⁷⁹ The profit model of Google Ads is based on three types of bids. The first is cost-per-click (CPC), in which an advertiser pays for every click their ad receives. The second is cost-per-mile (CPM) in which an advertiser pays for 1000 ad impressions. The third type of bid is cost-per-engagement (CPE) in which an advertiser pays a certain amount when a user demonstrates a predicted action with an advertiser's ad. When a user clicks or engages with an ad on a website that is part of the GDN, Google and the website owner split the profits.⁸⁰ Websites that are part of the GDN can be paid per-click or per-impression for the ads that appear on their web pages.⁸¹

When an advertiser wants to place a bid for an ad campaign, they can do so using an automated or manual bidding process. Manual CPC bidding lets advertisers set their own maximum cost per click for their ads. Advertisers can set a default bid amount for their entire ad group. They can also set separate bid amounts for individual ad keywords or placements, which they may have identified as more profitable. As an ad campaign runs, Google Ads will help advertisers evaluate and optimize their performance.

Advertisers who want even more detailed information can use Google Analytics.⁸² Google Analytics is a Google product that can track and report on website traffic to an ad.⁸³ It also feeds data back to Google, which can use the data to enhance its targeting practices and parameters.⁸⁴ According to the 2016 Englehardt and Narayanan study, Google Analytics is present on almost 70

percent of websites.⁸⁵ During automated bidding, an advertiser can choose from a number of automated bidding strategies, such as maximize clicks, maximize conversions, and maximize conversion value. Conversion is when a user engages with an ad and then performs an action an advertiser deems valuable such as making a purchase or communicating with the business.⁸⁶

During the automated bidding process, an advertiser will set a maximum budget that they are willing to spend. Google will take into account this maximum budget and adjust an advertiser's bid based on competing bids in order to optimize the advertiser's selected bid strategy.

A bid is not the only factor that influences who wins an ad auction, however. Google determines how to rank and place ads by calculating an Ad Rank value. The higher an advertiser's Ad Rank value, the higher its ads will rank on the list of ads to place.⁸⁷ Ad Rank is calculated by multiplying the maximum bid placed by an advertiser with an advertiser's Quality Score (QS). The QS is an estimate of how relevant an advertiser's ads, keywords, and landing page are to a user who sees an advertiser's ad. A higher QS typically results in lower costs for an advertiser and a better ad position.⁸⁸ This is because a user who engages with relevant content is more likely to continue using Google's services. As a result, Google is not only interested in the price advertisers pay, but it is also in Google's business interests to ensure relevant ads are displayed to users.⁸⁹ The QS is calculated by measuring the quality of an ad using a range of factors. These include:

1. **The Landing Page Experience:** An estimate of how relevant and useful an advertiser's landing page is to users who click on their ad. This metric accounts for how well the content on an advertiser's landing page matches a user's search query and how easy it is for users to navigate the landing page.
2. **Expected Clickthrough Rate (CTR):** A measure of how likely it is that an advertiser's ad will be clicked on when shown. This score is based on the past clickthrough performance of an advertiser's ads. It does not consider factors that could affect the visibility of an ad, such as ad position, and extensions.
3. **Ad Relevance Measure:** A measure of how closely an advertiser's keyword matches the message in their ads. A below average ad relevance score can indicate that advertiser's ads are too general or too specific to answer a user's query, or that the keyword is not relevant to the advertiser's business.⁹⁰ The higher an ad's QS, the higher it will be ranked and placed online.⁹¹

Advertisers can also opt to use responsive search ads. This is when an advertiser provides Google with a set of 15 headline texts and four variations of ad copy. Over time, using its analytics and insights, Google Ads will automatically test different combinations of these headlines and ad copy⁹² using A/B testing. It will then identify, select, and run which of these combinations performs the best.⁹³

Google AdSense is the primary mechanism that online publishers such as bloggers or website owners can use to host ads on Google. The AdSense platform is distinct from the Google Ads platform and was launched in 2003. It enables contextual advertising by letting online publishers opt in to allowing Google to place ads on their websites, which are targeted to the publishers' content or audience.⁹⁴ When a user enters search criteria into a website that is using AdSense, AdSense will place ads that it deems relevant to the user's search on the search results web page. If an advertiser, also known as a "sponsor," is interested in placing an ad online using AdSense, they provide Google with search criteria and potential ads that Google can deliver if a relevant match occurs. Sponsors will also provide Google with a bid they are willing to pay if a reader ends up clicking on an ad that was delivered to them.⁹⁵

In order to determine which ads should be displayed to a user during a search query, Google runs an automatic bid auction in real-time for the same search criteria. Typically, the ad with the highest bid will be displayed first, and the remaining ads will be ranked sequentially. If a bid is too low or if displaying an ad will result in passing a threshold related to an advertiser (e.g. exceeding the maximum account total for the sponsor), then Google may not show the ad.⁹⁶ AdSense automatically places targeted ads, which are created and paid for by advertisers, on a participating publisher's website.⁹⁷ Google pays these web publishers based on the number of clicks or impressions ads receive.⁹⁸ Since advertisers pay different prices for different ads, the amount that a web publisher can earn through using AdSense varies.⁹⁹

According to a 2016 study by Englehardt and Narayanan, DoubleClick is present on almost 50 percent of websites¹⁰⁰ thanks to "cookie syncing," which enables different trackers to share user identifiers with each other. This makes DoubleClick one of the most promiscuous cookie-syncing third parties, sharing 108 different cookies with 118 other third parties.¹⁰¹

Controversies Related to Google's Advertising Platform

Google's ad targeting and delivery practices have been scrutinized for their ability to produce harmful results. Many researchers have documented how the platform's targeting parameters can enable advertisers to generate discriminatory outcomes. As previously noted, the process of establishing targeting parameters involves discriminating or differentiating among categories

of users as it requires advertisers to include and exclude users based on their personal traits, behaviors, and preferences. While not all types of discrimination against users is problematic, in cases where the discrimination is based on protected classes or other sensitive categories, this can result in harmful and even illegal outcomes that negatively affect protected classes and exacerbate hidden societal biases.

According to a Google spokesperson, Google’s ad targeting policies prohibit targeting based on traits such as race, ethnicity, religion, disability, and negative financial standing. Despite this formal policy, however, advertisers are able to target ads based on related parameters such as address, which can be a proxy for race due to segregated housing patterns.¹⁰² In addition, as described below, even beyond deliberate use of proxies, targeted ads may lead to discriminatory patterns by detecting and perpetuating hidden biases in underlying data.

For example, in 2017 *BuzzFeed News* revealed that Google’s ad platform permitted advertisers to target users using discriminatory terms. In addition, the investigation found that Google’s ad platform also suggested discriminatory keywords in response to such harmful terms. These suggestions were generated by automated tools. For example, entering “white people ruin” as a potential advertising keyword returned automated suggestions to run ads next to searches including “black people ruin everything.” Similarly, entering “why do Jews ruin everything” returned automated suggestions to run ads next to searches including “the evil jew” and “jewish control of banks.”

Following an inquiry to the company about the keywords *Buzzfeed* used in its investigation, Google disabled all of the keywords except “blacks destroy everything.” The company stated that just because a keyword is eligible, it does not mean an ad campaign will necessarily be run using it. In addition, the company shared that it has language that informs advertisers when their ads are offensive and subsequently rejected. However, given they failed to catch these particular keywords, this investigation revealed how the use of automated tools during the ad targeting phase can result in the exacerbation of societal biases and discriminatory outcomes.¹⁰³

Numerous researchers have also identified instances of discriminatory outcomes during the ad delivery phase while using Google’s ad platform. A 2013 study on discrimination in online ad delivery found that when the researcher conducted a Google search for a person’s name, Google AdSense would return personalized ads for public records on the SERP. The study found that when the search terms included names that were commonly assigned to African American children, the resulting public record ads suggested that the individual had an arrest record 81 percent to 86 percent of the time. When the search terms included names that were primarily assigned to white children, the word “arrest” typically did not appear in the text of the public records ads that were delivered.

This study provided a dramatic example of how the algorithmic decision-making of Google AdSense's delivery system detected and perpetuated societal biases. Although historical bias in the American criminal justice system has resulted in disproportionate arrest rates for African Americans, the AdSense algorithm was also detecting and perpetuating the societal notion that African American names would be associated with an arrest record in public records.¹⁰⁴ In fact, however, in many of these cases, although the ads were suggestive of an arrest record, after actually clicking through, a user would find that these individuals had no actual arrest record.

In addition, a 2015 study found that when users self-identified their gender in their ads settings, Google was less likely to show female job seekers ads for coaching services for high paying jobs, compared to male users. In order to test this, the researchers built an automated system that mimicked 17,370 jobseekers, and indicated the gender of each fake user in the ad settings.

The profiles were then shown 600,000 ads that the researchers tracked and evaluated. The study highlighted how Google's ad delivery system made inferences about what types of ads were relevant to a user based on the user's identified gender. This raised significant concerns, as these results reflected existing societal biases and highlighted how such tools could be used to further exacerbate such biases, by for example, furthering the gender pay gap.¹⁰⁵

In some cases, AdSense has also raised issues when advertisers have found their content placed next to controversial or harmful content such as hate speech. This often results in public backlash as well as advertisers pulling their ads from Google's ad platform.¹⁰⁶ It has also highlighted the need for brand safety, and advertisers are increasingly looking to advertising platforms to ensure that their ads are placed next to brand-safe content.

Given the range of controversies Google has faced related to its ad practices, the company has come under significant pressure to provide greater transparency and accountability around its ad operations. The company has taken some steps to remedy these issues. However, given the vast reach and influence of its tools, and because often times discriminatory outcomes can result from societal biases that are hidden in the data that Google's algorithms are processing, efforts to date have been insufficient to address these serious and complex problems.

Google's ad targeting and delivery practices have also raised numerous privacy concerns. In the last several years, Google has varied its practices regarding what data it collects and combines for purposes of ad targeting. Prior to 2016, Google explicitly stated that unless users opted-in, the company would not combine users' personal information collected from Google products such as Gmail and YouTube with the information it collected on users' browsing activity, such as IP address and search queries,¹⁰⁷ using DoubleClick.¹⁰⁸ Despite this, the company introduced this practice in June 2016. Google could now deliver personalized ads

to users based on data such as the keywords¹⁰⁹ they used in Gmail or receipts they had in their Gmail account.¹¹⁰ This change raised a number of privacy concerns and signaled a significant increase in the company's invasive data collection practices.¹¹¹ In June, 2017, the company announced that it would end this practice before the end of that year, and Gmail content would not be used or scanned for any ads personalization.¹¹²

Google also buys data on consumers' credit card purchases and other offline activity in order to enhance its ad targeting and delivery offerings.¹¹³ The company asserts that by harnessing all of these data points, it will be able to deliver users with ads that are more useful and relevant to their interests.

Steps Google Has Taken to Adjust Its Advertising Platform Since the 2016 U.S. Presidential Election

Following the controversy around the use of political advertising to influence and suppress voters during the 2016 U.S. presidential election, Google created a database of paid political advertising from its ad platform in its transparency report. The database includes information on political ads in the United States, the European Union, and India. Information in the database includes the total ads run and the total ad spend since May 31, 2018, ad spend by geography (by state or country), the top advertisers, and the top keywords. This data is updated weekly and is available for download in a CSV (comma separated values) file. In addition, each ad in the database is accompanied by information on the ad's format (e.g. video), the amount spent on the ad, a range of the impressions the ad accrued, and a date range that the ad ran in.

Although these data points are a valuable starting point, the database does not include granular information on how and to whom these ads were targeted and the reach and engagement they had. For example, Google currently provides a range figure of the impressions an ad received, but this does not include detailed engagement information such as the number of likes, shares, and video views the ad received. In addition, the database does not include information on the size and characteristics of the user categories that the ad was delivered to. Transparency around this information is important because it is often required for researchers in order to identify instances of discrimination.

Further, Google's database and transparency report only cover three regions—the United States, the European Union, and India. In December 2019, the company expanded its election advertising transparency in the United States to include state-level candidates and officeholders, ballot measures, and ads that mention federal or state political parties.¹¹⁴ However, these transparency efforts are limited to certain markets, and are not applicable to all of Google's users and ad operations around the world, despite the fact that these services are used extensively globally.¹¹⁵

As a result, this database falls short in its attempts to provide meaningful transparency and accountability around Google’s political advertising operations to all of its users.

Additionally, Google does not provide a comprehensive database of other types of ads outside political ads on its platforms, even though they may also yield discriminatory outcomes. In March 2019, the Department of Housing and Urban Development (HUD) charged Facebook for discriminatory ad practices related to housing. Shortly thereafter, HUD announced it would also be investigating Google and Twitter’s ad practices, suggesting that discriminatory advertising around housing may also be a concern on Google’s ad platform. Investing resources into expanding this ad transparency database to include all categories of ads would therefore be in the best interest of both the public and the company.

In addition, Google’s transparency report does not provide any data on how the platform enforces its ad content and targeting policies, and how many ads and accounts are subsequently removed or suspended. This demonstrates a significant gap in transparency and accountability around how the platform’s ad content and targeting policies are enforced. In December 2019, France’s competition watchdog, l’Autorité de la Concurrence, fined the company €150 million for using and randomly enforcing “opaque and difficult to understand” operating rules on its ad platform. The authority has ordered the company to provide further clarity around how it creates its Google Ads rules and how it enforces these rules to suspend accounts.¹¹⁶

Google has taken some limited steps to address concerns regarding transparency and accountability around their ad targeting practices. In November 2019, the company announced that it would limit targeting options for verified advertisers seeking to run election ads so that these advertisers can only target users based on geographic location (down to the postal code level), age, and gender. Advertisers will be able to continue targeting users with contextual targeting (e.g. delivering ads to users reading or watching a story about the economy). The policy change applies to ads on Google Search, YouTube, and through the GDN. Google initially rolled this policy out in the United Kingdom ahead of the General Election in December 2019, and then expanded it to the European Union at the end of 2019, and globally in the beginning of 2020.¹¹⁷

As the debate around appropriate rules for political ads online continues, many civil society organizations and researchers are calling for internet platforms to ban political ads altogether. In March 2019, Google banned political ads in Canada ahead of the country’s federal elections. This followed Canadian lawmakers’ enactment of a bill in December 2018 that will go into effect in June 2020 and will require online platforms to keep a record of all political and partisan ads they publish directly or indirectly during an election season. According to Google, complying with this law would necessitate major alterations to its advertising system, as it deploys a real-time bidding procedure.

Given these technical barriers and the fact that both political and partisan ads are hard to define and thus hard to monitor, Google decided to ban political ads in the country rather than risk non-compliance, which could result in fines or jail time.¹¹⁸

Other prominent voices, such as Ellen L. Weintraub, the chair of the Federal Election Commission, have suggested that instead of banning political ads altogether, platforms should prohibit the practice of microtargeting political ads. Microtargeting is a marketing and advertising strategy that harnesses user data, such as their interests, online connections, purchase history and demographic data, in order to segment users into smaller, more precise groups for ad and content targeting.¹¹⁹ According to Weintraub, the practice of political ad microtargeting, enables advertisers to identify vulnerable groups and deliver political misinformation to them with little accountability, as these ads are never seen by the larger public.¹²⁰ Numerous researchers, technologists, and civil society and advocacy organizations have joined this call for a prohibition on political ad microtargeting.¹²¹ Others, however, have pushed back, stating that limiting microtargeting prevents grassroots candidates from establishing themselves and puts them at a disadvantage.¹²²

Further, following controversy over Facebook's decision to exempt political ads from its fact-checking policy in October 2019, Google clarified that it prohibits making false claims in political ads.¹²³ However, some political leaders in the United States have criticized the company for limiting this ban only to false claims that could result in voter suppression or threaten election integrity, and not applying it to cover false claims regarding specific candidates.¹²⁴

User Controls on Google's Advertising Platforms

Google has extended some limited features to its users in order to give them greater control over their ad experience. For example, Google users can adjust certain ad preferences and settings through the Google Accounts page. This includes being able to turn ad personalization off. It also includes being able to stop the platform from using activity and information from Google services to personalize ads. This applies to ads on websites and apps that partner with Google to show ads. When enabled, this feature stores data from these partner websites and apps in a user's Google Account. This page also includes information on how a user's ads were customized, and it explains that ads are customized based on personal information in a user's Google Account, data from advertisers that partner with Google, and Google's inferences about a user's interests. This last category includes a list of factors, such as apparel or combat sports, that Google's ad system believes are relevant to the user and thus considers when targeting and delivering ads to them. Users have the option to remove any of these factors from the list of active factors considered by Google's

systems. This page also links users to an install page for a browser plugin that enables users to maintain their preferences for opting out of personalized ads from Google. The plugin enables users to retain these preferences even after they have cleared their cookies. However, these features and controls are only available to Google users who are logged in, leaving non-logged in users or users without a Google account with fewer controls.

According to Google, there were 2.5 billion visits to the Google Accounts page in 2018, but only 20 million people per month visited the ad settings page.¹²⁵ This raises concerns that although Google offers a limited set of controls to users, users may be either unaware of these options or the settings may be hard to access or understand.

Users and non-users can also adjust ad settings using AdChoices, where they can control ads from other ad networks. This tool provides a list of all of the participating companies including which of them run personalized ads for a user's browser. The tool also enables users to opt out of data collection and use from participating companies. These participating companies abide by the Digital Advertising Alliance's (DAA) self-regulatory principles¹²⁶

Additionally, in May 2019, following decisions by browser companies such as Mozilla and BRAVE to restrict the use of cookies, Google implemented similar restrictions on Google Chrome, noting that it wanted to provide users with greater controls over how these cookies were deployed to track their behaviors. The changes included requiring website developers to flag when cookies can be used for cross-site tracking purposes. The changes also allowed users to clear cross-site tracking cookies, while enabling single domain cookies to continue operating, thus maintaining user logins and settings.¹²⁷ Also in May, Google introduced an auto delete feature that complemented the existing manual deletion controls. The new feature enables users to set up automatic deletion of their location, app, and web history. Users are able to choose a period of time—either three months or 18 months—that Google can save their data. Data that falls outside this period is deleted automatically. Users also continue to have the option to pause Google's web and app activity tracking.¹²⁸

Further, Google introduced a browser extension that users can install to receive more information about the ads they see across Google products and services and from its ad network partners. This information includes the names of third-party companies that helped with the ad targeting process, as well as the names of companies with trackers embedded in ads.¹²⁹

Despite Google's efforts to promote greater transparency and accountability around its online ad practices, the company falls short. As described, it offers users a limited set of controls over their online ad experience, but these are difficult to understand and use. In addition, non-users or logged out users do not have access to these controls and are therefore unable to manage and opt out of

receiving personalized and targeted ads. Further, although Google publishes a political advertising database and transparency report, this report fails to provide granular information on the targeting parameters of these ads as well as related engagement metrics. This report also does not provide this information on an adequate global scale, despite the fact that Google has a global user base. This report also does not provide a comprehensive overview of total advertising on Google's ad platform, as it only covers political ads. Given that the platform is one of the largest digital advertising players in the industry, it needs to do more to provide adequate and meaningful transparency and accountability around its ad targeting and delivery operations to its users and the public.

Case Study: Facebook



Source: PixieMe / Shutterstock

A General Overview of Facebook's Advertising Platform

Facebook is the world's largest social media platform, with approximately 1.59 billion daily active users as of June 2019.¹³⁰ The platform ranks third in global internet engagement on Alexa rankings.¹³¹ Facebook is the second largest player in the United States and global digital advertising industries, after Google. Projections from earlier in 2019 estimated that the platform would have approximately 22 percent of the U.S. market share¹³² and generate \$67.37 billion in net digital ad revenues in 2019.¹³³ Advertising accounts for the majority of the company's revenue. In 2018, the company generated \$55 billion in ad revenues, a significant increase from 2017, when it generated almost \$40 billion in ad revenues.¹³⁴

One of the primary reasons that Facebook has been able to successfully grow its advertising operations is because it collects a vast amount of data on its users.¹³⁵ Facebook collects this data using social plug-ins and cookies that track Facebook users, regardless of whether they are logged into the platform or not.¹³⁶ According to a 2016 report by Steven Englehardt and Arvind Narayanan, Facebook has placed trackers on approximately 30 percent of the top 10,000 sites listed on Alexa ranking's top site list and 32 percent of the top 500 Android apps.¹³⁷

Facebook’s expansive advertising operations have yielded concerning results, however. In early 2018, it was revealed that British political consulting firm Cambridge Analytica had acquired the personal data of millions of Facebook users without their consent and used it for political advertising. Following the scandal, Facebook imposed some limits on data sharing with outside third parties. For example, in March 2018, the company asserted that it would shut down its Partner Categories program. This program had enabled advertisers to use data from third-party data brokers like Axiom for ad campaign targeting. In exchange for this data, data brokers would receive a portion of the profit of an advertising sale.¹³⁸ News outlets speculated that shutting down the Partner Categories program was in direct response to Cambridge Analytica’s mishandling of user data, though the explicit link was never confirmed by Facebook.¹³⁹

A Technical Overview of Facebook’s Advertising Platform

The data Facebook collects on users includes granular data on factors such as user likes, engagements, interests, and behaviors. In some cases, users can also specify their own interests.¹⁴⁰ Given that this data is detailed, its use can result in concerning ad targeting and delivery practices on the platform, including user manipulation. For example, in 2018, the *Intercept* acquired a confidential internal document from Facebook that outlined an internal advertising service that enables advertisers to target users based on predictions related to their behaviors, thoughts, and purchases. This artificial intelligence-powered prediction tool was first publicly revealed in 2016, and is known as “FBLearner Flow.” The document outlined how the platform could, for example, identify users who were “at risk” of switching to a competing brand, and then aggressively target them with advertising that could prevent this shift. In addition, in 2017, the Guardian acquired a leaked internal presentation for advertisers that claimed the company could track when teens felt “insecure,” “worthless,” and “stressed,” among other things. This information could be collected by monitoring posts, interactions, and photos, and could be used by advertisers to deliver ads to these users during “moments when young people need a confidence boost.”¹⁴¹

Although Facebook asserts it does not sell or widely share the sensitive and precise data that it collects with outside parties, the company is able to monetize the data for advertising purposes. In particular, this data enables advertisers who use the Facebook advertising platform to target users based on very specific characteristics. These characteristics can be broken down into three broad categories:

1. **Demographics and attributes:** Like many other advertising platforms, Facebook lets advertisers select and target audiences based on demographic factors. These include age, gender, location, profile

information, activity on the Facebook platform, and data from third parties. Advertisers can also target users based on specific interests, which are inferred based on signals such as the Pages a user likes and engages with and the ads that a user clicks on. They can also target users based on behaviors such as how they connect to Facebook (e.g. through a browser or mobile device).¹⁴² Research suggests that Facebook offers over 1,000 well-defined attributes and hundreds of thousands of free-form attributes to advertisers.¹⁴³

2. **Personal information:** Advertisers can also specify particular individual users that they would like to target by uploading information on these users to Facebook along with some of their personal information (e.g. name, address, date of birth). Once this information has been uploaded, Facebook can use it to identify the corresponding users on the Facebook platform, and target them with ads. Advertisers can then also target users based on various attributes such as their political affiliation or whether they visited a company's website.¹⁴⁴ Audiences that are created using either of these mechanisms are known as Custom Audiences on the Facebook platform.¹⁴⁵
3. **Similar users:** In 2013, Facebook introduced the Lookalike Audiences feature. This feature lets advertisers target users who demonstrate similar characteristics to users they have previously chosen to target during advertisement campaigns. The Lookalike Audiences tool identifies similarities between the users in an initial Custom Audience list that the advertiser has constructed. It then creates a new audience list that includes other users who share similar qualities. The Lookalike Audiences feature offers advertisers the opportunity to create granular and segmented target audience lists. For example, advertisers can upload a list of users' personal information in the original Custom Audience list and also use demographic and attribute factors to more precisely identify and target an audience.¹⁴⁶

Advertisers can buy ads using a range of tools on Facebook, including the Ad Manager tool, and Pages.¹⁴⁷ When an advertiser wants to run an advertising campaign on Facebook's advertising platform, they are presented with a number of objectives (e.g. ad impressions or engagements) to choose from. Each of these objectives aims to optimize a campaign at different stages of the marketing funnel¹⁴⁸ and they are known as "optimization events." The optimization events available to advertisers include "awareness," in which a campaign is optimized for the most impressions or views; "consideration," in which a campaign is optimized for the most clicks or engagements; and "conversion," in which a campaign is optimized to generate a valuable outcome for an advertiser, such as sales.¹⁴⁹ Once an advertiser selects an optimization event, they bid on the

objective itself. This means that if an advertiser selected the “consideration” option, they would bid on the number of ad clicks or engagements. These bids and objectives are considered by an advertising platform if and when an advertiser’s ad campaign is delivered.¹⁵⁰

A bid can take many formats, but it typically includes information such as the start and end period of an advertising campaign and an overall or daily budget cap for the campaign. Advertisers can either manually create a custom bid or opt for automatic bidding. When Facebook has ad slots available, it will run an automated ad auction among the active ads that are bidding for a specific user and the limited space on their newsfeed and pages they view. Using the budget information provided by an advertiser, Facebook will place bids in ad auctions for the advertiser. Advertisers also have the option to opt for a per-bid cap.¹⁵¹ This limits the amount that Facebook’s automated system will bid on their behalf during automatic bidding for a single optimization event during the ad auction. However, ad placement is not solely based on the bids placed by an advertiser and how the ad fares in the auction process.¹⁵²

The ad that eventually wins an auction and is shown to a user is the one that represents the highest total value. This value is derived based on more than just how much an advertiser is willing to bid to show their ad. It can be characterized by the following formula:¹⁵³

$$\text{Total value} = \text{advertiser bid} \times \text{estimated action rate} + \text{user value}$$

The platform also calculates an advertiser value score, using the following formula:

$$\text{Advertiser value} = \text{advertiser bid} \times \text{estimated action rate}$$

In these formulas, the estimated action rate and user value are both calculated using machine learning models. The estimated action rate is the likelihood that a user will perform the action the advertiser set as their objective (e.g. for conversion). This is calculated by an algorithm that considers a user’s behavior on Facebook, such as what types of ads and pages they click on and like. The algorithm also considers a user’s off-Facebook activity, which is often provided by advertisers through Facebook business tools such as Pixel and SDK. The estimated action rate is also influenced by the content of the ad itself. If one user performs the desired action when engaging with an ad, the algorithm will use pattern recognition to identify commonalities between the user and other users, and reinforce the model. In this way, the automated tools are refined in almost real-time based on the results of an auction.¹⁵⁴

The user value is the machine learning model's prediction of the quality of an ad, as opposed to the likelihood that a user will perform a certain action in relation to the ad. The quality of an ad is based on a range of data sources including feedback from users who view or hide ads, and Facebook's assessments of low-quality characteristics in an ad, such as too much text in an image, sensationalized language, and engagement bait.¹⁵⁵ According to the company, there is a significant amount of human review during this ad review process, as they seek to build models that are able to predict high-quality ads as accurately as possible.¹⁵⁶

In addition, factors such as ad relevance are considered when delivering ads. On Facebook, ad relevance can be measured by assessing an ad's quality and an ad's estimated action rate.¹⁵⁷ Each ad is assigned a relevance score between 1 and 10 on the platform, with 10 being the highest. Advertisers can view this score using Facebook's ad reporting tools. The relevance score is constantly updated and is based on the positive and negative feedback Facebook expects an ad to receive from its target audience. The more predicted positive interactions, such as video views, conversions, etc., the higher the ad's relevance score will be. The more predicted negative interactions, such as hiding or reporting an ad, the lower its relevance score will be. Ads that are guaranteed delivery, such as those bought through reach and frequency,¹⁵⁸ are not affected by relevance score calculations. Relevance score calculations also affect the cost and delivery of brand awareness campaigns to a lesser degree, as these ads are optimized to reach users, rather than drive a specific action like purchases or app downloads.¹⁵⁹ Generally, the higher an ad's relevance, the less it will cost to be delivered.¹⁶⁰

On the Facebook platform, users can see ads through paid or organic distribution. When an ad is delivered through paid distribution, it will be marked as "Sponsored." When an ad is delivered based on organic or non-paid distribution (e.g. because a user's friend shared the ad or because a user visited a company's Page), the ad is not marked as "Sponsored." An ad can continue to circulate on the Facebook platform through organic distribution even after its paid distribution budget has been depleted.¹⁶¹

Facebook offers its advertisers a range of insights and data on how well their advertising campaigns are performing, including a detailed interface and a dedicated application programming interface (API) for launching and monitoring ads. An API is a set of programming code that simplifies data transmission between two pieces of software. Both the interface and API provide semi-live updates on campaign delivery. This includes data on the number of impressions and optimization events occurring as the ad is delivered. Advertisers can request more granular data on their ads, including that performance insights be broken down based on characteristics such as age, gender, and location (known as a Designated Market Area or DMA region). Importantly, the interface and API do

not break down this data based on race, as this could lead to racial discrimination.¹⁶²

Although Facebook provides advertisers with a range of insights on their campaign's performance, the reliability of these metrics has come into question. For example, a study found that a large number of ads that brands were paying for were being viewed by "bots," or computer programs that imitate the behavior of internet users. In response, Facebook conceded that it had made measurement errors related to engagement metrics.¹⁶³ Although researchers were able to uncover this error, Facebook generally provides little transparency around how the company calculates such metrics. Without greater transparency around what they are measuring in the first place, it is difficult for the public to hold the company to account when mistakes like this happen.

Facebook has also created tools that help advertisers create and deliver ads. For example, in May 2019, Facebook launched the Automated Ads platform, a tool designed for small and medium businesses that may not have large advertising budgets and technical experience. The tool walks advertisers through a set of questions about their company and the goals of their advertising campaign. Based on this information and the details on the company's Facebook page, the tool recommends call-to-action buttons (e.g. shop now, book an appointment, donate now), advertising content such as text, and details of the ad creative based on the business' Facebook page. The tool also automatically generates up to six different versions of an ad using these recommendations that an advertiser can then run across Facebook and its other services: Instagram, Messenger, and Audience Network. Audience Network is an off-Facebook, in-app advertising network for mobile apps that enables advertisers to deliver ads to users of other mobile sites and apps. This tool offers advertisers with ad targeting recommendations and budget guidance based on their stated goals of the advertising campaign. If an advertiser opts to set their own budget for the campaign, the tool estimates results. Once an ad campaign is live, the Automated Ads platform optimizes the campaign by delivering the "best performing" ads going forward. Advertisers using this platform similarly receive reports and insights on how their campaign is performing and how they can improve results.

¹⁶⁴

Controversies Related to Facebook's Advertising Platform

As previously highlighted, targeted advertising is founded on discrimination, as it enables advertisers to delineate which specific categories of users they want to target during campaigns. Depending on the characteristics being used to distinguish among users, at times this discrimination can be harmless, such as distinguishing between users interested in baseball gear versus users interested in football gear. In addition, there can be legitimate instances of targeting users based on characteristics such as race and gender, such as, for example, targeting

students who attend Historically Black Colleges and Universities, or professional societies for women, to advertise events focused on benefitting these groups. However, this type of targeting can also be harmful, even though these harms are often difficult to detect.¹⁶⁵

The current structure of the Facebook advertising platform has enabled problematic and discriminatory outcomes to result from both the targeting and delivery phases. For example, in 2016, *ProPublica* revealed that they had been able to run housing-related ads on Facebook that explicitly prevented certain categories of users from receiving them, based on their “ethnic affinity.”¹⁶⁶ In November 2017, *ProPublica* further demonstrated that they had been able to buy numerous home-rental ads on Facebook and target the audience so that they specifically excluded “African Americans, mothers of high school kids, people interested in wheelchair ramps, Jews, expats from Argentina and Spanish speakers.”¹⁶⁷

In 2017, *ProPublica* also revealed that Facebook’s automated targeting category generator was generating antisemitic categories of users based on topics user’s had previously expressed interest in. These included phrases such as “How to burn Jews” or “Jew hater.” Although these categories were so small that the platform would not have enabled them to run an ad campaign on this basis, this case study demonstrated how automated tools can suggest categories of users that advertisers should target in ways that reflect societal biases and exacerbate overt, harmful biases.¹⁶⁸

In December 2017, *ProPublica* and the *New York Times* reported that Facebook also permitted numerous major brands and companies to run recruitment advertisements for specific age groups, including those ages 25-36. At the time of the controversy, Facebook pushed back against this criticism, asserting that age-based targeting in advertising was “an accepted industry practice”¹⁶⁹ as it enables employers to recruit candidates of all ages to find work.¹⁷⁰ However, under the Age Discrimination in Employment Act (ADEA), it is illegal to discriminate against people over the age of 40 in employment.¹⁷¹

Facebook has also been the subject of a series of formal complaints and lawsuits alleging that its advertising practices can result in illegal housing and employment discrimination. In March 2018, the National Fair Housing Alliance (NFHA) and three of its member organizations filed a lawsuit against Facebook, claiming that its advertising platform can allow housing providers to engage in illegal housing discrimination as it allows advertisers to segment which audiences see housing-related ads based on characteristics such as race, religion, sex, and disability.¹⁷² In August 2018, HUD filed a formal complaint against Facebook based on similar grounds.¹⁷³

In March 2019, Facebook agreed to pay approximately \$5 million to settle numerous lawsuits that claimed the company’s advertising platform enabled

discrimination in housing, employment, and credit ads. Shortly after the settlement, HUD formally charged Facebook for violating the Fair Housing Act. HUD Secretary Ben Carson released a statement noting that “using a computer to limit a person’s housing choices can be just as discriminatory as slamming a door in someone’s face.”¹⁷⁴ Similarly, in September 2018, the American Civil Liberties Union (ACLU), Outten & Golden LLP, and the Communications Workers of America filed a class action employment discrimination suit alleging that Facebook permitted employers to target ads based on categories including race, national origin, age, and gender.¹⁷⁵ Under the Fair Housing Act, as amended, which is part of the Civil Rights Act of 1968, individuals are protected from discrimination in housing based on an individual’s race, color, disability, religion, sex, familial status, or national origin.¹⁷⁶

In October 2019, 54-year-old Neuhtah Opiotennione filed a class-action lawsuit in a federal district court in San Francisco. The suit alleged that Facebook’s ad platform permitted financial services advertisers to target audiences based on their age and gender, thus preventing users who are female and over the age of 40 across the United States from accessing financial services products such as loans, insurance coverage, and bank accounts. The lawsuit alleges that this violates California’s Unruh Civil Rights Act, which states that all people in the state, regardless of sex, age, race, medical condition, or other characteristics, are entitled to “full and equal accommodations, advantages, facilities, privileges, or services in all business establishments of every kind whatsoever.” The suit also alleges that Facebook has taken no steps to prevent financial services advertisers from excluding users based on their age or gender, except in the case of credit opportunities.¹⁷⁷ Although we have not had the ability to assess the merits of these various lawsuits, the vast array of discrimination cases that have been brought against Facebook demonstrate that it is important to ensure that anti-discrimination statutes cover discrimination in the online ecosystem, including online advertising. This is an important area for policymakers and civil society to focus on in the future.¹⁷⁸

Whereas the ad targeting phase involves input from both the advertiser and the advertising platform, the ad delivery phase largely falls under the control of the advertising platform. On Facebook, the use of algorithmic decision-making in the ad delivery phase has especially yielded problematic and discriminatory outcomes. In an April 2019 study on how housing and employment ads were delivered to a group of target users on Facebook, researchers from universities and a digital rights organization found that Facebook’s ad delivery algorithm was skewed to deliver these ads to specific gender and racial groups. According to the researchers, when they ran five ads for jobs in the lumber industry and tried to deliver them to a large and inclusive audience, the ad delivery algorithm instead delivered it to over 90 percent male users and over 70 percent white users in aggregate. In addition, when the researchers ran five ads for janitor employment opportunities, the ad delivery algorithm delivered the ad to over 65 percent

female users and over 75 percent black users in aggregate. Similarly, when the researchers ran an advertising campaign related to houses for sale and targeted the campaign to a broad audience, they found that the ad delivery algorithm delivered the ads to 75 percent white users, while rental ads were delivered to a more balanced group of users.¹⁷⁹

These skews in ad delivery occurred because Facebook's ad delivery algorithm made inferences on what categories of users would be more likely to engage with the ads, based on data it had about engagements and impressions thus far. It made these skewed deliveries despite the fact that the advertisers had specified they wanted to target a large, inclusive group of users.¹⁸⁰ In fact, the researchers found that even when choosing to target an underrepresented or marginalized group, the ad delivery system would still not deliver those ads to this category of users, and in some cases, it would not deliver the ads at all. This is because Facebook's ad delivery system is geared to optimize ad delivery by delivering ads to users whom the algorithm has identified as the most likely to engage with the ads, regardless of the advertisers intentions and desires.¹⁸¹ What ads a user engages, or doesn't engage with, can reflect their own personal biases and experiences, including related to their race, age, and so on. When an algorithm identifies patterns in such engagements, it will seek to optimize ad campaigns based on these trends in order to drive further return on investment.

This means that in some cases, algorithms can perpetuate biases based on inferences made from historical data, which could suggest, for example, that women don't want to be lumberjacks based on the fact that there are few women in the profession. However, in doing so, these algorithms can exacerbate societal inequities and create barriers to equal outcomes and opportunities by excluding groups from having the chance to engage with such ads.¹⁸² This prevents potentially new patterns in engagement data from being created and as a result, the algorithm can further skew the data against the underrepresented group and make false interpretations on the basis of missing data.

The researchers also concluded that the skew in ad delivery occurred due to market and financial optimization effects. According to their findings, the budget of an advertiser and the creative of an ad (e.g. headline, image, etc.) to a specific gender or race of users also affected how an ad would be delivered.¹⁸³ When the researchers used ads that only included links to sites such as bodybuilding.com, which is typically targeted toward male users, or elle.com, which is typically targeted toward female users, the ad was delivered relatively equally to users of both genders. However, when the researchers added images of a weightlifting man or of makeup brushes, ad delivery was skewed so that the bodybuilding ad's distribution was 75 percent male and the distribution of the cosmetics ad was over 90 percent to females.¹⁸⁴ While this is a highly specific example, it showcases how the use of images can increase how automated tools skew ad targeting toward a specific user attribute.

Advertising systems work as designed, but that doesn't always generate expected results.. Platforms such as Facebook make money when people click on ads. But an individual's tendency to click on certain types of ads (and not others) often reflects deep-seated social inequities: the neighborhood they live in, where they went to school, how much money they have, etc. An ad system that is designed to maximize clicks and profits for Facebook can therefore reinforce these social inequities and so serve as a barrier to equal opportunity.¹⁸⁵

In response to many of these criticisms, Facebook has taken steps, though limited, toward providing more meaningful transparency and accountability around its ad targeting and delivery practices.

As previously mentioned, in March 2019, the company agreed to settle numerous lawsuits that claimed the company's advertising platform enabled discrimination in housing, employment, and credit ads. Under the terms of the settlement, Facebook will no longer allow landlords, employers, creditors, and similar advertisers to target or exclude users based on age, gender, race, zip code, religion, and other sensitive targeting characteristics.¹⁸⁶ The company also agreed to study the potential for algorithmic bias in ads¹⁸⁷ and to consult with the ACLU and the other plaintiffs in the March 2019 settlement before adding future targeting categories.¹⁸⁸

Further, after *ProPublica* revealed that Facebook's automated tools had recommended antisemitic targeting categories to certain advertisers, Facebook introduced "new guardrails" to prevent this in the future.¹⁸⁹ Prior to this settlement, Facebook's policy simply required that advertisers had to review and self-certify that they would comply with Facebook's non-discrimination policy for ads.¹⁹⁰ When advertisers use the Ads Manager tool, they see a continuous, non-dismissable prompt that reminds them about the platform's non-discrimination policy before they begin using targeting tools.¹⁹¹ Both the prompt and policy have been in place since before the settlement. In 2018, in response to claims that Facebook's ad targeting options enabled advertisers to discriminate against certain categories of users, Facebook removed thousands of targeting options from its advertising platform.¹⁹²

Following its settlement with the Fair Housing Groups and the ACLU's coalition, Facebook also introduced new features to attempt to remedy targeting-related discrimination issues in the future, including a separate advertising portal for housing, employment, and credit ads that offers less precise targeting options.¹⁹³ Limiting the targeting options for these categories of ads is intended to minimize the likelihood of harmful discrimination in these contexts. This new system requires advertisers to flag if their ad involves housing, employment, or credit opportunities. If it does, then the advertiser will be routed to this new system. According to representatives from the company, Facebook uses machine learning tools in the background to ensure that advertisers do not intentionally or accidentally bypass this system by failing to clearly delineate their ads as falling

under these categories. In addition, the company uses human review to ensure that these machine learning systems are improving, and are able to identify repeat offenders.¹⁹⁴

However, research has demonstrated that automated tools are not always effective at detecting certain forms of content. This is especially true in cases where the definitions for the categories of content are vague, or when determinations of which category a piece of content should fall under depend on subjective and complex forms of speech or text.¹⁹⁵ In addition, it is still unclear how Facebook plans to comprehensively enforce this policy.¹⁹⁶ According to representatives from the company, the platform offers an appeals process to advertisers who believe their ads were mislabeled as a housing, employment, or credit ad by the company's machine learning systems or human reviewers. However, little is known about the scope, scale, and impact of this appeals process.¹⁹⁷

Further, despite these changes, researchers have found that even if an advertiser uses the new special ads portal, the composition of audiences can still be skewed toward specific demographic groups, thus yielding potentially discriminatory outcomes. The researchers found that this was because Facebook's modified algorithm relies on proxy characteristics that correlate with factors such as age and gender. As a result, they found that removing a select number of protected features has little impact on the presence of overall bias.¹⁹⁸

Additionally, the Lookalike Audiences feature has revolutionized how advertisers can target users on the Facebook platform, and there is little transparency around how these audience lists are developed. Facebook has not disclosed any information related to which factors its algorithms use to determine which audiences should be in a Lookalike Audience. This is because Facebook does not have access to the characteristics of the initial audience that an advertiser has and wants to duplicate. In addition, the advertiser will not be given the identities of the users who end up on the Lookalike Audience list. However, independent experiments have revealed that Facebook's Lookalike Audience algorithm can replicate the racial affinities and political affiliations of the initial audience in the Lookalike audience.¹⁹⁹

Steps Facebook Has Taken to Adjust Its Advertising Platform Since the 2016 U.S. Presidential Elections

As highlighted, Facebook's advertising platform has accrued significant value over the past few years, enabling advertisers of all sizes to target, influence, and engage with users based on precise characteristics. It has also produced numerous controversial and harmful impacts, many of which became evident after the 2016 U.S. presidential elections, when it was revealed that Russian operatives had used the Facebook advertising platform to influence voting

behaviors among users including to discourage African-Americans from voting.

²⁰⁰ One of the reasons such manipulation was able to take place is the fundamental lack of transparency and accountability guidelines around Facebook's advertising platform. As a result, Russian operatives were able to run influential advertisements on the platform without proper accountability, and users engaged with the ads without knowing the identity or objectives of those behind the campaigns.

Since then, the platform has come under heavy scrutiny and pressure to provide greater transparency around its ad operations and to develop stricter regulations governing who can buy ads, especially political ads.²⁰¹ In response to these pressures, Facebook introduced a range of new advertising transparency mechanisms. These include an ad database that aims to house all of the issue, electoral, or political ads on the platform, dubbed the "Ad Library." However, this new ad transparency tool is not user friendly and as a result searching, sorting through, and evaluating data in the database is challenging.²⁰² In this way, the tool fails to serve as a meaningful and useful transparency and accountability mechanism. In January 2020, the company made some improvements to the Ad Library's usability, enabling search based on exact phrases, introducing new search filters and better grouping similar ads.²⁰³

In addition, the database provides limited insight into what ads are run on the platform. Currently, the Ad Library enables users to search for political ads based on a name, topic, or organization. When a user enters a search term, they can view all related ads and filter them based on factors such as whether they are active or inactive; the impressions ads have received within the last day, 7 days, 30 days, and 90 days; relevant Pages; any disclaimers; and on which Facebook platform or product the ad was run. When a user clicks on a specific ad, they can view information such as an approximate number of impressions the ad received, an approximate figure for how much money was spent on the ad, a percentage breakdown of the genders and ages of those who viewed the ad, and a map indicating where in the United States the ad was shown, including a breakdown by percentage. Users can also see when an advertiser's page was created, if its name has been changed, the primary country location for people who manage the Page, and recent and total figures on how much the advertiser has spent on ads about social issues, elections, or politics.²⁰⁴

In January 2020, the platform expanded the metrics offered in the Ad Library to include Potential Reach,²⁰⁵ which highlights the estimated target audience size Facebook calculated for an advertiser, based on the targeting and ad placement options they selected.²⁰⁶

Although Facebook has expanded the metrics reported on in the Ad Library, the Ad Library's effectiveness as a tool is limited as it does not disclose vital metadata that is necessary for understanding the larger ad ecosystem, as well as the impact of ads run on the platform. For example, the database does not provide granular

information on the reach and engagement of an ad (e.g., how many likes, shares, and video views an ad received),²⁰⁷ how an ad was targeted, and granular information on the size and attributes of an audience that an ad was delivered to. This information is integral to identifying practices such as discrimination in advertising, as often times discrimination can only be identified when comparing the content of an ad with its targeting parameters.²⁰⁸ The company asserted that it cannot share this information publicly due to privacy concerns.²⁰⁹ However, many researchers have outlined effective safeguards that should be implemented in order to ensure these disclosures are made responsibly.²¹⁰

Further, although the company stated it aims to create a comprehensive database of political ads run on the platform, it fails to recognize the immense challenges and complexities involved in this process. The definition of a political ad is not clear, and the boundaries to enforce these definitions are vague, often subjective, and differ in each country and region that the platform operates in. Neither algorithmic nor human decision-making are likely to be able to make these determinations reliably and consistently, and as a result many political ads will be left out of the database. In addition, political statements and beliefs can easily be circulated on the platform through non-paid methods that rely on organic engagement to accrue virality.²¹¹

Additionally, although a database of political ads is a theoretically positive step toward providing more transparency and accountability around Facebook's advertising operations, its focus on political ads means that the public does not have insight into the broader landscape of potentially problematic ads on the platform, such as the previously discussed employment and credit ads.

To address this, Facebook has taken steps toward expanding its Ad Library to include other categories of ads. For example, in December 2019, Facebook expanded its Ad Library to include U.S. run housing ads. The Ad Library enables users to search for these ads by advertisers and by location, regardless of whether the user is part of the target audience of the ad campaign.²¹² This means that users can affirmatively search for and see all housing opportunities, even if they were not part of the demographic group targeted by an advertiser.

However, the data offered on housing ads is far less granular than the data offered on issue, electoral, and political ads. Currently, a user can filter ads using the same time-bound impression filters available for issue, electoral, and political ads, and by which Facebook product or platform the ad was run on. In addition, when a user clicks on an ad they are able to see the content of the ad and a brief description of the advertiser. They are unable, however, to see information on impressions, reach, and ad spend,²¹³ which is available on issue, electoral, and political ads. The company has stated it will develop similar tools for employment and credit opportunities in 2020.²¹⁴

In addition, according to U.K.-based nonprofit Privacy International, Facebook's efforts to provide greater transparency and accountability around its advertising platform have largely been focused in certain countries. The organization found that the company provides detailed information related to political advertising in only 35 countries and regions, such as the United States, the European Union, and India.²¹⁵ Facebook is the primary method used by political groups to target and influence political voters in approximately 80 percent of the world, and this therefore demonstrates a lack of meaningful transparency and accountability to all of Facebook's users.²¹⁶

Further, Facebook's transparency efforts do not involve the disclosure of any data on how the platform enforces its ad content and targeting policies, and how many ads and accounts are subsequently removed or suspended. Publishing such information would provide a vital transparency and accountability mechanism because it would highlight how, and if, a platform is enforcing its advertising-related policies.

Beyond these deficiencies in transparency, Facebook has also raised serious concerns through its October 2019 decision to exempt political ads from its fact-checking process and rules. According to Mark Zuckerberg, Facebook introduced this change based on the notion that private companies should not be able to censor politicians or the news. The company has maintained this stance despite pressure from policymakers, civil society, and the public.²¹⁷ However, researchers like Siva Vaidhyanathan have asserted that the policy change was introduced because Facebook is unable to enforce its ad fact-checking policies globally for political ads and because political advertising is a valuable source of revenue and online influence around the world.²¹⁸

Given the impact that political advertising can have on the mindsets and behaviors of users, this policy change by Facebook is concerning.²¹⁹ Many civil society organizations and researchers have called for Facebook to follow platforms such as Twitter and ban political ads altogether. However, some experts have suggested that in doing so Facebook would put first-time political candidates aspiring to enter politics at a disadvantage, as the internet platform has been found to create a greater level playing field for challengers than television.²²⁰ Facebook has also faced similar calls as Google, urging it to instead prohibit political ad microtargeting.²²¹ However, most recently they have disclosed that they will not be making changes to their microtargeting policy.²²²

User Controls on Facebook's Advertising Platform

Facebook has also introduced a limited set of controls and features that enable users to understand and control how ads are targeted and delivered to them. In 2014, Facebook introduced the "Why am I seeing this ad?" tool, which lets users see how factors such as their demographic information, interests, and website

visits can be used to deliver specific ads to them in their news feeds. The tool also provides additional details about the ads a user sees if they find that the information on an advertiser's list of target users matches information on a user's Facebook profile. This information includes when an advertiser uploaded the information, or if an advertiser worked with a marketing partner to deliver their ad campaign.²²³

Users also have the option to adjust the ads that they see by adjusting their ads preferences.²²⁴ Users can change preferences such as their interests; what Facebook-related information the platform uses to deliver ads to them; as well as whether Facebook should be able to deliver ads to them based on data from partners, activity on Facebook Company Products that a user sees elsewhere, and their social actions. Users can also choose to hide ads from a certain advertiser.²²⁵ In addition, in January 2020, the company announced that this year they would be introducing further controls, including allowing users to choose if they want to stop seeing ads from an advertiser who has created a Custom Audience from a customer list, or make themselves eligible to see ads if an advertiser used a list to exclude them. The company also stated it will introduce a feature that will enable users to choose if they would like to see fewer social and political ads on Facebook and Instagram.²²⁶

Although this new set of controls provides a positive step forward, an opt out mechanism—and one which requires—users to understand Facebook's ad platform, puts the onus on users to protect their own privacy. Similarly, requiring users to opt in to receiving ads they may have been excluded from also puts the onus on them to ensure their platform experience is equitable. These barriers are compounded because the controls are not easily comprehensible by the average user, and have been found to be difficult to understand and navigate. There are similar deficiencies with Facebook's new options for data portability. Although Facebook now offers users the option to download a vast range of the data the company has collected on them, including information on their likes, posts, and searches,²²⁷ this information is incomplete. For example, it does not include the data that advertisers and Facebook have collected and used for ad targeting and delivery.²²⁸

Case Study: LinkedIn



Source: Roman Pyshchuk / Shutterstock

A General Overview of LinkedIn's Advertising Platform

LinkedIn is a professional social network geared toward promoting business and employment-focused services. The company was founded in 2002²²⁹ and it ranks 58th for global internet engagement on Alexa rankings.²³⁰ In December 2016, Microsoft acquired LinkedIn for \$26.2 billion.

An examination of LinkedIn offers an interesting perspective on how automated tools can be used for ad targeting and delivery, because prior to LinkedIn's acquisition, advertising was not its primary source of revenue. The company also relied on other sources of revenue such as premium membership subscriptions and its learning solutions, which enable users to learn new skills and subject matter.

The company's ad operations transformed significantly after its acquisition. Before the deal, marketing and advertising solutions accounted for approximately 18 percent of LinkedIn's total revenue, at \$175 million. In comparison, talent solutions for recruiters accounted for approximately 65 percent of LinkedIn's total revenue at \$623 million and premium subscriptions accounted for approximately 17 percent of LinkedIn's total revenue at \$162 million.²³¹ Under Microsoft, little data has been available regarding the size and scope of LinkedIn's advertising business in comparison to other revenue streams.

However, LinkedIn shared that it expected to bring in approximately \$2 billion from its media business by the end of 2018, with the majority of this revenue coming from advertising.²³²

A Technical Overview of LinkedIn's Advertising Platform

Since its acquisition by Microsoft, LinkedIn has expanded its data collection practices. In May 2019, the company acquired Drawbridge, a machine-learning identity management company, with the aim of enhancing ad personalization²³³ through greater cross-device targeting.²³⁴ Prior to this, LinkedIn primarily relied on methods such as cookie-based retargeting. However, as browsers (such as Google Chrome) have limited this functionality, this has decreased the effectiveness of these methods.²³⁵ LinkedIn's expansion of data collection practices and advertising operations suggests that they are both vital components of growing the company's ad operations, and thus its revenue.

LinkedIn offers advertisers a range of methods through which they can engage with users: These include:²³⁶

1. **Text ads:** Text ads are text-based ads that link to an external website, such as a company's website. The ads appear in the right panel or top banner of numerous pages on LinkedIn.com, and they can include images (such as an advertiser's logo). Text ads can appear on both the desktop website and the mobile app.
2. **Sponsored Content:** Sponsored Content is a post that is promoted by a company that can appear on both the desktop website and the mobile app. Sponsored Content appears in a user's LinkedIn feed and it can feature text, an image or a video, and links to a website. Sponsored Content is always labeled as "Promoted" or "Sponsored."
3. **Sponsored InMail:** Sponsored InMail is a sponsored message that a user receives from an advertiser in their LinkedIn messaging inbox. The message is labeled as "Sponsored" in the subject line. These ads appear on both desktop and mobile devices.
4. **Dynamic ads:** Dynamic ads are personalized ads that are tailored based on a user's profile image, name, and job function. Dynamics ads can cover advertisers, jobs or professional content that a user may be interested in. Dynamic ads appear in the right panel of LinkedIn.com pages on desktop devices when a member is signed into their LinkedIn account.
5. **Display ads:** Display ads are image-based ads that appear in the right panel of LinkedIn.com pages on desktop devices.

6. LinkedIn Audience Network: Most of the ads that LinkedIn places appear on LinkedIn.com. However, an advertiser can also opt to join the LinkedIn Audience Network, which enables ads to appear on third-party publisher websites that LinkedIn characterizes as “brand-safe premium publisher apps and sites.” These are apps and sites that LinkedIn believes its members would benefit from viewing and that LinkedIn has deemed are in compliance with their brand safety protocols and requirements.²³⁷ According to LinkedIn, during this process, no data that can be used to identify users is transferred to the third-party websites.²³⁸

Just as on Facebook and Google, when an advertiser creates an ad campaign on LinkedIn, they have to select the type of ad they want to run, define their business goals, select their audience based on targeting parameters, and set a budget for their campaign. Advertisers can also choose to use LinkedIn’s Matched Audiences feature to engage with users who have shown an interest in their company before by visiting their website or sharing their email address.²³⁹ Once LinkedIn users see or interact with an ad, the advertiser receives an aggregated report with information on user impressions. LinkedIn asserts that it does not provide advertisers with data on specific users, and it does not provide data that could be used to identify specific users.²⁴⁰ Rather, the company states that it provides advertisers with aggregated “general demographic information” on all of the users who have seen an ad.²⁴¹ This information could include, for example, the cities that users who engaged with the ads live in.

LinkedIn’s advertising auction system operates similarly to those of Google and Facebook. Advertiser’s can bid per click (CPC) or per thousand impressions (CPM), depending on the advertiser’s campaign objectives. Bidding on LinkedIn’s advertising products operates in a second-price auction manner. This means that if an advertiser wins with a bid they placed, they only need to pay the minimum amount that is needed to beat the second-place bidder, no matter what their maximum bid was. LinkedIn also calculates and considers a relevance score when determining which ads a user should see. This score measures how engaging a given update is for the targeted audience, and it is calculated using a combination of click-throughs, comments, likes, shares, and other data points. By factoring in such a relevance score, LinkedIn aims to deliver relevant advertisements to its users, thus incentivizing them to spend more time on the platform.²⁴²

Since its acquisition in 2016, the company has introduced a host of new advertising features that aim to broaden targeting and delivery opportunities. For example, over the course of 2019, the company expanded its targeting options to include categories such as interest-based and Lookalike Audiences (similar to Facebook’s Lookalike Audiences feature). This is in addition to existing parameters such as company affiliation, demographics, education, job experience, and the advertiser’s own contact lists.²⁴³

Currently, the only targeting parameter LinkedIn requires advertisers to choose is location. Permanent and long-term location data is collected from the locations users specify in their profile. This information is used in advertising campaigns, such as Sponsored InMail campaigns. Short-term location data is inferred from a user's IP address.²⁴⁴ While continuous and granular location tracking can result in the collection of highly personal and sensitive data, collecting a user's single, static location raises fewer privacy concerns. Moreover, since a user's location in their profile is likely relevant to which jobs may interest them, this information has clear relevance to LinkedIn's primary function of connecting users with job opportunities, and it permits the platform to deliver specific employment-related results.

In addition, because LinkedIn is an employment-focused website, the company has only narrowly permitted targeting based on sensitive and protected characteristics. For example, the platform only permits gender-based targeting in specific instances, such as for women-oriented events and conferences. In addition, all ads that are based on gender-related targeting parameters go through a specialized review process, which can be automated or manual depending on the type of ad and the advertiser's history on LinkedIn's ad network.²⁴⁵

One component of this review process requires advertisers to self-certify on LinkedIn's Campaign Manager platform that, if an ad is related to employment, housing, education, or credit, it will not use LinkedIn to discriminate based on gender or other protected characteristics. Once an advertiser self-affirms that they are not intending to discriminate based on gender, the option to target their ads by gender becomes available to them.²⁴⁶ While this self-certification process aims to highlight the importance of—and ensure equitable procedures during—ad targeting, it does not provide strong safeguards against discrimination based on gender and other protected characteristics. An advertiser could easily self-certify based on false information or intentions, and continue to implement discriminatory targeting and delivery outcomes on the platform.

Controversies Related to LinkedIn's Advertising Platform

Since its acquisition, LinkedIn has adopted a more aggressive and streamlined targeted advertising model. These advertising practices have yielded some concerning outcomes. For example, a 2017 investigation by *ProPublica* and the *New York Times* found that major companies such as Amazon, T-Mobile, and Goldman Sachs were using platforms such as Facebook, Google, and LinkedIn to run ad campaigns that explicitly excluded audiences over the age of 40 years old from viewing and responding to employment-related ads.²⁴⁷ Such practices violate the Age Discrimination in Employment Act (ADEA), which prohibits employment discrimination against people over the age of 40.²⁴⁸

In response to the investigation, LinkedIn changed its ad targeting offerings so that age-related targeting for employment ads is only available to advertisers who affirm through a self-certification process that they are not running age-discriminatory ads.²⁴⁹ If the campaign's objective is to generate talents leads and recruitment, LinkedIn does not offer targeting based on age at all.²⁵⁰ Additionally, the self-certification process advertisers need to go through is similar to the self-certification process advertisers need to go through when targeting based on gender. Therefore, while it demonstrates an attempt to provide greater accountability around LinkedIn's ad targeting and delivery practices, it does not establish sufficient and concrete safeguards.

Another troubling example emerged in April 2019. Cynet Systems, a technology recruiting firm based in Virginia, posted an employment ad on LinkedIn and other platforms for an account manager. The ad listed a preference for a "Caucasian who has good technical background." The company also posted an ad for an account manager, soliciting "female candidates only."²⁵¹ According to a LinkedIn spokesperson, the company removed these ads once they were discovered. LinkedIn also stated that such postings were not frequent.²⁵² The platform's Ad Targeting Discrimination guidelines outline that ads that are flagged as discriminatory are instantly removed.²⁵³ However, there is little transparency around the scope and volume of ads that are removed for violating LinkedIn's advertising policies, as well as little transparency around why ads are removed.

Steps LinkedIn Has Taken to Adjust Its Advertising Platform Since the 2016 U.S. Presidential Elections

Compared to Google and Facebook, LinkedIn played a smaller role in hosting campaigns to influence voter behavior during the 2016 U.S. presidential elections. However, Russian operatives were found to have used the platform to create fake accounts, aiming to connect with individuals who were part of U.S.-based political organizations.²⁵⁴

In 2018, LinkedIn banned political advertising on its platform entirely. Although this is not an insignificant decision, the ramifications of this decision for the platform's overall functionality and advertising revenue stream are small compared to the impact that platforms such as Google and Facebook would face if they made a similar decision. This is because LinkedIn's services focus on professional—rather than personal—social networking, and as a result political advertising is primarily useful for politically-driven groups looking to use the platform to recruit members rather than expressly spread their message. In addition, in May 2019, the platform also introduced an Ad Transparency tool for LinkedIn company pages that enables LinkedIn users to view all of the Sponsored

Content that advertisers have run on LinkedIn in the past six months when visiting a company's LinkedIn page²⁵⁵

The company does not currently have an ad transparency database for any type of ads. In addition, although the company does publish a transparency report that highlights the scope and scale of its content moderation efforts to an extent, it does not disclose any data on how the platform enforces its ad content and targeting policies. Nor does it disclose how many ads and accounts are subsequently removed or suspended. As previously discussed, this is a vital mechanism for ensuring transparency and accountability around the enforcement of a platform's ad-related policies.

User Controls on LinkedIn's Advertising Platform

Like Google and Facebook, LinkedIn offers its users a limited set of ad controls. These include whether a user receives targeted ads based on the websites they have visited, whether a user sees promoted jobs and ads on websites and applications that are separate from LinkedIn, how certain ads are targeted and delivered based on a user's profile data, and whether a user views promoted jobs and ads based on their activity and similar members' activities on both LinkedIn and Bing (Microsoft's search engine).

Users also have the ability to control what data LinkedIn collects on them, to an extent. For example, users can opt out of having the company collect and subsequently target and deliver ads to them based on factors such as their connections, location, demographics, companies they follow, groups, education, and employer.

Users can also opt out of seeing targeted and promoted ads based on information provided to third parties and based on prior engagements by the user with ads. However, there is no single universal control for opting out of receiving advertising on the platform. In addition, LinkedIn states that even after a user opts out of certain forms of advertising, they "may continue to see ads that feature an option to manage your advertising preferences, including to opt out of interest-based advertising."²⁵⁶

Promoting Fairness, Accountability, and Transparency Around Ad Targeting and Delivery Practices

As described in this report, the process of targeting and delivering ads on internet platforms often involves the use of automated tools and algorithmic decision-making. Although internet platforms assert that these practices enable them to deliver more relevant, personalized, and interesting content to users, the use of automated tools in these instances can also result in inequitable and discriminatory outcomes and practices. The use of automated tools in these situations also relies on massive amounts of user data, which requires numerous highly-invasive data collection methods. In addition, there is a significant lack of transparency and accountability around how these ad targeting and delivery practices are implemented. Because digital advertising currently serves as the primary source of revenue for many large internet companies, these practices, however concerning, continue.

Going forward, internet platforms, civil society, and researchers should consider the following set of recommendations in order to promote greater fairness, accountability, and transparency around algorithmic decision-making in this space. This section also includes recommendations for U.S. policymakers in this regard. However, 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, this report provides limited recommendations for action by policymakers.

Recommendations for Internet Platforms

Internet platforms need to provide greater transparency and accountability around their advertising operations, both at the user level and in the larger advertising ecosystem. In order to achieve this, internet platforms should:

- 1) **Publish comprehensive and comprehensible descriptions of advertising content policies.** These advertising content policies should clearly delineate what categories of ads, types of ad content, and accounts are prohibited on the platform. They should also explain what tools and processes (e.g. automated tools) the platform uses to identify ads and accounts that violate its advertising content policies. These advertising content policies should be easy to access and understand. Further, the platform should provide a change log or archive of past advertising content policies so users can see how these policies have been altered over time. Whenever an internet platform changes its advertising content policies, users should be notified.

2) Publish comprehensive and comprehensible descriptions of advertising targeting policies. These ad targeting policies should clearly outline what information the platform and advertisers can use to target ads to users (e.g. location information), which targeting parameters are prohibited on the platform, and what tools and processes (e.g. automated tools) the platform uses to identify ads and accounts that violate its ad targeting policies. Furthermore, these policies should explain what review process, if any, algorithmically generated ad targeting categories undergo. We recommend that algorithmically generated ad targeting categories be reviewed by humans in order to prevent discriminatory and harmful outcomes. All of an internet platform's ad targeting policies should be easy to access and understand. The platform should provide a change log or archive of past ad targeting policies so users can see how these policies have been altered over time. Whenever an internet platform changes its ad targeting policies, users should be notified.

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. These include categories such as race, ethnicity, religion, disability, and socioeconomic background, and known proxies such as neighborhood of residence. Companies should also establish comprehensive enforcement mechanisms to implement their anti-discrimination policies, as outlined in recommendations four through six below.

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. This review process should be comprehensive and involve pre-review and approval of ads by the internet platform. During this review process, the internet platform should ensure that all ads and accounts advertising in these categories are in compliance with the platform's advertising content policies and its ad targeting policies. Because processes in which an advertiser simply "self-certifies" that they are not engaging in prohibited discriminatory practices can easily be abused, they are not sufficient. In addition, although some companies are investing in developing automated tools that can identify certain categories of ads (such as political advertising) these tools are unable to fully understand subjective definitions and uses of human speech. As a result, this review process should always involve human review prior to a final decision. Internet platforms should publicly disclose the scope and methods of this review process, including to what extent automated and human review are involved, and should disclose the guidelines that human reviewers use to make determinations on whether or not an ad, account, or targeting parameter should be permitted on the platform. Such comprehensive review and authentication procedures can help reduce instances of bias and discriminatory outcomes in the ad targeting and delivery process.

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.** In particular, these audits should evaluate the algorithms that are used to generate ad targeting categories. Internet platforms should conduct these audits proactively, as well as in response to credible allegations of violations of user privacy, freedom of expression, freedom of information, or cases of discrimination surfaced by community partners, civil society organizations, activists, etc. Companies should use the results of these audits to refine and improve ad targeting algorithms and make them more fair, accountable, and transparent. These audits should be conducted by an external third party, and companies should make summaries publicly available.

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.** These audits should particularly evaluate how automated ad delivery systems can inappropriately prevent certain categories of users from viewing ads due to inferences made on the basis of collected data. Internet platforms should conduct these audits proactively, as well as in response to cases of violations of user privacy, freedom of expression, freedom of information, or cases of discrimination surfaced by community partners, civil society organizations, activists, etc. Companies should use the results of these audits to refine and improve ad delivery and optimization algorithms and make them more fair, accountable, and transparent. These audits should be conducted by an external third party, and companies should make summaries publicly available.

7) **Empower users with comprehensive tools that help them understand how and why ads are targeted and delivered to them.** Internet platforms should enable users to understand how and why they are seeing certain ads. This information should be easy to access and understand. In particular, when an individual user clicks on a specific ad, they should be able to see the following:

- What factors (e.g. demographic information, interests, browsing history, etc.) about them were considered when targeting and delivering the ad to them.
- Whether the user appears on that advertiser's list of target users.

When a user clicks on a certain ad to learn more, they should also be able to access further information on:

- Which advertiser(s) lists a user appears on. This information should include whether an advertiser provided specific information on the user to the advertising platform, or whether it was acquired through a marketing partner or data broker.
- Whether an ad was delivered to a user because the advertising platform identified them as potentially interested in the content of the ad (e.g. whether they were part of a Lookalike Audience).

8) Explain to users why the platform collects, infers, and shares user data.

This information should outline the purposes and scope of each of these practices. It should also include an explanation of the risks associated with such data collection, inference, and sharing practices. This information should be easy to access and understand, and help inform users' decisions regarding the user controls recommended in recommendation 9 below.

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. These controls should be easy to access and understand. They should also be available to all users of the service, whether logged in, logged out, or not associated with a particular account. At a minimum, users should be able to:

- Adjust their advertising preferences to select and change the factors (e.g. demographic information, interests, browsing history, etc.) that advertisers may consider when targeting ads to them, and that the advertising platform may consider when delivering ads to them. These preference settings should include the ability to completely opt out from having any of these factors considered.
- Change their preferences related to whether they are shown ads that are targeted and delivered based on data from partners or their activity on related products or websites.
- Manage how all advertisers contribute to their online ad experience, including having the ability to hide and opt out of receiving ads from specific advertisers.
- Control whether and how internet platforms collect data on them. This should include the ability to automatically delete their browsing, app, and location history. Ideally, users should have to opt in to such data collection practices. At a minimum, companies should provide users with controls that enable them to pause and/or opt out entirely from practices that

involve collecting data beyond the basic information needed to establish an account.

- Determine whether or not they would like an internet platform to be able to make inferences about them based on data they have collected or provided.
- Decide whether they want to receive targeted ads. Ideally, users should have to opt in to having ads targeted and delivered to them on any platform. At a minimum, users should have access to controls that enable them to fully opt out of the ad targeting and delivery process. Users should have easy to use controls that let them opt out of all practices at once.

10) **Provide clear labels for sponsored and paid content across all of the platform's products, services, and ad networks.** Companies should label all forms of paid advertising as sponsored content. Companies should ensure that these labels are appended to ads across all of a company's products and services and should apply to all categories and mediums of advertising, whether political, employment related, text-based, video-based, etc.

11) **Create a publicly available online database of all of the ads that a company has run on its platform.** This can help explain a platform's ad operations in a comprehensive manner and can also enable meaningful trend analysis and research. This database should include ads from all categories of ads on the platform, including categories of ads that could have significant real-life consequences such as political ads, housing ads, employment ads, and credit ads. It should also be user-friendly. In particular, this database should include search functionality. In order to protect user privacy, the information in this database should not enable the identification of users that received the ad. At a minimum, this database should disclose the following information about each of the ads in the database:

- The format of the ad (e.g. text, video, etc.)
- The name of the advertiser
- What region the ad was run in
- How much the ad spend for the ad was
- The time period during which an ad was active
- Granular engagement and interaction information, such as how many users saw the ad, and the number of likes, shares, and views that an ad received

- What targeting parameters the advertiser specified to the advertising platform
- What categories of users the ad was eventually delivered to (i.e. what targeting parameters did the ad delivery system eventually select and optimize for)
- Whether the ad was delivered to a custom set of users or one generated by an automated system (e.g. Lookalike users)

12) Publish a transparency report that provides a granular overview of the platform’s advertising operations across all regions that it operates in.

This transparency report should be published along a consistent timeline (e.g. annually, quarterly, etc.). All of the data in the transparency report should be available in a structured data format (e.g. comma separated values), rather than or in addition to a flat PDF file. This is helpful to researchers who want to make use of the report data, as it simplifies the data extraction process and makes reports more accessible.

At a minimum, this transparency report should disclose the following general information for every reporting period:

- The total number of ads that a platform ran
- The total number of ads that a platform ran in each country in which it operates
- The total amount of ad spend across the platform
- The total amount of ad spend in each country in which it operates
- The top advertisers in each country
- The top keywords in each country

In addition, at a minimum, this transparency report should separately disclose the following information for ads that have been flagged or removed from the platform for every reporting period:

- The total number of ads flagged for violating the platform’s advertising content policies
- The total number of ads removed for violating the platform’s advertising content policies

- The total number of ads flagged for violating the platform’s ad targeting policies
- The total number of ads removed for violating the platform’s ad targeting policies
- A separate breakdown of the ads and accounts flagged and removed for violating the platform’s advertising content policies by:
 1. The advertising content policy they violated
 2. The format of the ad’s content (e.g. text, audio, image, video, live stream)
 3. The country of the advertiser
 4. For companies that operate more than one platform, the product or service on which the ad was run
 5. The detection method used (e.g. user flag, automated tool). Note that the identity of individual flaggers should not be revealed.
- A separate breakdown of the ads and accounts flagged and removed for violating the platform’s ad targeting policies by:
 1. The ad targeting policy they violated
 2. The format of the ad’s content (e.g. text, audio, image, video, live stream)
 3. The country of the advertiser
 4. For companies that operate more than one platform, the product or service on which the ad was run
 5. The detection method used (e.g. user flag, automated tool). Note that the identity of individual flaggers should not be revealed.

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.** These notice procedures are particularly important where ads are run by individuals or civil society organizations, as erroneous removal or moderation of their ads could particularly infringe on freedom of expression. In addition, given the lack of clear definitions around categories of ads such as political and

issue ads, such notice processes are important to protect freedom of expression. All notices should be available in a durable form that is accessible even if an advertiser's account is suspended or terminated. In addition, users who flag ads should have a log of ads they have reported and the outcomes of the review process. At a minimum the notice provided to advertisers who have had their ads or accounts flagged should include:

- The URL, content excerpt, and/or any other relevant information which would allow the party to identify the ad or account flagged or removed
- The specific portion of the advertising content policy and/or ad targeting policy that the ad or account was found to violate
- How the ad or account was detected or removed (e.g. user flag, government request, automated tool)
- An explanation of the process through which the relevant party can appeal the decision

14) **With regard to categories of ads that could have significant real-life consequences, such as political ads, housing ads, employment ads, and credit ads, 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.** These appeal procedures are particularly important where ads are run by individuals or civil society organizations, as erroneous removal or moderation of their ads could particularly infringe on freedom of expression. In addition, given the lack of clear definitions around categories of ads such as political and issue ads, such appeals processes are important to protect freedom of expression. At a minimum, the appeals process available to advertisers who have had their ads or accounts flagged and users who have flagged ads or accounts for these consequential categories of ads should include:

- Human review by a person or panel of people that were not directly involved in the initial decision-making process.
- The opportunity to provide additional context or information that should be considered in the appeals process.
- Meaningful notice which details the results of the appeals process. This notice should provide a comprehensive explanation of why the final decision was made.

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. In particular, this research should include platforms which have been less studied to date, such as Twitter, Reddit, Snapchat, etc. Additionally, new research should explore potentially consequential categories of ads beyond the ones discussed in this report as ads that could have significant real-life consequences such as political and employment ads. Other similarly consequential categories of ads may include education and healthcare, and these areas would benefit from further research as they can also yield discriminatory and harmful outcomes.

Recommendations for Civil Society and Researchers

In order to promote greater understanding and safeguards around how automated tools are used during the ad targeting and delivery processes, members of civil society 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.**

In particular, this research should focus on platforms that have been less studied to date, such as Twitter, Reddit, Snapchat, etc. As discussed above, this report has highlighted some categories of ads that could have significant real-life consequences, but there are other similarly consequential categories of ads, such as education and healthcare, that need to be further researched, as they can also yield discriminatory and harmful outcomes.

- 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.**

This will help ensure that users are adequately educated and aware of company practices, have a range of meaningful controls at their disposal, and know how to use them. It will also promote greater accountability around these algorithmic decision-making practices.

Recommendations for Policymakers

The recommendations for policymakers in this report are focused on U.S. policymakers. This is both because the platforms discussed are U.S. companies and also because the First Amendment of the U.S. Constitution imposes unique constraints on the extent to which U.S. policymakers can regulate how companies decide which content to permit on their platforms.

In order to help prevent discrimination and harmful outcomes as a result of algorithmic decision-making in ad targeting and delivery online, U.S. policymakers should:

1. **Clarify that all offline anti-discriminatory statutes apply in the digital environment.** Policymakers should ensure that anti-discrimination statutes apply online to the same extent as offline. Statutes such as the Civil Rights Act of 1964, the Fair Housing Act (part of the Civil Rights act of 1968), the Age Discrimination in Employment Act, and California's Unruh Civil Rights Act should be applied online as they are applied offline. Where necessary to fill gaps or clarify the applicability of such laws, Congress and state legislatures should enact appropriate legislation.
2. **Enact rules to require greater transparency from online platforms regarding their ad targeting and delivery practices.** The U.S. government is limited in the extent to which it can direct platforms how to decide what content to permit on their sites. However, Congress could improve accountability mechanisms by requiring greater transparency around these content policies and practices.

Notes

- 1 Jasmine Enberg, "Global Digital Ad Spending 2019," eMarketer, last modified March 28, 2019, <https://www.emarketer.com/content/global-digital-ad-spending-2019>
- 2 eMarketer Editors, "US Digital Ad Spending Will Surpass Traditional in 2019," eMarketer, last modified February 19, 2019, <https://www.emarketer.com/content/us-digital-ad-spending-will-surpass-traditional-in-2019>
- 3 Although Amazon's market share is small in absolute terms, compared to that of Google and Facebook, it has begun growing its advertising offerings and influence through a product-based ad business that is founded on product searches. <https://www.theverge.com/2018/10/1/17924422/google-search-targeted-advertising-youtube-expansion>
- 4 Enberg, "Global Digital," eMarketer.
- 5 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (PublicAffairs, 2019).
- 6 Dipayan Ghosh and Ben Scott, *Digital Deceit: The Technologies Behind Precision Propaganda on the Internet*, January 23, 2018, <https://www.newamerica.org/public-interest-technology/policy-papers/digitaldeceit/>
- 7 Nathalie Maréchal, "Targeted Advertising Is Ruining the Internet and Breaking the World," *VICE*, November 16, 2018, https://www.vice.com/en_us/article/xwjden/targeted-advertising-is-ruining-the-internet-and-breaking-the-world
- 8 Ranking Digital Rights, *Consultation Draft: Human Rights Risk Scenarios: Targeted Advertising* (2019).
- 9 Ranking Digital Rights, *Consultation Draft*. Amnesty International, *Surveillance Giants: How the Business Model of Google and Facebook Threaten Human Rights*, November 21, 2019, <https://www.amnesty.org/download/Documents/POL3014042019ENGLISH.PDF>
- 10 Amnesty International, *Surveillance Giants*.
- 11 Louise Matsakis, "Facebook's Targeted Ads Are More Complex Than It Lets On," *WIRED*, April 25, 2018, <https://www.wired.com/story/facebooks-targeted-ads-are-more-complex-than-it-lets-on/>
- 12 Matsakis, "Facebook's Targeted".
- 13 Zuboff, *The Age of Surveillance*.
- 14 Zuboff, *The Age of Surveillance*.
- 15 Muhammad Ali et al., *Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead to Skewed Outcomes*, April 3, 2019, <https://arxiv.org/pdf/1904.02095.pdf>
- 16 Ronan Shields, "How Serious Is Amazon's Challenge to Google and Facebook?," *AdWeek*, June 3, 2019, <https://www.adweek.com/programmatic/how-serious-is-amazons-challenge-to-google-and-facebook/>
- 17 A walled garden environment online is one that controls what content and services a user can access and which they cannot.
- 18 Programmatic advertising is the automated buying and selling of online advertising.
- 19 Ingrid Lunden, "AT&T Confirms It Is Buying Ad Platform AppNexus, Reportedly For Between \$1.6B-\$2B," *TechCrunch*, June 25, 2018, <https://techcrunch.com/2018/06/25/att-appnexus/>
- 20 Ronan Shields, "Verizon Media to Shutter Oath Ad Server," *AdWeek*, March 4, 2019, <https://www.adweek.com/programmatic/verizon-media-to-shutter-oath-ad-server/>
- 21 Ghosh and Scott, *Digital Deceit*.

- 22 Geoffrey Fowler, "Think You're Anonymous Online? A Third of Popular Websites Are 'Fingerprinting' You.," *The Washington Post*, October 31, 2019, <https://www.washingtonpost.com/technology/2019/10/31/think-youre-anonymous-online-third-popular-websites-are-fingerprinting-you/>
- 23 Dan Goodin, "Now Sites Can Fingerprint You Online Even When You Use Multiple Browsers," *Ars Technica*, February 13, 2017, <https://arstechnica.com/information-technology/2017/02/now-sites-can-fingerprint-you-online-even-when-you-use-multiple-browsers/>
- 24 Anthony Nadler, Matthew Crain, and Joan Donovan, *Weaponizing the Digital Influence Machine: The Political Perils of Online Ad Tech*, October 17, 2018, https://datasociety.net/wp-content/uploads/2018/10/DS_Digital_Influence_Machine.pdf
- 25 Nadler, Crain, and Donovan, *Weaponizing the Digital*.
- 26 Brendan Sullivan et al., Socioeconomic Group Classification Based on User Features, US Patent 20180032883, filed July 27, 2016, and issued Feb. 1, 2018, <http://appft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PG01&p=1&u=%2Fnehtahtml%2FPTO%2Fsrchnum.html&r=1&f=G&l=50&s1=%2220180032883%22.PG.NR.&OS=DN/20180032883&RS=DN/20180032883>
- 27 Ghosh and Scott, *Digital Deceit*.
- 28 Matsakis, "Facebook's Targeted".
- 29 Matsakis, "Facebook's Targeted".
- 30 PHD Media, "New Beauty Study Reveals Days, Times And Occasions When U.S. Women Feel Least Attractive," news release, October 2, 2013, <https://www.prnewswire.com/news-releases/new-beauty-study-reveals-days-times-and-occasions-when-us-women-feel-least-attractive-226131921.html>
- 31 Ghosh and Scott, *Digital Deceit*.
- 32 Charles Duhigg, "How Companies Learn Your Secrets," *The New York Times*, February 16, 2012, <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>
- 33 David Shepardson, "Trump Signs Repeal of U.S. Broadband Privacy Rules," *Reuters*, April 3, 2017, <https://www.reuters.com/article/us-usa-internet-trump-idUSKBN1752PR>
- 34 Steven Englehardt and Arvind Narayanan, *Online Tracking: A 1-million-site Measurement and Analysis*, http://randomwalker.info/publications/OpenWPM_1_million_site_tracking_measurement.pdf
- 35 Ali et al., *Discrimination Through*.
- 36 Ali et al., *Discrimination Through*.
- 37 Ali et al., *Discrimination Through*.
- 38 Ali et al., *Discrimination Through*.
- 39 Ali et al., *Discrimination Through*.
- 40 Ali et al., *Discrimination Through*.
- 41 Ali et al., *Discrimination Through*.
- 42 Ali et al., *Discrimination Through*.
- 43 Ali et al., *Discrimination Through*.
- 44 Ali et al., *Discrimination Through*.
- 45 A/B testing compares two versions of a variable by testing a user's response to variable A against variable B, and establishing which of the two variables is more effective.
- 46 Ghosh and Scott, *Digital Deceit*.
- 47 Ghosh and Scott, *Digital Deceit*.

- 48 Ghosh and Scott, *Digital Deceit*.
- 49 Ghosh and Scott, *Digital Deceit*.
- 50 Kashmir Hill, "How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did," *Forbes*, February 16, 2012, <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#6003b3f96668>
- 51 Duhigg, "How Companies".
- 52 Zuboff, *The Age of Surveillance*.
- 53 Maréchal, "Targeted Advertising".
- 54 Maréchal, "Targeted Advertising".
- 55 Maréchal, "Targeted Advertising".
- 56 Ranking Digital Rights, *Draft Indicators: Transparency and Accountability Standards for Targeted Advertising and Algorithmic Decision-Making Systems*, October 18, 2019, https://rankingdigitalrights.org/wp-content/uploads/2019/10/RDR-Index-Draft-Indicators_-Targeted-advertising-algorithms.pdf
- 57 Chris Gilliard, "Friction-Free Racism," *Real Life*, October 15, 2018, <https://reallifemag.com/friction-free-racism/>
- 58 Ali et al., *Discrimination Through*.
- 59 Tiffany Hsu, "Voter Suppression and Racial Targeting: In Facebook's and Twitter's Words," *The New York Times*, December 17, 2018, <https://www.nytimes.com/2018/12/17/business/russia-voter-suppression-facebook-twitter.html>
- 60 Ranking Digital Rights, *Draft Indicators: Transparency and Accountability Standards for Targeted Advertising and Algorithmic Decision-Making Systems*, October 18, 2019, https://rankingdigitalrights.org/wp-content/uploads/2019/10/RDR-Index-Draft-Indicators_-Targeted-advertising-algorithms.pdf
- 61 Tony Romm and Isaac Stanley-Becker, "Twitter to Ban All Political Ads Amid 2020 Election Uproar," *The Washington Post*, October 30, 2019, <https://www.washingtonpost.com/technology/2019/10/30/twitter-ban-all-political-ads-amid-election-uproar/>
- 62 TechFreedom, "Well-Intentioned 'Honest Ads' Bill Raises Serious Free Speech Concerns," TechFreedom, last modified October 19, 2017, <https://techfreedom.org/well-intentioned-honest-ads-bill-raises-serious-free-speech-concerns/> Sam Brodey, "Sen. Lindsey Graham Takes Heat from Conservatives for Backing John McCain's Election Meddling Bill," *Daily Beast*, May 10, 2019, <https://www.thedailybeast.com/lindsey-graham-takes-heat-from-conservatives-for-backing-mccains-election-meddling-bill> IFS Staff, "Analysis: 'Honest Ads Act' Targets Americans, Not Foreign Actors," Institute for Free Speech, last modified November 1, 2017, <https://www.ifs.org/news/analysis-honest-ads-act-targets-americans-not-foreign-actors/>
- 63 "Search Engine Market Share Worldwide - November 2019," StatCounter, <https://gs.statcounter.com/search-engine-market-share>
- 64 Alexa rankings ranks sites based on their global popularity with Alexa Toolbar users
- 65 "Google.com Competitive Analysis, Marketing Mix and Traffic," Alexa Internet, <https://www.alexa.com/siteinfo/google.com>
- 66 eMarketer Editors, "US Digital," eMarketer.
- 67 Richard Nieva, "Google's Ad Chief Pushes Privacy To Gain Back Frayed Trust," *CNET*, May 13, 2019, <https://www.cnet.com/features/googles-head-of-ads-pushes-privacy-to-gain-back-frayed-trust-prabhakar-raghavan/>
- 68 Zuboff, *The Age of Surveillance*.

- 69 Zuboff, *The Age of Surveillance*.
- 70 Zuboff, *The Age of Surveillance*.
- 71 Maréchal, "Targeted Advertising".
- 72 Nieva, "Google's Ad Chief".
- 73 Todd Spangler, "Google Killing Off DoubleClick, AdWords Names in Rebranding of Ad Products," *Variety*, June 27, 2018, <https://variety.com/2018/digital/news/google-doubleclick-adwords-rebranding-1202859088/>
- 74 Solon, "Google's Ad Tracking".
- 75 Solon, "Google's Ad Tracking".
- 76 Maréchal, "Targeted Advertising".
- 77 Latanya Sweeney, *Discrimination in Online Ad Delivery*, January 28, 2013, <https://arxiv.org/ftp/arxiv/papers/1301/1301.6822.pdf>
- 78 Christina Perricone, "The Ultimate Guide to Google Ads for 2019," HubSpot, <https://blog.hubspot.com/marketing/google-adwords-ppc>
- 79 Sweeney, *Discrimination in Online*.
- 80 Sweeney, *Discrimination in Online*.
- 81 Perricone, "The Ultimate," HubSpot.
- 82 Perricone, "The Ultimate," HubSpot.
- 83 Englehardt and Narayanan, *Online Tracking*.
- 84 Martin Moore, "How the Online Business Model Encourages Prejudice," *The Guardian*, October 28, 2018, <https://www.theguardian.com/technology/2018/oct/28/how-target-ads-threaten-the-internet-giants-facebook>
- 85 Englehardt and Narayanan, *Online Tracking*.
- 86 Google, "About Automated Bidding," Google Ads Help, <https://support.google.com/google-ads/answer/2979071>
- 87 Google, "Ad Rank," Google Ads Help, <https://support.google.com/google-ads/answer/1752122?hl=en>
- 88 Google, "About Quality Score," Google Ads Help, <https://support.google.com/google-ads/answer/7050591>
- 89 Google, "Check your Quality Score," Google Ads Help, <https://support.google.com/google-ads/answer/2454010?hl=en>
- 90 Google, "Check your," Google Ads Help.
- 91 Perricone, "The Ultimate," HubSpot.
- 92 Perricone, "The Ultimate," HubSpot.
- 93 Google, "About Responsive Search Ads," Google Ads Help, <https://support.google.com/google-ads/answer/7684791?hl=en>
- 94 Google, "How AdSense works," AdSense Help, <https://support.google.com/adsense/answer/6242051?hl=en>
- 95 Sweeney, *Discrimination in Online*.
- 96 Sweeney, *Discrimination in Online*.
- 97 Google, "How AdSense," AdSense Help.
- 98 Google, "How AdSense," AdSense Help.
- 99 Google, "How AdSense," AdSense Help.
- 100 Englehardt and Narayanan, *Online Tracking*.
- 101 Englehardt and Narayanan, *Online Tracking*.
- 102 Adi Robertson, "HUD Reportedly Also Investigating Google and Twitter In Housing Discrimination Probe," *The Verge*, March 28, 2019,

<https://www.theverge.com/2019/3/28/18285899/housing-urban-development-hud-facebook-lawsuit-google-twitter>

103 Alex Kantrowitz, "Google Allowed Advertisers To Target People Searching Racist Phrases," *BuzzFeed News*, September 15, 2017, <https://www.buzzfeednews.com/article/alexkantrowitz/google-allowed-advertisers-to-target-jewish-parasite-black#.js5Ynl3la3>

104 Sweeney, *Discrimination in Online*.

105 Samuel Gibbs, "Women Less Likely To Be Shown Ads For High-Paid Jobs on Google, Study Shows," *The Guardian*, July 8, 2015, <https://www.theguardian.com/technology/2015/jul/08/women-less-likely-ads-high-paid-jobs-google-study>

106 Mark Sweney and Alex Hern, "Google Ad Controversy: What The Row Is All About," *The Guardian*, March 17, 2017, <https://www.theguardian.com/technology/2017/mar/17/youtube-and-google-search-for-answers>

107 Nieva, "Google's Ad Chief".

108 Olivia Solon, "Google's Ad Tracking Is As Creepy As Facebook's. Here's How to Disable It," *The Guardian*, October 21, 2016, <https://www.theguardian.com/technology/2016/oct/21/how-to-disable-google-ad-tracking-gmail-youtube-browser-history>

109 Solon, "Google's Ad Tracking".

110 Todd Haselton and Megan Graham, "Google Uses Gmail to Track a History of Things You Buy — And It's Hard to Delete," *CNBC*, May 17, 2019, <https://www.cnn.com/2019/05/17/google-gmail-tracks-purchase-history-how-to-delete-it.html>

111 Julia Angwin, "Google Has Quietly Dropped Ban on Personally Identifiable Web Tracking," *ProPublica*, October 21, 2016, <https://www.propublica.org/>

[article/google-has-quietly-dropped-ban-on-personally-identifiable-web-tracking](https://www.propublica.org/article/google-has-quietly-dropped-ban-on-personally-identifiable-web-tracking)

112 Diane Greene, "As G Suite Gains Traction In The Enterprise, G Suite's Gmail And Consumer Gmail To More Closely Align," *The Keyword* (blog), entry posted June 23, 2017, <https://blog.google/products/gmail/g-suite-gains-traction-in-the-enterprise-g-suites-gmail-and-consumer-gmail-to-more-closely-align/>

113 Maréchal, "Targeted Advertising".

114 Scott Spencer, "An Update On Our Political Ads Policy," *The Keyword* (blog), entry posted November 20, 2019, <https://blog.google/technology/ads/update-our-political-ads-policy>

115 Google, *Political Advertising on Google*, <https://transparencyreport.google.com/political-ads/home>

116 Natasha Lomas, "France Slaps Google With \$166M Antitrust Fine For Opaque and Inconsistent Ad Rules," *TechCrunch*, December 20, 2019, <https://techcrunch.com/2019/12/20/france-slaps-google-with-166m-antitrust-fine-for-opaque-and-inconsistent-ad-rules/>

117 Spencer, "An Update," *The Keyword* (blog).

118 Tom Cardoso, "Google to Ban Political Ads Ahead of Federal Election, Citing New Transparency Rules," *The Globe and Mail*, March 4, 2019, <https://www.theglobeandmail.com/politics/article-google-to-ban-political-ads-ahead-of-federal-election-citing-new/>

119 Dipayan Ghosh, "What Is Microtargeting and What Is It Doing In Our Politics?," *Internet Citizen* (blog), entry posted October 4, 2018, <https://blog.mozilla.org/internetcitizen/2018/10/04/microtargeting-dipayan-ghosh/>

120 Ellen L. Weintraub, "Don't Abolish Political Ads On Social Media. Stop Microtargeting.," *The Washington Post*, November 1, 2019, <https://www.washingtonpost.com/news/technology/wp/2019/11/01/dont-abolish-political-ads-on-social-media-stop-microtargeting/>

www.washingtonpost.com/opinions/2019/11/01/dont-abolish-political-ads-social-media-stop-microtargeting/

121 Ashley Boyd, "Statement On Disallowing Microtargeting Of Political Ads," Mozilla, last modified November 8, 2019, <https://foundation.mozilla.org/en/blog/statement-disallowing-microtargeting-political-ads/>

122 Donie O'Sullivan, "Democrats Blast Google's New Rules For Political Ads," *CNN*, November 22, 2019, <https://www.cnn.com/2019/11/22/tech/democrats-google-ads/index.html>

123 Spencer, "An Update," *The Keyword* (blog).

124 O'Sullivan, "Democrats Blast".

125 Nieva, "Google's Ad Chief".

126 The DAA Principles are self-regulatory rules that major online advertising trade associations have agreed to follow. Currently the Principles cover political advertising, control to data used across devices, online behavioral advertising and multi-site data. "DAA Self-Regulatory Principles," Your Ad Choices, <https://youradchoices.com/principles>

127 Stephen Shankland, "Google Chrome Cracks Down On Cookies That Track You Online," *CNET*, May 7, 2019, <https://www.cnet.com/news/google-chrome-cracks-down-on-cookies-that-track-you-online/>

128 Eli Blumenthal and Jessica Dolcourt, "Google Will Now Let You Automatically Delete Location And Activity History. Here's How," *CNET*, May 2, 2019, <https://www.cnet.com/how-to/google-will-now-let-you-automatically-delete-location-and-activity-history-heres-how-to-use-it/>

129 Nieva, "Google's Ad Chief".

130 Dan Noyes, "The Top 20 Valuable Facebook Statistics – Updated July 2019," Zephoria Digital

Marketing, <https://zephoria.com/top-15-valuable-facebook-statistics/>

131 "Facebook.com Competitive Analysis, Marketing Mix and Traffic," Alexa Internet, <https://www.alexa.com/siteinfo/facebook.com>

132 eMarketer Editors, "US Digital," eMarketer.

133 Enberg, "Global Digital," eMarketer.

134 J. Clement, "Facebook's Advertising Revenue Worldwide From 2009 to 2018 (In Million U.S. Dollars)," Statista, last modified February 4, 2019, <https://www.statista.com/statistics/271258/facebook-ads-revenue-worldwide/>

135 Zuboff, *The Age of Surveillance*,

136 Jack Morse, "Mark Zuckerberg Doesn't Want To Talk About Tracking Users Who've Logged Out Of Facebook," *Mashable*, April 10, 2018, <https://mashable.com/2018/04/10/mark-zuckerberg-testimony-tracking-users/>

137 Englehardt and Narayanan, *Online Tracking*.

138 Taylor Hatmaker, "Facebook Will Cut Off Access To Third Party Data For Ad Targeting," *TechCrunch*, March 28, 2018, <https://techcrunch.com/2018/03/28/facebook-will-cut-off-access-to-third-party-data-for-ad-targeting/>

139 Josh Horwitz, "After Cambridge Analytica, Some Targeted Ads On Facebook Will Become Less Targeted," *Quartz*, March 29, 2018, <https://qz.com/1240369/after-cambridge-analytica-some-targeted-ads-on-facebook-will-become-less-targeted/>

140 Julia Carrie Wong, "Revealed: Facebook Enables Ads To Target Users Interested In 'Vaccine Controversies,'" *The Guardian*, February 15, 2019, <https://www.theguardian.com/technology/2019/feb/15/facebook-anti-vaccination-advertising-targeting-controversy>

- 141 Sam Levin, "Facebook Told Advertisers It Can Identify Teens Feeling 'Insecure' And 'Worthless,'" *The Guardian*, May 1, 2017, <https://www.theguardian.com/technology/2017/may/01/facebook-advertising-data-insecure-teens>
- 142 Private meeting with Facebook representative, January 10, 2020
- 143 Ali et al., *Discrimination Through*.
- 144 Aaron Rieke and Miranda Bogen, *Leveling the Platform: Real Transparency for Paid Messages on Facebook*, May 2018, <https://www.upturn.org/static/reports/2018/facebook-ads/files/Upturn-Facebook-Ads-2018-05-08.pdf>
- 145 Ali et al., *Discrimination Through*.
- 146 Ali et al., *Discrimination Through*.
- 147 Private meeting with Facebook representative, January 10, 2020
- 148 A commonly used visualization in the marketing and sales industries for understanding the process of converting potential leads into customers
- 149 Private meeting with Facebook representative, January 10, 2020
- 150 Ali et al., *Discrimination Through*,
- 151 Ali et al., *Discrimination Through*,
- 152 Ali et al., *Discrimination Through*,
- 153 Private meeting with Facebook representative, January 10, 2020
- 154 Private meeting with Facebook representative, January 10, 2020
- 155 Facebook, "About Ad Auctions," Ads Help Center, <https://www.facebook.com/business/help/430291176997542?id=561906377587030>
- 156 Private meeting with Facebook representative, January 10, 2020
- 157 Facebook, "About Ad Auctions," Ads Help Center.
- 158 Reach is the number of people that are exposed to an advertiser's message. Frequency is the number of times an advertiser touches each person with their message.
- 159 Facebook for Business, "Showing Relevance Scores for Ads on Facebook," Facebook for Business, last modified February 11, 2015, <https://www.facebook.com/business/news/relevance-score>
- 160 Ali et al., *Discrimination Through*.
- 161 Rieke and Bogen, *Leveling the Platform*.
- 162 Ali et al., *Discrimination Through*.
- 163 Sweney and Hern, "Google Ad Controversy".
- 164 Amy Gesenhues, "Facebook Gives Small Businesses New Advertising And Engagement Tools," *Marketing Land*, last modified May 7, 2019, <https://marketingland.com/facebook-gives-small-businesses-new-advertising-and-engagement-tools-260517>
- 165 Gilliard, "Friction-Free Racism".
- 166 Julia Angwin and Terry Parris Jr., "Facebook Lets Advertisers Exclude Users by Race," *ProPublica*, October 28, 2016, <https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race>
- 167 Julia Angwin, Ariana Tobin, and Madeleine Varner, "Facebook (Still) Letting Housing Advertisers Exclude Users by Race," *ProPublica*, November 21, 2017, <https://www.propublica.org/article/facebook-advertising-discrimination-housing-race-sex-national-origin>

- 168 Julia Angwin, Madeleine Varner, and Ariana Tobin, "Facebook Enabled Advertisers to Reach 'Jew Haters,'" *ProPublica*, September 14, 2017, <https://www.propublica.org/article/facebook-enabled-advertisers-to-reach-jew-haters>
- 169 Kaya Yurieff, "HUD Charges Facebook With Housing Discrimination In Ads," *CNN Business*, March 28, 2019, <https://www.cnn.com/2019/03/28/tech/facebook-hud-ad-discrimination/index.html>
- 170 Julia Angwin, Noam Scheiber, and Ariana Tobin, "Facebook Job Ads Raise Concerns About Age Discrimination," *The New York Times*, December 20, 2017, <https://www.nytimes.com/2017/12/20/business/facebook-job-ads.html>
- 171 U.S. Equal Employment Opportunity Commission, "Age Discrimination," U.S. Equal Employment Opportunity Commission, <https://www.eeoc.gov/laws/types/age.cfm>
- 172 NFHA Staff, "Fair Housing Groups Settle Lawsuit with Facebook: Transforms Facebook's Ad Platform Impacting Millions of Users," news release, March 19, 2019, <https://nationalfairhousing.org/2019/03/18/national-fair-housing-alliance-settles-lawsuit-with-facebook-transforms-facebooks-ad-platform-impacting-millions-of-users/>
- 173 Yurieff, "HUD Charges".
- 174 Yurieff, "HUD Charges".
- 175 Brian Heater, "Facebook settles ACLU Job Advertisement Discrimination Suit," *TechCrunch*, March 19, 2019, <https://techcrunch.com/2019/03/19/facebook-settles-aclu-job-advertisement-discrimination-suit/>
- 176 U.S. Department of Housing and Urban Development, "Housing Discrimination Under the Fair Housing Act," U.S. Department of Housing and Urban Development, https://www.hud.gov/program_offices/fair_housing_equal_opp/fair_housing_act_overview
- 177 Ahiza Garcia, "Facebook Denied Financial Services Opportunities To Women And Older People, Lawsuit Alleges," *CNN Business*, October 31, 2019, <https://www.cnn.com/2019/10/31/tech/facebook-class-action-women-older-people-financial-services/index.html>
- 178 Gaurav Laroia and David Brody, "Privacy Rights Are Civil Rights. We Need to Protect Them.," Free Press, last modified March 14, 2019, <https://www.freepress.net/our-response/expert-analysis/insights-opinions/privacy-rights-are-civil-rights-we-need-protect-them>
- 179 Ali et al., *Discrimination Through*.
- 180 Ali et al., *Discrimination Through*.
- 181 Ali et al., *Discrimination Through*.
- 182 Aaron Rieke and Corrine Yu, "Discrimination's Digital Frontier," *The Atlantic*, April 15, 2019, <https://www.theatlantic.com/ideas/archive/2019/04/facebook-targeted-marketing-perpetuates-discrimination/587059/>
- 183 Ali et al., *Discrimination Through*.
- 184 Ali et al., *Discrimination Through*.
- 185 Rieke and Yu, "Discrimination's Digital".
- 186 Rieke and Yu, "Discrimination's Digital".
- 187 Facebook et al., "Summary of Settlements Between Civil Rights Advocates and Facebook," National Fair Housing Alliance, <https://nationalfairhousing.org/wp-content/uploads/2019/03/3.18.2019-Joint-Statement-FINAL-1.pdf>
- 188 Megan Rose Dickey, "Facebook Civil Rights Audit Says White Supremacy Policy Is 'Too Narrow,'" *TechCrunch*, June 30, 2019, <https://techcrunch.com/2019/06/30/facebook-civil-rights-audit-says-white->

supremacy-policy-is-too-narrow/?
utm_medium=TCnewsletter&tpcc=TCdailynewslette
r

189 Wong, "Revealed: Facebook".

190 Private meeting with Facebook representative,
January 10, 2020

191 Private meeting with Facebook representative,
January 10, 2020

192 Yurieff, "HUD Charges".

193 Facebook, "Updates To Housing, Employment
and Credit Ads in Ads Manager," Facebook for
Business, last modified August 26, 2019, [https://
www.facebook.com/business/news/updates-to-
housing-employment-and-credit-ads-in-ads-manager](https://www.facebook.com/business/news/updates-to-housing-employment-and-credit-ads-in-ads-manager)

194 Private meeting with Facebook representative,
January 10, 2020

195 Spandana Singh, *Everything in Moderation: An
Analysis of How Internet Platforms Are Using Artificial
Intelligence to Moderate User-Generated Content*,
July 22, 2019, [https://www.newamerica.org/oti/
reports/everything-moderation-analysis-how-
internet-platforms-are-using-artificial-intelligence-
moderate-user-generated-content/](https://www.newamerica.org/oti/reports/everything-moderation-analysis-how-internet-platforms-are-using-artificial-intelligence-moderate-user-generated-content/)

196 Singh, *Everything in Moderation*.

197 Private meeting with Facebook representative,
January 10, 2020

198 Piotr Sapiezynski et al., *Algorithms that "Don't
See Color": Comparing Biases in Lookalike and
Special Ad Audiences*, December 17, 2019, [https://
arxiv.org/pdf/1912.07579.pdf](https://arxiv.org/pdf/1912.07579.pdf)

199 Rieke and Bogen, *Leveling the Platform*.

200 Hsu, "Voter Suppression".

201 Ranking Digital Rights, *Draft Indicators*.

202 Rieke and Bogen, *Leveling the Platform*.

203 Rob Leathern, "Expanded Transparency and
More Controls for Political Ads," news release,
January 9, 2020, [https://about.fb.com/news/
2020/01/political-ads/](https://about.fb.com/news/2020/01/political-ads/)

204 Facebook, "Facebook Ad Library," Facebook Ad
Library, [https://www.facebook.com/ads/library/?
active_status=all&ad_type=all&country=US&impress
ion_search_field=has_impressions_lifetime](https://www.facebook.com/ads/library/?active_status=all&ad_type=all&country=US&impression_search_field=has_impressions_lifetime)

205 Leathern, "Expanded Transparency."

206 Facebook, "About Potential Reach," Ads Help
Center, [https://www.facebook.com/business/help/
1665333080167380?id=176276233019487](https://www.facebook.com/business/help/1665333080167380?id=176276233019487)

207 Rieke and Bogen, *Leveling the Platform*.

208 Rieke and Bogen, *Leveling the Platform*.

209 [https://www.facebook.com/ads/manage/
customaudiences/tos.php](https://www.facebook.com/ads/manage/customaudiences/tos.php). This is a customized
settings page that is available to each logged in user.

210 Rieke and Bogen, *Leveling the Platform*.

211 Privacy International, "Social Media Companies
Have Failed To Provide Adequate Advertising
Transparency To Users Globally," Privacy
International, last modified October 3, 2019, [https://
privacyinternational.org/long-read/3244/social-
media-companies-have-failed-provide-adequate-
advertising-transparency-users](https://privacyinternational.org/long-read/3244/social-media-companies-have-failed-provide-adequate-advertising-transparency-users)

212 Facebook, "Doing More to Protect Against
Discrimination in Housing, Employment and Credit
Advertising," Facebook Newsroom, last modified
March 19, 2019, [https://about.fb.com/news/2019/03/
protecting-against-discrimination-in-ads/](https://about.fb.com/news/2019/03/protecting-against-discrimination-in-ads/)

213 Facebook, "Facebook Ad Library," Facebook Ad
Library.

214 Dickey, "Facebook Civil".

215 Privacy International, "Social Media," Privacy International.

216 Privacy International, "Social Media," Privacy International.

217 Mike Isaac and Cecilia Kang, "Facebook Says It Won't Back Down From Allowing Lies in Political Ads," *New York Times*, January 9, 2020, <https://www.nytimes.com/2020/01/09/technology/facebook-political-ads-lies.html>

218 Siva Vaidhyanathan, "The Real Reason Facebook Won't Fact-Check Political Ads," *The New York Times*, November 2, 2019, <https://www.nytimes.com/2019/11/02/opinion/facebook-zuckerberg-political-ads.html>

219 Alex Hern, "Facebook Exempts Political Ads From Ban On Making False Claims," *The Guardian*, October 4, 2019, <https://www.theguardian.com/technology/2019/oct/04/facebook-exempts-political-ads-ban-making-false-claims>

220 Isaac Stanley-Becker, "Ban Political Ads On Facebook? Upstart, Anti-Trump Candidates Object.," *The Washington Post*, November 10, 2019, https://www.washingtonpost.com/politics/ban-political-ads-on-facebook-upstart-anti-trump-candidates-object/2019/11/09/12be4f1a-fffd-11e9-8501-2a7123a38c58_story.html

221 Boyd, "Statement On Disallowing," Mozilla.

222 Alexandra S. Levine and Zach Montellaro, "Facebook Sticking With Policies On Politicians' Lies and Voter Targeting," *Politico*, January 9, 2020, <https://www.politico.com/news/2020/01/09/facebook-sticking-with-policies-on-politicians-lies-and-voter-targeting-096597>

223 Ramya Sethuraman, "Why Am I Seeing This? We Have an Answer for You," Facebook Newsroom, last modified March 31, 2019, <https://about.fb.com/news/2019/03/why-am-i-seeing-this/>

224 Facebook, "Why Am I Seeing Ads From An Advertiser On Facebook?," Facebook Help Center, <https://www.facebook.com/help/794535777607370>

225 Facebook, "How Does Facebook Decide Which Ads To Show Me?," Facebook Help Center, https://www.facebook.com/help/562973647153813?helpref=faq_content

226 Leathern, "Expanded Transparency."

227 Facebook, "Accessing & Downloading Your Information," Facebook Help Center, https://www.facebook.com/help/1701730696756992/?helpref=hc_fnav

228 Ross Schulman and Eric Null, "The Data Portability Act: More User Control, More Competition," New America's Open Technology Institute, last modified August 19, 2019, https://newamericadotorg.s3.amazonaws.com/documents/OTI_Data_Portability_Act_Appendix_A_Aug_2019.pdf

229 "About LinkedIn," LinkedIn, <https://about.linkedin.com/>

230 "LinkedIn.com Competitive Analysis, Marketing Mix and Traffic," Alexa Internet, <https://www.alexa.com/siteinfo/linkedin.com>

231 LinkedIn, "LinkedIn Announces Third Quarter 2016 Results," LinkedIn Newsroom, last modified October 27, 2016, <https://news.linkedin.com/2016/linkedin-announces-third-quarter-2016-results>

232 Sara Fischer, "LinkedIn Expects Media Biz To Bring In \$2 Billion In 2018," *Axios*, November 13, 2018, <https://www.axios.com/linkedin-expects-media-biz-to-bring-in-2-billion-in-2018-1542131728-009aba73-50d5-4acb-8282-035139ffcb4f.html>

233 Alex Wilhelm and Natasha Mascarenhas, "LinkedIn Snaps Up Drawbridge, Its Second Known Acquisition In 8 Months," *Crunchbase*, May 28, 2019,

<https://news.crunchbase.com/news/linkedin-snaps-up-drawbridge-its-second-known-acquisition-in-8-months/>

234 Kerry Flynn, "'Focusing On Pain Points': LinkedIn Adds More Ad-Targeting Options," *Digiday*, July 1, 2019, <https://digiday.com/marketing/focusing-pain-points-linkedin-improves-ad-platform-even-targeting-options/>

235 Flynn, "'Focusing On Pain'".

236 LinkedIn, "LinkedIn Ads and Marketing Solutions - Overview," LinkedIn Help, <https://www.linkedin.com/help/linkedin/answer/85787>

237 LinkedIn, "LinkedIn Audience Network - Quality & Brand Safety," Marketing Solutions Help, <https://www.linkedin.com/help/lms/answer/83627>

238 LinkedIn, "LinkedIn Ads and Marketing," LinkedIn Help.

239 LinkedIn, "What You Can Do With Matched Audiences," LinkedIn Marketing Solutions, <https://business.linkedin.com/marketing-solutions/ad-targeting/matched-audiences>

240 LinkedIn, "LinkedIn Ads and Marketing," LinkedIn Help.

241 LinkedIn, "LinkedIn Ads and Marketing," LinkedIn Help.

242 Andrew Kaplan, "How the LinkedIn Advertising Auction Works," *LinkedIn Marketing Solutions Blog*, entry posted March 21, 2017, <https://business.linkedin.com/marketing-solutions/blog/linkedin-b2b-marketing/2017/how-the-linkedin-advertising-auction-works>

243 Flynn, "'Focusing On Pain'".

244 LinkedIn, "Targeting Options for LinkedIn Advertisements," LinkedIn Help, <https://www.linkedin.com/help/linkedin/suggested/722/>

targeting-options-for-linkedin-advertisements?lang=en

245 LinkedIn, "Ad Targeting Discrimination," Marketing Solutions Help, <https://www.linkedin.com/help/lms/answer/86856/ad-targeting-discrimination?lang=en>

246 LinkedIn, "Ad Targeting," Marketing Solutions Help.

247 Julia Angwin, Noam Scheiber, and Ariana Tobin, "Dozens of Companies Are Using Facebook to Exclude Older Workers From Job Ads," *ProPublica*, December 20, 2017, <https://www.propublica.org/article/facebook-ads-age-discrimination-targeting>

248 U.S. Equal Employment Opportunity Commission, "Age Discrimination," U.S. Equal Employment Opportunity Commission.

249 Angwin, Scheiber, and Tobin, "Dozens of Companies".

250 LinkedIn, "Ad Targeting," Marketing Solutions Help.

251 Kari Paul, "Tech Firm Apologizes After Job Ad Seeks 'Preferably Caucasian' Candidates," *The Guardian*, April 29, 2019, <https://www.theguardian.com/technology/2019/apr/29/tech-firm-apologizes-after-job-ad-seeks-preferably-caucasian-candidates>

252 Paul, "Tech Firm".

253 LinkedIn, "Ad Targeting," Marketing Solutions Help.

254 Michael Kan, "LinkedIn Ousts Fake Accounts Targeting US Political Groups," *PC Mag*, August 8, 2018, <https://www.pcmag.com/news/362997/linkedin-ousts-fake-accounts-targeting-us-political-groups>

255 Mindaou Gu, "Making Advertising More Transparent on LinkedIn," *LinkedIn Marketing*

Solutions Blog, entry posted May 21, 2019, <https://business.linkedin.com/marketing-solutions/blog/linkedin-news/2019/making-advertising-more-transparent-on-linkedin->

256 LinkedIn, "LinkedIn Ads and Marketing," LinkedIn Help.



This report carries a Creative Commons Attribution 4.0 International license, which permits re-use of New America content when proper attribution is provided. This means you are free to share and adapt New America’s work, or include our content in derivative works, under the following conditions:

- **Attribution.** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

For the full legal code of this Creative Commons license, please visit creativecommons.org.

If you have any questions about citing or reusing New America content, please visit www.newamerica.org.

All photos in this report are supplied by, and licensed to, [shutterstock.com](https://www.shutterstock.com) unless otherwise stated. Photos from federal government sources are used under section 105 of the Copyright Act.