OPEN RIGHTS GROUP

PROFILING BY PROXY
HOW META'S DATA DRIVEN ADS FUEL DISCRIMINATION

June 2025

ABOUT ORG

Open Rights Group (ORG) is a UK based digital campaigning organisation working to protect fundamental rights to privacy and free speech online. We are a grassroots organisation with supporters and local groups across the UK.

Our work on data protection and privacy includes challenging the immigration exemption to UK data protection law, defending the General Data Protection Regulation (GDPR) from attempts to water down its provisions, and challenging uncontrolled and unlawful data sharing by online advertisers.

openrightsgroup.org

Author **Anna Dent**

EXECUTIVE SUMMARY

Data profiling, or behavioural profiling, is one of the most powerful tools in digital advertising, and generates billions of dollars of revenue for the sites that host ads. In 2024, Meta earned over $160 billion from advertising revenue, over 95% of its total global revenue.[1] Their net income was over $62 billion in the same year.[2] Platforms such as Meta's Facebook and Instagram collect, store and analyse thousands of data points about each individual user, gleaned from the things we do while on their platforms and elsewhere on the Internet. Our individual profiles can contain demographic, behavioural and geographical data as well as the kind of devices we use to access social media.

Profiling is the recording and categorisation of these thousands of data points into a digital profile. Algorithms look for patterns in our data which provide information about our shopping habits, topics we are interested in, places we like to visit and much more. These characteristics enable advertisers to find audiences that are most likely to respond to their adverts; they can select from extensive lists of interests, behaviours and other characteristics which have been automatically assigned to an individual user by Meta's AI. This is ad targeting.

Much of the content of a data profile is not information that a user has directly disclosed, but that has been inferred by their online behaviours, and collated with other data into a composite picture. It is this combination that underpins much consumer unease, and creates the potential for ad targeting to cause discrimination and harm.

The combining of multiple data sources and subsequent analysis creates what are called proxies: where the platform or advertiser does not hold data that directly relates to a characteristic they wish to target, other characteristics can stand in for the desired one. Proxies can be created by training an algorithm with historical data about user activity, such as clicking on an advert for a particular product. The algorithm can then 'learn' what other characteristics are associated with that user action.

Proxies are powerful tools for advertisers, but can also reveal sensitive information and data which is protected by anti-discrimination and data processing laws. These laws mean that companies like Meta are not allowed to process certain types of data, or make decisions based on that data which may discriminate against certain users. This can happen accidentally, or intentionally by unscrupulous advertisers or platforms, but either way, proxies mean that restrictions on data processing and ad targeting which are meant to prevent discrimination or harm can never be fully effective.

Profiling and ad targeting create huge revenues for the platforms that host adverts, and many advertisers believe that they improve audience engagement and conversion into sales. There is some evidence to support this, though it is not conclusive, and there appears to be increasing dissatisfaction and unease among advertisers with the business and social impacts of ad targeting. At the moment however they have few other options.

Many consumers are concerned about being profiled, and feel uncomfortable about 'creepy' adverts following them around the Internet: less than a fifth of UK Internet users said they are happy with their data being used in exchange for a free or personalised service for example. But at the same time, consumers are faced with a lack of other options, either having to accept profiling and targeting or stay off the platforms that have become an essential part of many of our lives.

Many examples exist of discrimination and harm caused by profiling and ad targeting. They include gender discrimination in who sees job ads; racial discrimination in housing and education ads; predatory data collection and targeting by online gambling sites; inappropriate products being

1 Statista, 2025, Annual advertising revenue of Meta Platforms worldwide from 2009 to 2024

2 Statista, 2025, Annual revenue and net income generated by Meta Platforms from 2007 to 2024

promoted to under 18s including alcohol and pharmaceutical drugs; discriminatory and predatory advertising of credit; sharing of sensitive health information between NHS websites and Meta; and the targeting of scams to more vulnerable users.

With the introduction of Generative AI to Meta's ad tools it is likely that existing problems with opacity and lack of accountability will worsen, and discriminatory targeting could increase.

It is likely that many users of Meta platforms do not fully understand how data profiling works, the volume of information collected and inferred about them, or the uses to which it is put, because so much of the process is opaque. There are some options within the platforms for users to reduce the amount of data collected and processed, but these do not switch off profiling or targeting completely.

Outside of Meta's own settings, as consumers and users of sites like Facebook we are limited in our ability to opt out of data processing and ad targeting. Under the UK GDPR we have a right to object to our data being processed, which applies absolutely in the case of direct marketing. In March, 2025 Meta settled in legal action brought human rights campaigner Tanya O'Carroll, and said that they would no longer process her personal data for targeted advertising.[3] Since then thousands of people in the UK have requested that Meta stop profiling them for advertising.[4]

In theory we should be able to use sites without giving up our personal information for advertising purposes, but in reality Meta is not respecting our rights, and as consumers we have no way of forcing them to do so. Other models of digital advertising are emerging, and there has been some progress on ad transparency, but there is a long way to go before users can truly choose how much they share and how it is used by platforms and advertisers. These next steps are recommended to begin the journey:

Respect people's right to consent to targeted advertising

Every user of a site or platform which uses profiling and ad targeting should only see targeted ads if they have consented for their data to be processed for this purpose. People should also be able to simply and effectively use their right to opt out at any time. It should not be a paid-for privilege, but a universally available right. Opting users out of data profiling and targeting should be the default for sites like Facebook, with users who prefer targeted ads able to opt in if they wish.

Improve ad transparency

The transparency introduced by the Meta Ad Library should be built on and strengthened, with all ads subject to a stronger minimum level of transparency, and access to the Library should be freely available without logging into a Meta account.

Develop and support new models of adtech

Ad targeting doesn't have to be done through data profiling: contextual advertising can achieve similar results without collecting personal data and violating user privacy. This and other models of privacy-preserving online advertising should be developed and supported by advertisers and platforms.

User switching and Interoperability

While the market has an incentive for attention, which has little impact on user retention, the same harms are likely to emerge. For a better ad market to emerge, users need to be able to disengage with platforms with ease. Interoperability and user switching can help markets become more responsive, as users can move to choose better user experiences, including more truthful and less exploitative advertising environments, without losing their friends and contacts.

3 https://www.bbc.co.uk/news/articles/c1en1yjv4dpo

4 https://www.openrightsgroup.org/press-releases/metas-consent-or-pay-must-not-allow-stalker-ads-report

INTRODUCTION

Platforms like Facebook make billions of pounds every year from showing adverts to their users. In fact, for Meta, Facebook's parent company, online advertising is the primary source of revenue.

As Internet users we've probably all experienced the creepy feeling of looking at a product on one website, only for it to follow us around every other site we visit, or when an ad appears on our Instagram feed which is weirdly appropriate to our life or interests.

The reason that Meta makes so much from advertising is the same reason that online ads seem to know us better than we know ourselves: the multi-billion dollar adtech industry. It involves harvesting thousands of data points on everyone that spends time online to create detailed profiles of who we are, what we enjoy doing, where we live, who we're friends with and crucially, what we like spending money on.

Some people are willing to accept the creepy feeling, preferring to see ads more relevant to them rather than a random selection, but many are unhappy with the trade-off between using a platform and the capture and use of personal data. If you want to opt out your rights and choices are seriously limited, and if you're part of a community that experiences discrimination or is vulnerable in some way, the data profiling and ad targeting that are at the heart of Meta's ad offer can cause serious harm.

Targeting does not need to be achieved through the use of sensitive or restricted personal data, it can be done using contextual browsing information, but behavioural profiling is deeply entrenched in the mainstream adtech model.

Meta knows far more about us than we consciously or proactively disclose. By bringing together data from across its own platforms and everywhere else we spend time online, and analysing it using powerful AI, it has the ability to 'learn' who we are and which ads we are most likely to respond to. In creating these complex user profiles, seemingly innocuous data can reveal deeply personal information about us, or allow ads to be shown in ways that unlawfully exclude or target people.

Outside of a handful of investigations by civil society and reporters, most of this happens without our knowledge, with no straightforward way to challenge it. Our rights on paper are not enforceable in reality because Meta does not respect them. There needs to be a rebalancing of rights between platforms and users: users must be able to enforce their data rights without having to resort to legal action, or giving up the platforms and sites that have become central to our digital lives.

WHAT IS DATA PROFILING?

Data or behavioural profiling (referred to simply as profiling from hereon) is one of digital advertising's most powerful advertising tools. Major platforms and websites including Facebook, Google, YouTube and Instagram use it as part of their wider approach to adtech, from which they generate a significant proportion of their revenues.[5] In 2024, Meta earned over $160 billion from advertising revenue, over 95% of its total global revenue.[6] Their net income was over $62 billion in the same year.[7] This explanation of how profiling works focuses on Meta, but the same fundamental method applies to other platforms and companies as well.

Through the information we actively give to Meta through our profiles, our activity on their platforms, and their ability to track our online activity when online elsewhere, it has a huge amount of data about their users' interests, behaviours and preferences. In the early days of Facebook, the platform really only knew what we chose to share about ourselves consciously and intentionally. Now, Meta can create detailed profiles of every user, which contain far more information than we might realise or actively choose to share. A data profile can include a huge range of different types of data including demographic data, interests and attitudes, behavioural data, geographical data, the kinds of devices and software used by the individual, and data derived from all of this such as someone's likely credit score.[8]

Every post on Facebook that a user interacts with, every link they click, person they are friends with and advert they engage with all build up the 'on-site' data which contributes to their Meta user profile. Add this to 'off-site' data supplied to Meta by data brokers, cookies, mobile IDs and the Meta Pixel, which can be embedded in any website to automatically send user data directly to Meta,[9] and Meta holds thousands of data points about every user.[10] [11]

Joining up this data and using algorithms to analyse it creates a hugely rich digital profile picture of each individual user. It includes both directly known characteristics, for example if you have included your age, location or gender on your user profile, as well as inferred or assumed characteristics, which advertisers and others can then use to learn all sorts of things about us. Meta has evolved its AI tools from relatively simple machine learning systems to a much more sophisticated, complex and powerful model of AI analysis and decision-making to fine-tune profiling, improve their ad targeting and bring in more revenue.

Algorithms look for patterns in data which suggest, for example, a person with these characteristics tends to like X, visit Y place, or spend money on Z things, thereby allowing advertisers to target their adverts at people who are most likely to be receptive to them.[12] Someone who reads a lot of articles about supercars, is a member of Facebook groups for car owners, and shares a lot of posts about their driving holidays in Europe is highly likely to be an enthusiastic car owner and driver, and respond to motoring adverts. Meta will be able to discern this even if the individual has never explicitly ticked a box to declare this.

---

5 For more detail on their advertising business model see Consent without Paying: Alternatives to Meta's surveillance advertising models, Open Rights Group, 2025

6 Statista, 2025, Annual advertising revenue of Meta Platforms worldwide from 2009 to 2024

7 Statista, 2025, Annual revenue and net income generated by Meta Platforms from 2007 to 2024

8 Buchi et al, 2019, The chilling effects of algorithmic profiling: Mapping the issues, Computer Law and Security Review 36

9 The Meta Pixel can harvest details from other websites such as the items someone has added to their online shopping cart, or their use of an online mental health service, which can in turn feed sensitive data to Facebook to aid their profiling

10 Which?, 2023, Are you still following me?

11 For more detail on the range of data sources used in digital advertising, see AWO for the European Union, 2023, Study on the impact of recent developments in digital advertising on privacy, publishers and advertisers

12 Buchi et al, 2019, The chilling effects of algorithmic profiling: Mapping the issues, Computer Law and Security Review 36

It is these digital profiles that are the basis of Meta's advertising offer. Advertisers that wish to place an advert on a Meta platform can select from an enormous list of user characteristics and interests to choose the type of audience they wish to see their advert. Facebook analyses user profiles using AI tools to select the users most likely to engage with the advert, click links and ultimately spend with the advertiser. They also use AI to analyse ad performance, tracking how many users, with which characteristics, click on them and convert into a customer. This is called ad targeting.

This is what Meta says about what determines the adverts you see on their platforms:

To decide which ads to show you, we use advertiser audience selections and the results of our ad auction[13] to determine the best ad to show you at a given point in time.

Our machine learning models are part of our ads delivery system that learns as it receives new information, without being explicitly programmed.

This allows the machine learning models to carry out tasks quickly and efficiently, like delivering ads that might be relevant to you.

The two main factors that we use to determine which ads to show you are:

- **Advertiser audience selection**
- **Results of our ad auction**

And about how they analyse ad performance:

As more people view an ad, share feedback or click through to make a purchase on an advertiser's website, our models get better at predicting the estimated action rate and ad quality of an ad.

Much of this targeting happens through inferred characteristics: the user profile does not necessarily include data which directly discloses an interest in the subject of the advert, but through the combination of all the individual datapoints an accurate assumption can often be made.[14] By identifying patterns within data, profiling makes 'probabilistic assumptions' about whether an individual belongs to a larger group, such as 'people likely to default on a loan' or 'people likely to purchase high-end skincare'.[15]

This is one of the primary sources of consumer unease about profiling: joining up multiple separate data points into a composite. We might be comfortable sharing one or two pieces of personal data with individual organisations, but when this is brought together into a detailed digital profile of ourselves, this not only can feel more unsettling, it can lead to more problematic outcomes.[16]

On paper, UK and European users of Meta platforms have the 'right to object'.[17] Under the UK GDPR, individuals have the right to object to the processing of their personal data in particular circumstances. In the case of direct marketing, this is an absolute right, i.e. the data processing organisation has no basis on which to turn down an individual's request. The right to object, where it applies, should be exercisable at any time. Meta contends that the direct marketing objection does not apply to them, in contrast to other sectors of the media where it is actively implemented. This not only gives them an unfair advantage, but also confuses the landscape for consumers, and means their rights are unevenly enforceable.

As well as the right to object, there are anti-discrimination and data processing laws which platforms must abide by. Direct and indirect discrimination, such as might result through targeted advertising, is unlawful on the basis of: age, disability, gender

13 Every time an advert is going to be placed on a Meta platform an ad auction is held: advertisers specify an amount they are willing to pay, and Meta uses algorithms to determine the most relevant audience(s) as well as the quality of the advert. More details on the Meta website

14 Derek E Bambauer, 2025, Target(ed) Advertising, UC Davis Law Review Vol. 58

15 Mann, M and Matzner, T, 2019, Challenging algorithmic profiling: The limits of data protection and anti-discrimination in responding to emergent discrimination. Big Data & Society, 6(2)

16 Derek E Bambauer, 2025, Target(ed) Advertising, UC Davis Law Review Vol. 58

17 ICO, no date, Right to object

reassignment, marriage or civil partnership (in employment), pregnancy and maternity, race, religion or belief, sex, or sexual orientation.

Under the GDPR, there is also 'special category data' which reveals someone's racial or ethnic origin; political opinions; religious or philosophical beliefs; trade union membership; genetic data; biometric data (where used for identification purposes); data concerning health; sex life; or data concerning a person's sexual orientation. Processing of special category data is restricted because collecting and using it has a higher risk of causing discrimination or other harms.[18]

PROXIES

The composite profiling and inference outlined above can lead to the 'creepy' feeling of being followed around the Internet by platforms and adverts that know a disturbing amount about us. It can also lead to discrimination or harm. It does not rely on capturing or processing highly personal and often legally protected data such as gender, age, sexual orientation or race, but it can be used to profile people just as effectively. Our browsing history, membership of Facebook groups and many other seemingly neutral or innocuous online activities can reveal sensitive information about us just as much as protected and special category data can.[19]

The combining and algorithmic sorting of data creates what are called proxies: where the platform or advertiser does not hold data that directly relates to a characteristic they wish to target, or where targeting (or exclusion) is unlawful, other characteristics can stand in for the desired one. Proxies can be created by training an algorithm with historical data about a user action or activity which is of interest to the platform or advertiser, such as clicking on an advert for a particular product. The algorithm can then 'learn' what other characteristics are associated with that user action, which when combined serve as proxies.[20]

Proxies mean that while Meta can and has banned ad targeting on the basis of protected characteristics, and the processing of special category data is prohibited, discriminatory outcomes can still occur. A 2020 experiment showed that creating proxies by bringing together two or more characteristics can significantly skew ad targeting towards or away from particular age groups or genders.[21] Because Meta's policies restrict targeting via individual characteristics, it is not clear how much these policies have an impact on targeting via proxies, which are composite categories and may not be made up of any special category or protected characteristic data at all. In many ways, proxies enable the replication of protected characteristics, meaning discrimination can still occur.

An advertiser may be prohibited from targeting an advert for housing at a particular age group for example, but could put together a range of different indicators to create an effective proxy: people who like a particular genre of music, and live in a particular area and post regularly about grandchildren are likely to be of an older age profile. An advertiser of high end products might want to target people with a lot of disposable income. This category may not be available to select directly, but they could use other characteristics such as 'goes on multiple foreign holidays' and 'interested in luxury goods' together with location information to create an audience likely to be receptive to their ads.

Proxies that reveal protected or sensitive characteristics can be created intentionally or unintentionally. An ad targeting algorithm might combine postcode data, browsing history and analysis of a user's Facebook friends and through this reveal protected and sensitive characteristics. Once created, proxies can also be used intentionally or unintentionally as a stand-in for protected characteristic data. An

18 ICO, 2024, What is special category data?

19 Sandra Wachter, 2019, Algorithmic bias within online behavioural advertising means public could be missing out; AWO for the European Union, 2023, Study on the impact of recent developments in digital advertising on privacy, publishers and advertisers

20 Barocas, Solon and Selbst, Andrew D., 2016, Big Data's Disparate Impact, 104 California Law Review 671

21 Giridhari Venkatadri, Alan Mislove, 2020, On the Potential for Discrimination via Composition

advertiser might set out to do this to get around legal restrictions, or it may happen by accident, even when efforts have been made to avoid it.[22] Discrimination achieved through proxies can therefore be invisible to consumers and advertisers.

An example of proxies in action is one of Meta's advertising tools, 'Lookalike Audiences', which enables advertisers to reach a new audience that 'looks like' i.e. shares attributes with their existing audience. It can enable targeting which is not legally allowed, because the amount of data it holds through its user profiling can accurately predict protected characteristics like race, sex or age.[23] For this reason it is not allowed in Meta adverts for housing, credit or housing opportunities, but could still be used to achieve discriminatory targeting in other ad categories.

HOW ELSE IS PROFILING USED?

This algorithmic profiling does not just happen on Meta platforms, it can take place across a wide range of products and services in the public and private sectors, from marketing to finance, policing to employment. Famously, profiling was used in political advertising in the case of Cambridge Analytica, a company which used ad targeting to try to manipulate US voters. They had received thousands of data points on millions of Facebook users, which they claimed could enable extremely detailed ad targeting by political parties.[24] For more on political advertising, see Open Rights Group's report on disinformation, division and fraud on Meta's platforms.[25]

WHY DO PLATFORMS AND ADVERTISERS USE PROFILING AND AD TARGETING?

Advertisers use profiling and ad targeting to maximise the chances that their online adverts are not only seen by the most relevant audiences, but also by people that are mostly likely to click on links and ultimately make a purchase. Of course, there is nothing new in this; marketing and advertising of all kinds aim to identify who is most likely to engage and purchase a product or service. The difference with digital marketing and advertising is the depth and breadth of data that it holds and uses about us as individuals.

This is how Meta describes the process of 'audience selection' that an advertiser can go through:

When creating an ad campaign, advertisers first choose their desired audience through our business tools.

Advertisers can create audiences based on age, location, interests and categories. For example, some information you provide us, combined with actions you take, might suggest to us that you're interested in something, such as cooking or fitness, or that you might be part of a larger group (called a category), such as a mobile user.

Advertisers can also use information they have about their audiences, such as a list of email addresses or people who've visited their website, to build a Custom Audience or a Lookalike Audience.

22 Barocas, Solon and Selbst, Andrew D., 2016, Big Data's Disparate Impact, 104 California Law Review 671

23 Buchi et al, 2019, The chilling effects of algorithmic profiling: Mapping the issues, Computer Law and Security Review 36

24 Amnesty International, 2019, 'The Great Hack': Cambridge Analytica is just the tip of the iceberg

25 Open Rights Group, 2025, Bad Ads: Targeted Disinformation, Division and Fraud on Meta's platforms

Note:

- **Advertisers can include teens in their audience based only on age and location.**
- **When advertisers show ads about credit, employment or housing opportunities, we limit the categories they can choose to create an audience.**

This is what the interface looks like for someone creating an audience for their advertising:

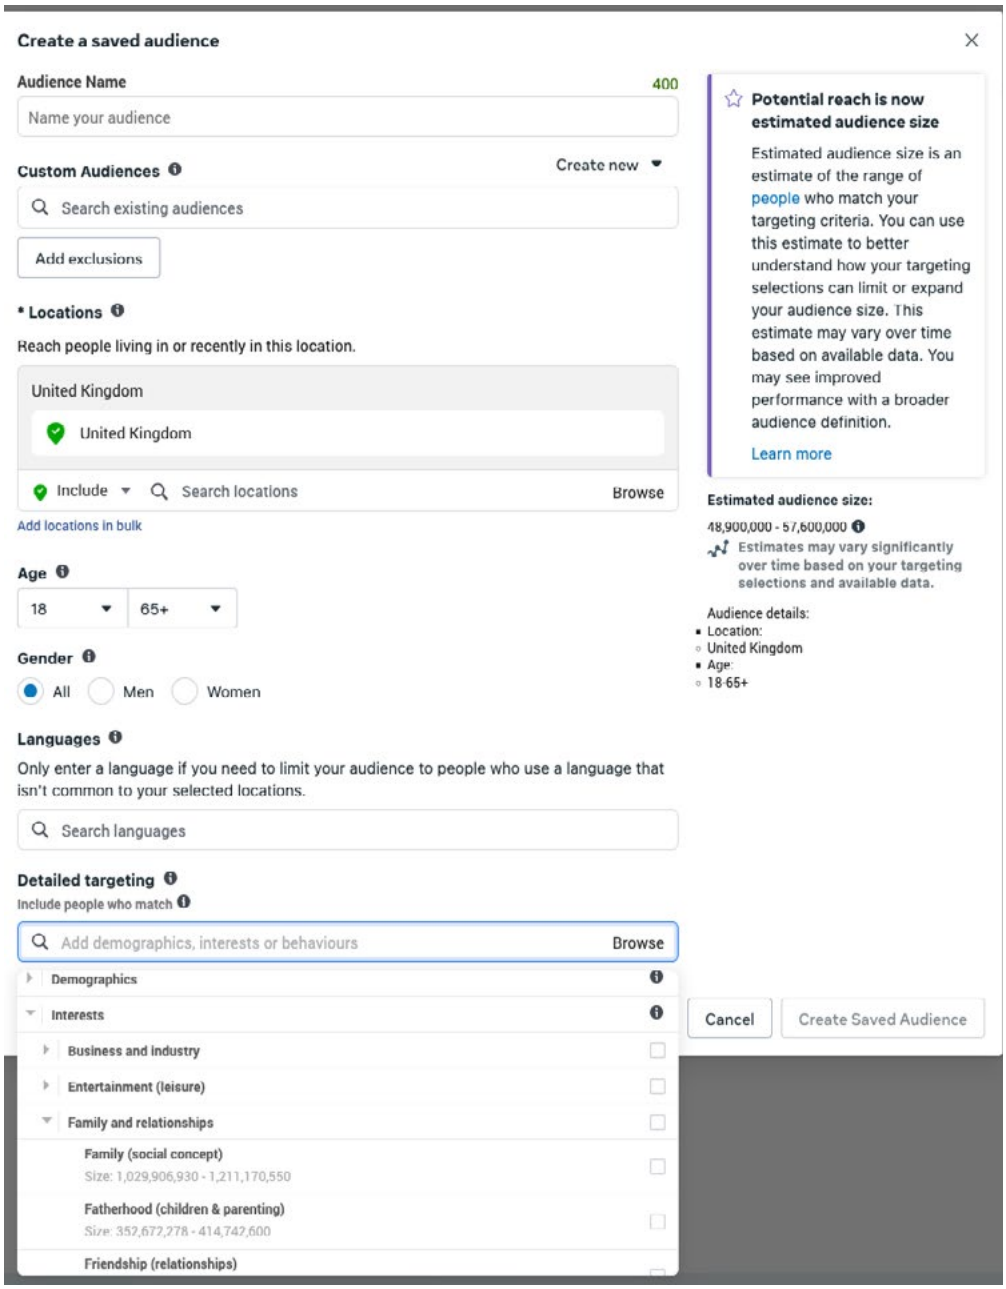

Figure 1: Screengrab from Meta's ad portal showing some of the options available to advertisers when choosing their ad audience

As far back as 2010 there was some evidence that targeted advertising improved the effectiveness of online advertising, with less tailored adverts resulting in less purchases than those targeted towards a specific audience.[26] A 2010 survey of advertising networks also found that customers who clicked on targeted adverts were twice as likely to buy as those who clicked on non-targeted ads, and sites that host targeted adverts were able to charge more for hosting them.[27]

In general, customers spend less time searching for a product when targeted marketing is used, and ad targeting can mean advertisers have to use less adverts to achieve the same results. However, if platforms can charge higher rates to host targeted adverts this may on balance not be a benefit for the advertiser.[28]

While evidence does suggest that targeted advertising is more effective than non-targeted, it is not overwhelming or entirely consistent, and effectiveness does appear to vary according to the quality and accuracy of the targeting (partly driven by the volume and quality of the data underpinning it), the products or services and company that is advertising.

It should also be noted that ad targeting does not have to be achieved through the type of data / behavioural profiling that Meta uses. Contextual advertising, serving ads based on the content that someone looks at rather than their personal information (see Section 8) can be just as effective a way to target ads.[29] Contextual advertisers like Kobler and Opt Out Tracking have even reported higher conversion rates than behavioural adverts.[30]

It is also worth remembering that platforms benefit significantly from the revenues generated by offering targeted advertising; some, like Facebook, rely on it for a vast majority of their income.[31] Platforms that host ads are incentivised to keep people on their platform for as long as possible: the more time people spend on Facebook the more ads they will see, the more advertisers will use it and pay for ad services. This can lead to a misalignment of incentives: profit versus user privacy or safety. The evidence provided by whistleblower Frances Haughen suggests that Meta achieves maximum user engagement by promoting harmful or divisive content, keeping this content circulating specifically to maximise the eyes on adverts and therefore their ad revenue.[32]

Meta's own research shows that young people attribute some of their mental health challenges to their use of Instagram, for example being shown ads and other content relating to weight loss can exacerbate or encourage disordered eating. But in order to maintain their business model, keeping as many users on their platforms for as long as possible is fundamental, and takes precedence over safety concerns.[33]

Advertisers may not wish to use profiling or ad targeting,[34] or advertise on Meta platforms at all, but because most people now spend the majority of their time online on platforms, firms that wish to advertise online have little choice. There is evidence that not all advertisers are happy with the service or results they get from Meta: a class action lawsuit in the US alleges that Meta overcharged advertisers based on artificially inflated numbers of potential ad viewers. If the lawsuit is successful potential damages could reach over $7 billion. This pales in comparison to Meta's ad revenue however, which was over $100 billion in 2024 alone.[35] As long as Meta platforms retain most of their users and the majority of them are not able to opt out of profiling and ad targeting, advertisers are stuck.

---

26 Avi Goldfarb, Catherine E Tucker, 2010, Privacy Regulation and Online Advertising, Management Science Vol 57 Issue 1

27 Forbes, 2010, Behavioral Ads Offer A Windfall For Marketers, Publishers

28 Yan Lau for the Federal Trade Commission, 2020, A brief primer on the economics of targeted advertising

29 Brave, 2020, Update (Six Months of Data): lessons for growing publisher revenue by removing 3rd party tracking

30 CITATION TO ADD

31 Derek E Bambauer, 2025, Target(ed) Advertising, UC Davis Law Review Vol. 58

32 CBS News, 2021, Whistleblower: Facebook is misleading the public on progress against hate speech, violence, misinformation

33 Privacy International, 2021, To Reduce Facebook's Harms to Teens, Target its Data-Hungry Business Model

34 AWO for the European Union, 2023, Study on the impact of recent developments in digital advertising on privacy, publishers and advertisers

35 Digwatch, 2025, Meta faces massive advertiser lawsuit after US Supreme Court decision

WHAT DO CONSUMERS THINK?

Consumers often hold a variety of potentially contradictory feelings about targeted advertising, simultaneously finding completely untargeted adverts annoying and also not wanting to be subject to intensive digital surveillance and data collection.

The majority of users may not be aware of the full range of ways that companies collect and use their personal data, and tend to have limited understanding of how online targeting works.[36] As consumers' understanding grows, their level of comfort with targeting decreases: in a survey of UK Internet users, the proportion saying that ad targeting in exchange for using a website for free is acceptable almost halved when they were given more information about adtech and how it works.[37] In another survey, less than a fifth of UK Internet users said they are happy with their data being used in exchange for a free or personalised service. However, they seem resigned to the trade-off, feeling disempowered by a lack of choice: either accept the ads or don't use the website.[38]

Consumer concerns are particularly notable when they are asked about sensitive personal information. A survey conducted in France and Germany in 2021 found that many people were uncomfortable with being digitally profiled and targeted with adverts. Nearly 90% of respondents said that adverts should not be allowed to be targeted by information about income or health, with sexual orientation, religious views and race or ethnicity all scoring around 80%. The majority (57%) said they did not want to see any targeted advertising at all, with only 11% saying they were happy with their personal data being used to target them with adverts.[39]

Which organisation is using online ad targeting has a big effect on how comfortable consumers feel about it: over 80% of people in a survey said they would find it acceptable for the NHS to target people with ads for the flu jab. This compares to less than 20% feeling that targeting gambling ads towards those most likely to place a bet was acceptable.[40]

On balance, many enjoy using platforms such as Facebook for free and put up with data profiling and ad targeting, despite their concerns about privacy.[41] It is likely that many users are simply not aware of the scale of data that Meta holds on them, or the extent to which it can be used to build sophisticated profiles that go way beyond the data we knowingly share.

---

36 Ofcom, 2022, Adults' Media Use and Attitudes Report; Centre for Data Ethics and Innovation, 2020, Online targeting: Final report and recommendations

37 ICO / Ofcom, 2019, Adtech Market Research Report

38 Ofcom, 2022, Adults' Media Use and Attitudes Report

39 Global Witness, 2021, Do people really want personalised ads online?

40 Centre for Data Ethics and Innovation, 2020, Online targeting: Final report and recommendations

41 Derek E Bambauer, 2025, Target(ed) Advertising, UC Davis Law Review Vol. 58; Yan Lau for the Federal Trade Commission, 2020, A brief primer on the economics of targeted advertising

WHY SHOULD WE BE CONCERNED ABOUT DATA PROFILING AND AD TARGETING?

DATA PRIVACY AND TRANSPARENCY

One significant factor that leads to concerns among consumers and experts is how opaque the profiling and targeting processes and outcomes are. There is a huge imbalance in knowledge and understanding between users whose personal data is the fuel that powers ad targeting, and Meta which makes billions every year from advertisers. As a consumer or Internet user it is almost impossible to fully understand the extent to which our data is captured, analysed and used to generate profiles for use in ad targeting.

The inferences created through profiling may be shared directly with a user, through a recommendation to something they may like based on their profile; shared with advertisers and then replayed to users through the advertisements they see; or used for purposes which are never directly revealed to users, such as being shared with other platforms and companies via data brokers.[42]

Meta provides some information to users about the reasons they see particular adverts:

We use machine learning models that mirror our ads delivery models to provide the insights found in "Why am I seeing this ad?". These models create and surface how certain activity both on and off Meta technologies contributed to Facebook showing you a particular ad.

For example, liking a post, clicking on an ad or activity on websites, apps and products sent to us from advertisers. Meta's use of your activity off Meta technologies depends on your ad settings. While these insights don't represent all of the factors contributing to delivering an ad, we use them to provide you with an informative and understandable view into some of the factors that contributed to why you saw an ad."

Facebook has a series of settings which can be used to put some limits on the data held on a user's profile. Users are tagged with a list of 'Interests' which might be anything from construction to nursing, video games to dating, as well as 'Categories', which can include things like 'newlywed', or 'close friends with expats'. If the user is aware of and can find these lists, it is possible to remove them from their profile.[43] There are several other actions that users can take[44]:

- **Customise your ad preferences in Accounts Centre to influence the ads you see. You can also update your ad settings to choose whether we show you ads based on your activity on apps and websites off Meta technologies. Note: you must be logged in to change your ad preferences and ad settings.**
- **Hide an ad that isn't interesting or useful to you or review why you're seeing a particular ad.**
- **Update your profile settings and walk through Privacy Check-up to make sure that you're sharing your information with who you want.**
- **Visit Access your information to see and manage your Facebook information, or Download your Facebook information for review.**
- **Learn more about why you're seeing an ad with the "Why am I seeing this ad?" tool that uses advertiser selections and machine learning models to show you more information about the ads that you see."Why am I seeing this ad?" allows you to tap on ads in News Feed, get context on why they're appearing and take action to further personalise what you see. Learn more about why you may be seeing an ad.**
- **Review your activity off Meta technologies. Review and manage the information that businesses and organisations share with us about your interactions, such as visiting their apps or websites.**

42 Buchi et al, 2019, The chilling effects of algorithmic profiling: Mapping the issues, Computer Law and Security Review 36

43 Wired, 2018, Facebook's Targeted Ads Are More Complex Than It Lets On

44 How Facebook ads use machine learning

Turning off options like 'Ads based on data from partners' removes a lot of the options for targeted advertising, but that doesn't stop them entirely. The information you share without filling in any personal details in your profile is enough to receive targeted ads. Just by recording the links you click on, and posts you like or even linger over longer than others, Facebook can build a detailed profile of you, your interests, and your online behaviors and preferences. While you can go through and delete each of these, few people will know this or have the time to go through each individual interest and remove it from their profile.[45]

Taking all of these actions is not a quick task, and even when these controls are used, they do not appear to stop unwanted adverts or topics from appearing in someone's Facebook feed. A case study developed by the Panoptykon Foundation showed that despite using Facebook's ad control tools to opt out of health-related and parenting content and ads, a user continued to be shown hundreds of such ads, causing her significant anxiety and distress. Initially, when the profiling and ad controls were activated the number of unwanted adverts did decrease, but over time rose again, and new health and parenting 'interest categories' were attached to the user's profile, without their knowledge.[46]

Being included in a category which leads to discrimination or harm is often entirely invisible to the individual; it happens automatically, by algorithm or other AI, in a 'black box' which is hidden from consumers. Understanding what data is being processed, to what end, and what discrimination may have occurred is almost impossible for an individual. The highly automated nature of ad targeting makes it hard for people to know if they have been targeted, using what data, and what ads they may have seen or not seen as a result.[47]

There is currently no consistent, lasting and comprehensive way for a user of a Meta platform in the UK to entirety opt out of being profiled and shown targeted advertising. Even though Meta agreed to stop profiling human rights campaigner Tanya O'Carroll after she sued them, the company have not as yet provided a mechanism for allowing other users to opt out of profiling.

DISCRIMINATION AND HARMS

All major platforms that make revenue from digital advertising rely on enormous webs of data collection to be able to sell advertising space for the highest price. A Wired investigation revealed the extent of the industry that has grown up to deliver digital advertising, and the intimate information it holds about us. They found data brokers offering advertisers lists of thousands of US residents in categories such as likelihood of being in financial debt, enjoying gambling, and a wide range of medical conditions including asthma, cardiovascular conditions and diabetes.

These 'audience segments' can be used by advertisers to fine tune who to show their adverts to, opening up the possibility of intrusive advertising to people who did not wish to disclose a medical condition, gambling ads shown to people at risk of harmful or addictive gambling behaviour, or high interest loans offered to people in financial distress. Each user is given a unique ID which, if combined with other datasets, can be de-anonymised, risking personal privacy and safety.[48]

As with many instances of tech-enabled discrimination, harms and discrimination caused by ad targeting tend to affect people and groups that already experience exclusion and other challenges. The following section describes a wide range of documented problems caused by ad targeting, in which discrimination on the basis of race, ethnicity, age or gender has occurred, as well as harm caused by the use of sensitive personal information in ad targeting. These instances tend to replicate and intensify existing prejudices and discriminatory practices, be they in access to financial services, employment or housing, or targeting people who have challenges such as gambling addiction or mental health issues.

45 Digitaltrends, 2019, I turned off Facebook's ad trackers, and the ads only got more personalised

46 Panoptykon Foundation, 2021, Algorithms of trauma: new case study shows that Facebook doesn't give users real control over disturbing surveillance ads

47 European Digital Rights, 2021, How online ads discriminate: Unequal harms of online advertising in Europe

48 Wired, 2025, Google ad-tech users can target national security 'decision makers' and people with chronic disease

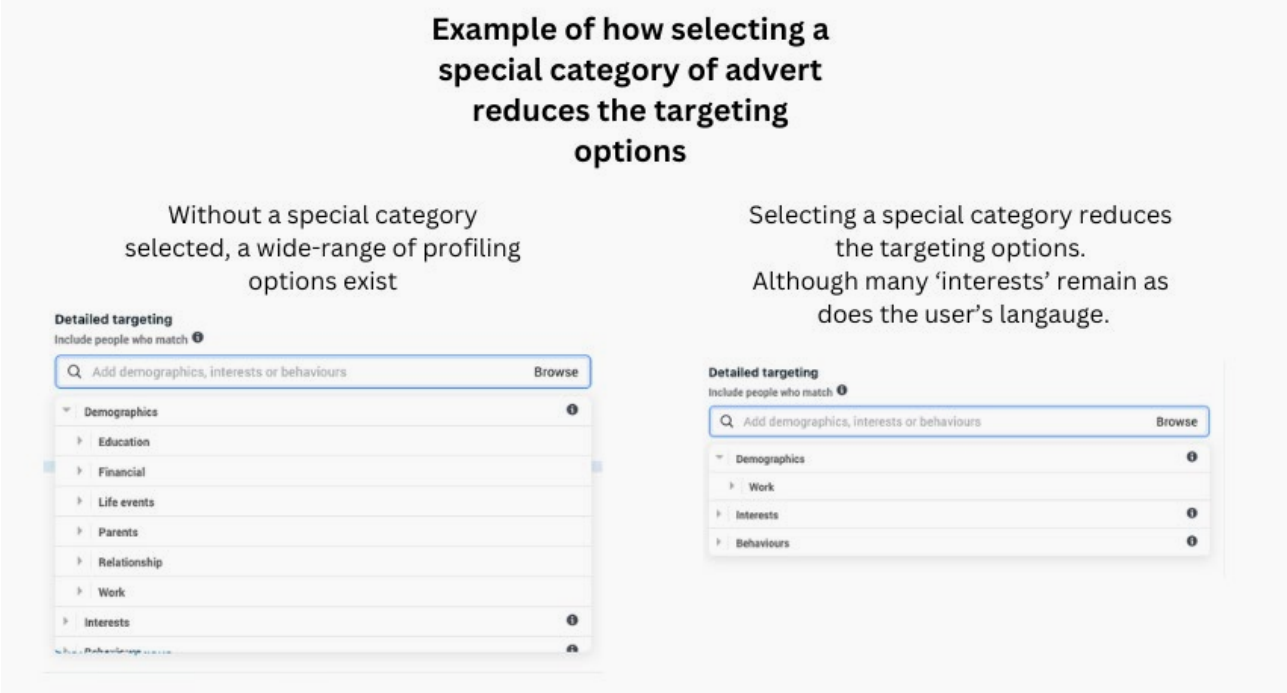

Figure 2: Screengrabs from Meta's ad portal showing the difference in targeting options between regular ads and 'special ad categories'

There may be inherent risks in how data profiling works which contribute to the discriminatory effects of ad targeting. A one-off visit to a website may not reveal much about our interests or preferences, but repeated and regular engagement with a particular topic or site will create more data and patterns of behaviour which contribute much more significantly to our data profiles. If our data profiles are largely influenced by the things we do the most, including our addictions, vulnerabilities, or chronic concerns such as health, these sensitive issues become central to how we are targeted, making the replication and intensification of harms more likely.

In some cases, different laws and rules apply in the US to Europe and the UK and other territories, however the end results are broadly similar. In Europe and the UK local anti-discrimination laws and GDPR apply, which in theory means that ads should not be discriminatory, that advertisers need explicit consent to process our data for the purposes of ad targeting, and that 'special category data' cannot be processed.[49] In practice, just as in the US, advertisers can reach their desired audiences through proxies, and proxy-enabled targeting can also still happen unintentionally. There are examples of ongoing discrimination and harm in the UK and Europe despite GDPR rules and anti-discrimination law,[50] as well as in the US.

PROFILING AND TARGETING RESTRICTIONS INTRODUCED BY META

Prior to 2019, when a settlement was made between Facebook and a group of US civil rights organisations, numerous examples of direct discrimination were found in Facebook advertising. Advertisers were able to explicitly target or exclude potential audiences on the basis of protected characteristics such as sex, age or race. The settlement meant that Facebook would no longer allow advertisers to directly discriminate in the advertising

49 Special Category Data in the UK GDPR is data which reveals someone's racial or ethnic origin; political opinions; religious or philosophical beliefs; trade union membership; genetic data; biometric data (where used for identification purposes); data concerning health; sex life; and data concerning a person's sexual orientation.

50 European Digital Rights, 2021, How online ads discriminate: Unequal harms of online advertising in Europe

of housing, employment or credit.[51] These products and services are now deemed 'special ad categories', which means that targeting options are more restricted.

Meta guidelines also now say that all ads must not discriminate against people based on 'personal attributes such as race, ethnicity, colour, national origin, religion, age, sex, sexual orientation, gender identity, family status, disability, medical or genetic condition'. Advertisers are not allowed to use audience selection tools to either target people for advertising or exclude people from seeing adverts in ways that result in discrimination on these grounds.

Meta provides guidance to advertisers in the use of the Meta Pixel as well as its own ad targeting tools, which includes the disclaimer below, suggesting that Pixel users should not share data in various sensitive categories:

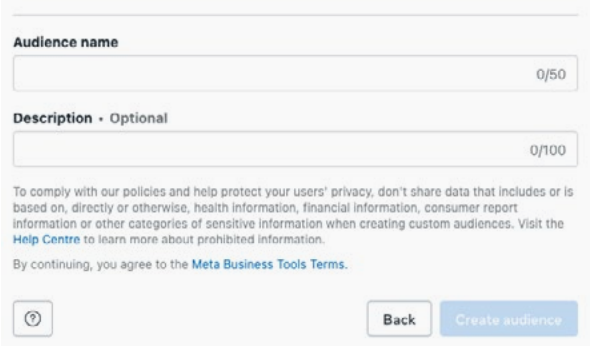

Figure 3: Screengrab from Meta's ad portal showing a disclaimer to companies setting up a Meta Pixel on their own websites

While a useful reminder not to share some categories of sensitive data, this puts the onus on the third party and puts Meta at arms length from any misuse, and does not necessarily stop these categories of data being used by unscrupulous advertisers. Although Meta has tools in place to attempt to stop this kind of data being used in targeting they have also admitted that their systems are not foolproof. The wording of the disclaimer also leaves room for interpretation: "data that includes or is based on, directly or otherwise" does not give a clear and definitive picture of what is and is not allowed, and could be circumvented quite easily with proxies.

THE PERSISTENCE OF DISCRIMINATION AND HARM

Despite these changes, discriminatory and harmful advertising still appears on Meta platforms, in part because the overall aims of profiling and ad targeting are misaligned with fairness.

Discrimination and harm, particularly resulting from proxy characteristics, are an unavoidable part of the adtech and platform ecosystem. The technology of platforms like Facebook is sophisticated enough to infer sensitive details about us and our lives from non-sensitive data, rendering legal or regulatory prohibition against discrimination and harm essentially useless.[52] Advertisers want to reach audiences who are most likely to engage with their ads and spend money with them, so Meta's priority is to facilitate this with its ad systems.

The data Meta receives from 'off-site' activity including the Meta Pixel, combined with the data users create through their 'on-site' activity, feeds Meta's advertising algorithms and other AI tools, which then automatically look for characteristics to improve ad engagement and follow through. Meta's systems 'learn' which kinds of users are most likely to engage with which types of adverts and subsequently prioritises them. Through the use of proxies this can occur even if it is in direct contradiction of efforts or laws to maintain equal access to services, jobs or goods. Fairness and equality of access are not the primary considerations.[53]

As Figure 2 screengrab shows, even when a 'special ad category' is selected, which should restrict ad targeting, a number of 'Interests' are available which could

51 European Digital Rights, 2021, How online ads discriminate: Unequal harms of online advertising in Europe

52 Derek E Bambauer, 2025, Target(ed) Advertising, UC Davis Law Review Vol. 58

53 MIT Technology Review, 2019, Facebook's ad-serving algorithm discriminates by gender and race, MIT Technology Review

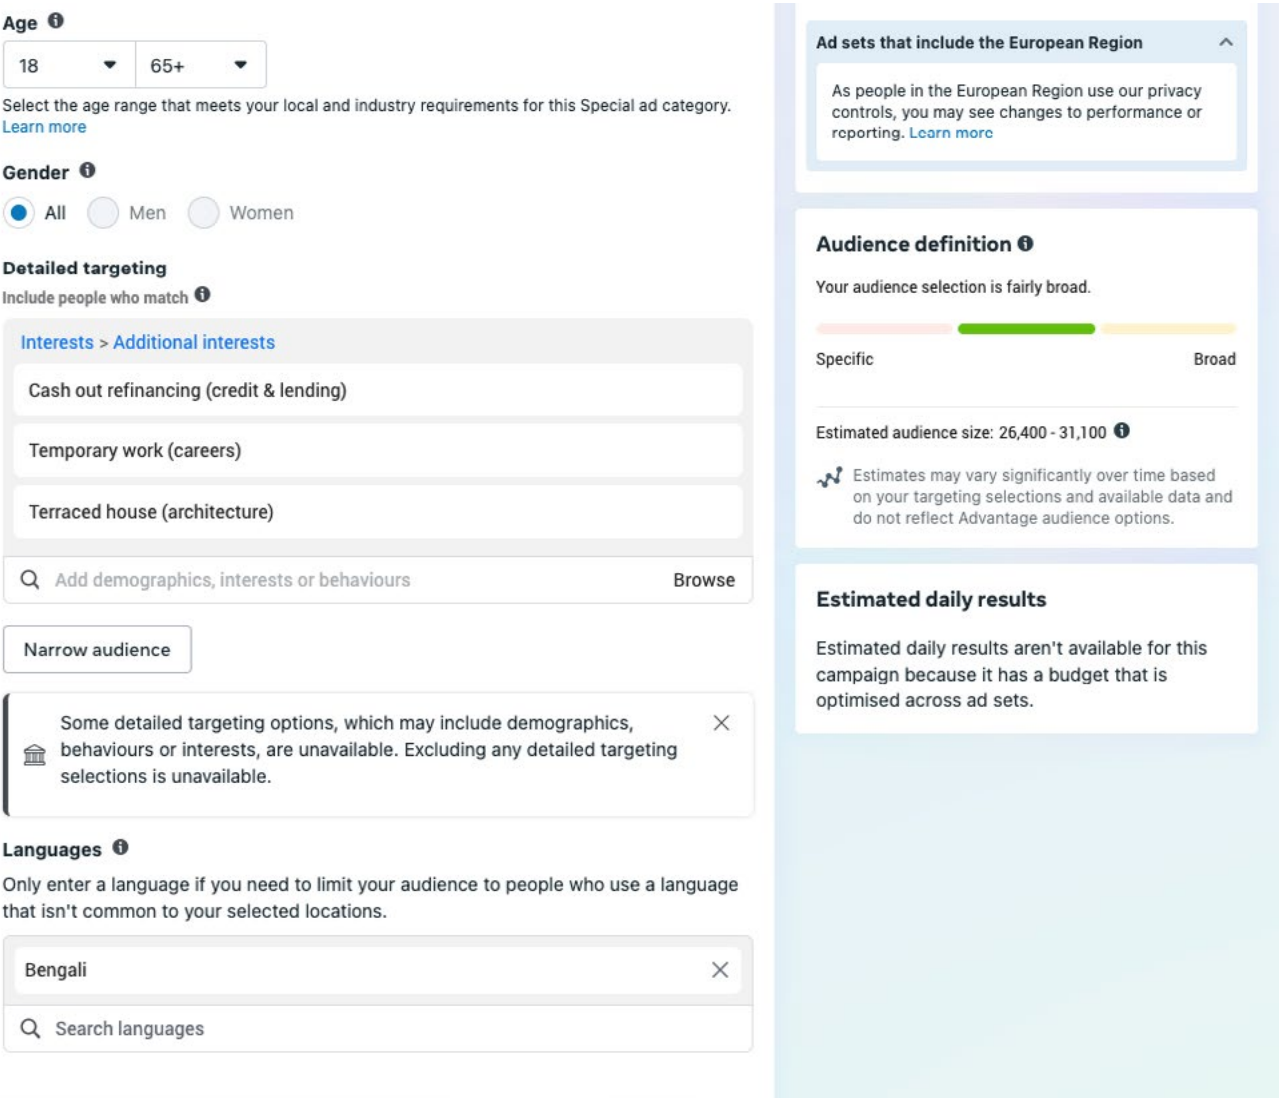

Figure 4: screengrab from Meta's advertising portal showing audience selection options when a 'special ad category' of Financial Products and Services has been selected

indicate someone on a lower income, as well as the ability to select an audience language which could easily act as a proxy for race, ethnicity, religion or national origin. In combination, these selections could very well enable an advertiser to unfairly or unlawfully target an already disadvantaged or marginalised group.

The following sections document a range of recorded discrimination and harms which have occurred as a result of profiling and ad targeting, primarily by Meta.

EMPLOYMENT

A ground-breaking investigation into racial targeting of job adverts in the US was carried out by Julia Angwin and Terry Parris Jr of ProPublica in 2016. Advertisers were at the time allowed to select the types of users they wanted their ads shown to in ways that breached discrimination law. A category called 'Ethnic Affinities' could be used to exclude racial groups in categories including (as described on the Facebook advertising portal) African American, Asian American and Hispanic. It was illegal to discriminate in both employment and housing advertising, and yet Facebook's ad portal enabled advertisers to do just that. Facebook said at the time that the categories were not actually sorting people by race, as they did not ask users for that information,

but that it was judged on the basis of pages and posts that users engaged with or liked.[54]

In 2020, Algorithm Watch carried out an experiment to see if Facebook was targeting job adverts in discriminatory ways in Europe. They placed adverts for machine learning developers, truck drivers, hairdressers, childcare workers, legal counsels and nurses, linking to real vacancies in five European countries. They did not specify that the adverts should be targeted at any groups of users in particular. Despite this, Facebook's ad algorithms still targeted the adverts significantly towards men or women: the truck driver advert in Germany was shown to over 4,500 men but under 400 women. The childcare worker ad was seen by nearly 6,500 women and only around 250 men. Facebook's algorithm had calculated that these people were the most likely to engage with the job adverts, regardless of the fact that Algorithm Watch had not asked for any gender-based targeting to be used. From their experiment, it seemed that one of the factors used in selecting the best audience was the images used in the adverts.[55]

This was found again in a 2021 study; an advert for a traditionally 'male' role which used an image containing men was more likely to be shown to men than women. The choice of image used to accompany an advert appears to be a way to intentionally or unintentionally skew who is more likely to see it.[56]

In 2021, Global Witness investigated age and gender discrimination in Facebook job adverts in the UK. They found that an advert for jobs at Facebook itself was predominantly served to men between the ages of 25 and 34, with only 3% of people seeing it being over 55 despite that age group making up nearly 20% of Facebook users. When the researchers posted their own ads, with no specific requirements for them to be shown to men or women, they were shown to heavily gendered audiences: 96% of people shown the mechanic jobs were men, and 95% seeing the nursery nurse job were women. The only criteria that the researchers used was to use Facebook's option to serve the adverts to 'people who are most likely to click on them'. This was potentially in breach of the UK's equality and data protection laws.[57]

This still appeared to be the case in Europe in 2023, when further research by Global Witness indicated that European users of Facebook were still being served adverts disproportionately based on their gender. Adverts for mechanics were mostly shown to men, and those for pre-school teachers mostly to women. The research was carried out in France and the Netherlands, and complaints were filed against Meta in both countries. Global Witness also carried out similar experiments in India, South Africa and Ireland and found similar problems in each country. The researchers took out job adverts linking to real vacancies, and did not specify any gender-related preferences in who the adverts were shown to, and yet they were split along clearly gendered lines.[58]

GAMBLING

The Meta Pixel, which is embedded in non-Meta websites and sends data about visitors back to Meta, has been sending data about people visiting gambling websites, without their consent. This can then be used by Meta in their ad targeting, potentially meaning people with problematic gambling habits could be flooded with gambling adverts on Facebook or Instagram. The Observer tested 150 gambling websites to check if they were all abiding by the requirement to obtain explicit permission from users to share data with Meta, and found many were not doing so. The data was transferred from the gambling sites to Meta as soon as someone opened the webpage, without giving them the chance to opt in or out of data sharing.

54 Julia Angwin, Terry Parris Jr, ProPublica, 2016, Facebook Lets Advertisers Exclude Users by Race

55 Nicolas Kayser-Bril, Algorithm Watch, 2020, Automated discrimination: Facebook uses gross stereotypes to optimize ad delivery

56 Adam Grybowski, Princeton School of Public and International Affairs, 2023, Despite progress, researchers find more potential discrimination in Facebook ads

57 Global Witness, 2021, How Facebook's ad targeting may be in breach of UK equality and data protection laws

58 CNN, 2023, People are missing out on job opportunities on Facebook because of gender, research suggests; Global Witness, 2023, New evidence of Facebook's sexist algorithm

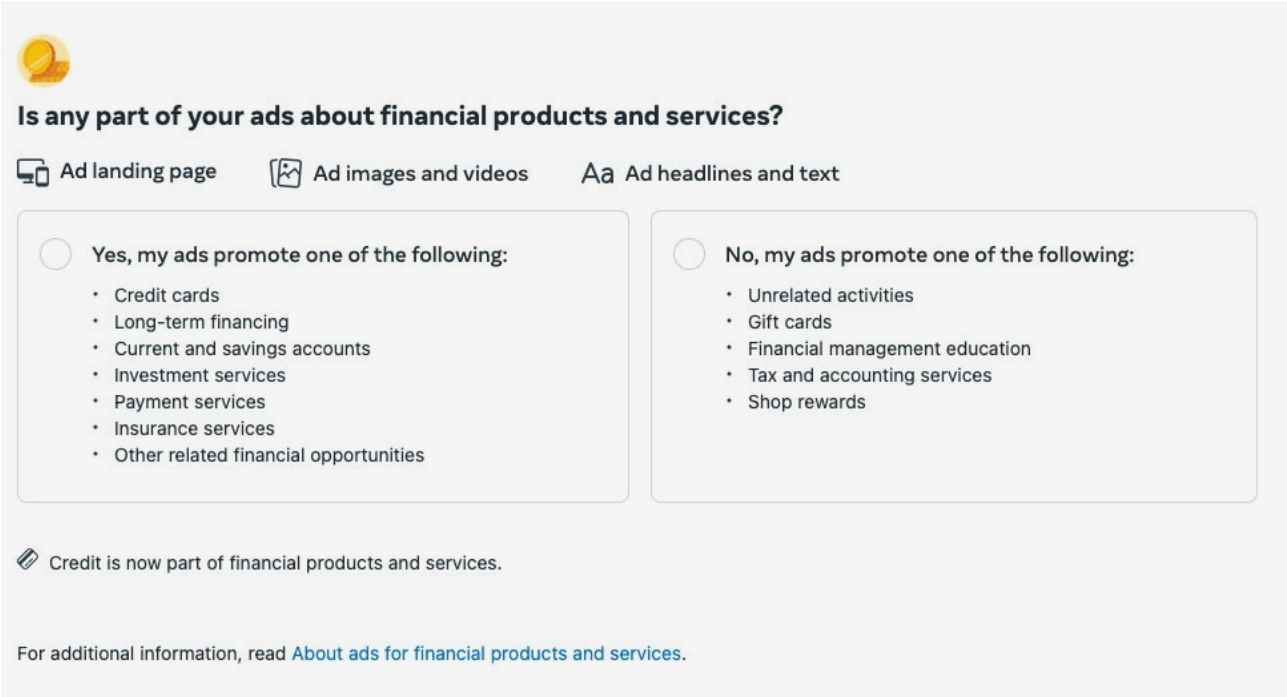

Figure 5: Screengrab from Facebook advertising portal showing the list of products and services which count as 'financial products and services'

Following the visits to the gambling sites, a user's Facebook page was filled with gambling adverts, highly likely to be as a result of the Meta Pixel data transfer. The adverts were not just from the sites that had shared data unlawfully, but a whole range of others as well. Some of the sites that shared data provided extremely detailed information, down to the areas of the site that had been clicked. This all means that someone who might not engage with any gambling content on Facebook, may never have actually placed a bet online or used any online gambling games can still find their feed full of gambling adverts. For someone trying to avoid online gambling this proliferation of adverts could cause serious harm.[59]

An investigation for UK charity Clean Up Gambling uncovered the extremely detailed web of data that flows through the online gambling industry. Anyone using an online gambling site has huge amounts of data captured about their online behaviour, which games they play, even how much they bet. The investigation visited 37 gambling sites, which led to over 2,000 transmissions of different pieces of data to 44 different companies, including Facebook and Google.

The data included the potential value to a gambling company if a lapsed user returned to its site, clearly incentivising the company to find ways to get them back as a player. Knowing how much a user is spending on what games, gives gambling sites exactly the right information to encourage users to gamble more or return to their sites. Some of this information is received by Meta, enhancing user profiling and making it more valuable to advertisers.[60]

The UK's Information Commissioner's Office (ICO) decided that one of the companies, Sky Bet, had breached the UK GDPR and reprimanded them for this unlawful conduct.[61] This is a rare, if limited, example of a regulator stepping in to take action against breaches of our data rights by platforms. The case also went to the High Court, which found in favour of an individual ex-gambling addict who

59 The Observer, 2025, Revealed: gambling firms secretly sharing users' data with Facebook without permission

60 Cracked Labs, for Clean Up Gambling, 2022, Digital profiling in the online gambling industry

61 AWO, 2024, Sky Bet brands have breached UK GDPR

asserted that his data was being used in ways he had not consented to, and which had the effect of encouraging him to gamble more.[62]

Gambling is not one of the products or services within the restricted 'Financial Products and Services' category, which limits the extent to which advertisers can target users:

Advertisers who want to run gambling ads do have to apply to Meta for specific permission, which applies only in the geographic location they have been approved in. Local laws about gambling advertising must be followed, and gambling ads are not allowed to be targeted at under 18s.

Notably, advertisers can run ads for 'social casino' and 'free-to-play' games without written permission from Meta. These are games which do not pay out prizes of monetary value, and may initially be free by virtue of a fixed number of in-game credits given to the user when they sign up. After the free credits are used up, players spend real money on new game credits and other in-app purchases, so while the prizes may not be in any legal tender, users still spend their own money on playing. Because of the restriction on cash or cash-equivalent prizes, players can never 'cash out'; their money is locked into the social casino game forever, even if they win a significant prize. Facebook is one of the most common platforms for people to find and play social casino games.

Gambling experts say despite not handing out cash prizes, social casinos may be as addictive as 'real' gambling, and that problem gamblers are still at risk of problematic use of these games.[63] Many social casino games also have in-game advertising and extensive data collection, further fuelling the profiling and ad targeting machine.[64] There is evidence that social casino companies are using Facebook user data to target their highest-spending users to keep them using their games and spending their money on game credits.[65] This disparity between 'real' gambling and social

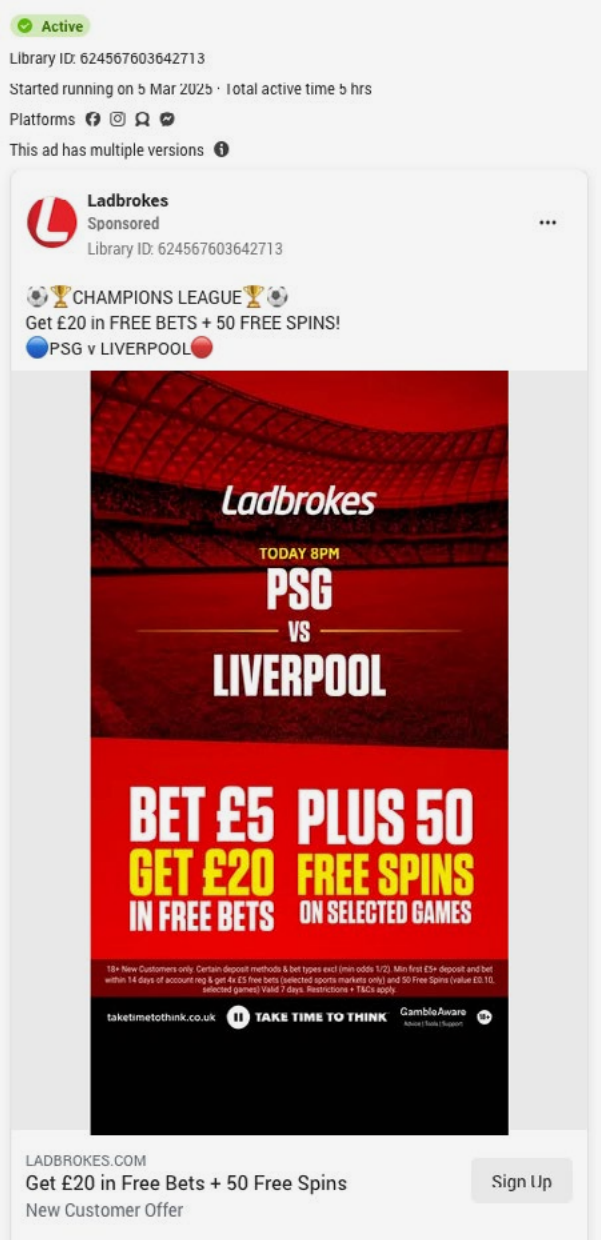

Figure 6: Screengrab from Meta's ad library showing where and when a gambling ad was shown on Meta platforms

gambling seems rife for abuse by gambling companies, allowing them to advertise harmful games without needing the same permissions, and also creates another source of data for profiling and targeting, which as discussed above can cause serious harm.

62 AWO, 2023, Landmark High Court ruling on GDPR consent to profiling and targeting in RTM v Bonne Terre - Analysis

63 ABC News, 2022, What are social casino games, and why do people become addicted?; Gainsbury et al, 2017, Virtual addictions: An examination of problematic social casino game use among at-risk gamblers

64 Business of Apps, 2024, Social gaming and big tech can work together for better data security

65 PBS, 2019, How social casinos leverage Facebook user data to target vulnerable gamblers

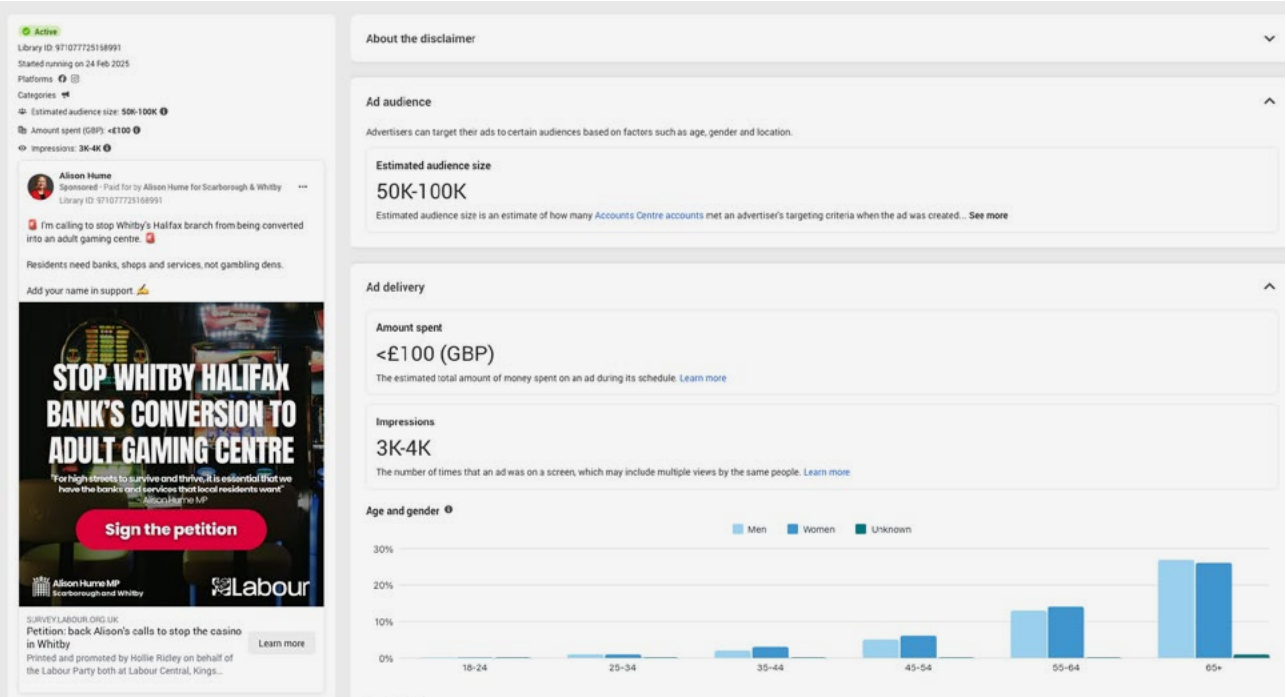

Figure 7: Screengrab from Meta's ad library, showing a more detailed breakdown of the audience who were shown a political ad campaigning against an adult gaming centre

Ironically, taking out an ad to push for political action on problem gambling has more conditions and restrictions attached than gambling ads: political ads on Meta are subject to more rigorous transparency requirements than regular ads, and political advertisers have less options available to them to target their ads. Meta's ad library shows the breakdown by age and gender of users who have been shown political ads, information which is not available for mainstream ads such as gambling:

CHILDREN AND YOUNG PEOPLE

When using Meta's ad portal, users under 18 years old can be selected as an audience, but the level of targeting is highly restricted. Selecting an audience with users under 18 means the choice to target by gender, demographics, interests, behaviours and language are all unavailable:

From 2021, advertisers were restricted in their ability to target teenagers on Meta platforms based on their use of other platforms or apps. In 2023, new controls were introduced to allow teenage users more control over the ads they see, and to limit advertisers' ability to target teenagers based on their gender.[66] Despite these restrictions, social media platforms make enormous amounts of money from the revenue they generate selling advertising space which is seen by young people. A 2023 study in the US estimated that Facebook, Instagram, Snapchat, TikTok, X and Youtube derived over $10 billion from adverts seen by under 18s in the US in one year.[67]

66 Cesar Cadenas, TechRadar, 2023, Meta's new ad policy further protects teen privacy and tackles discrimination

67 Harvard TH Chan School of Public Health, 2023, Social media platforms generate billions in annual ad revenue from US youth

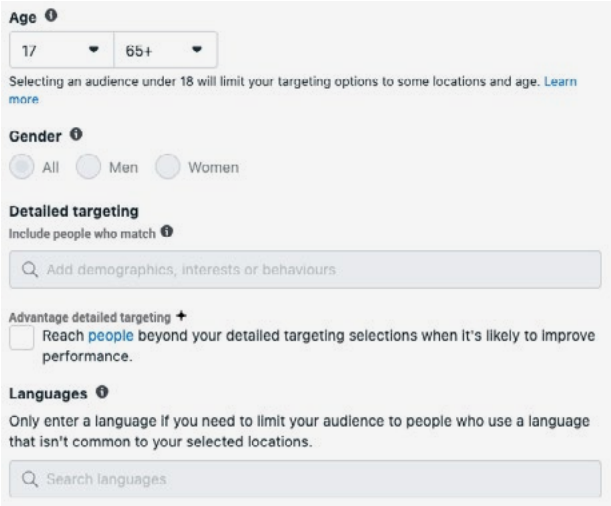

Figure 8: Screengrab from Meta's ad portal showing audience selection options unavailable for an audience that includes under 18s

In 2021, an American transparency group was able to gain approval for ads, targeted at children aged under 18, for products related to illegal drugs, eating disorders and dating services. The Tech Transparency Project had adverts approved with target audiences of hundreds of thousands of teenage Facebook users for alcohol, gambling, smoking, pharmaceutical drugs and weight loss products, within hours of submitting them. They selected from a range of 'interest categories' which Facebook provided, including 'Extreme weight loss', 'Pharmaceutical industry' and 'Online Gambling'. Within these categories they were able to select audiences aged 13-17 years old.[68]

They repeated the experiment in 2024, and once again received approval for ads promoting prescription drug abuse, dangerous diet tips, alcohol, vaping, dating and gambling to 13-17 year olds. This time, they created the adverts' imagery using Facebook's own GenAI tools. All of these advertisements would have violated Facebook's own policies on advertising to under 18s and /or advertising harmful products.[69]

In 2024, Meta and Google were accused of deliberately targeting 13-17 year olds with Instagram adverts, despite Google banning ad targeting to under 18s. Google apparently used proxy indicators of age such as online activity and app downloads to pinpoint 'with a high degree of confidence' users who were likely to be in their teens. The feature was allegedly rolled out in Canada and then the US, with plans to expand it internationally and add in adverts for Facebook. Google has cancelled the campaign and stresses that it does not allow targeted advertising to under 18s, but this reinforces the power of proxies to circumvent the rules, and platforms' continued interest in reaching younger Internet users.[70]

CREDIT AND FINANCE

There are restrictions on advertising financial products on Facebook, as they are considered a 'special ad' category. This means they cannot be targeted to users on the basis of age, gender or postcode, and tools like Lookalike Audiences cannot be used to find new ad audiences that match the characteristics of any existing audiences. Facebook banned the advertising of payday loans in 2015,[71] and the targeting restrictions, which apply in the UK and most of Europe as well as the US, were brought in in early 2025.[72]

Prior to this change there have been multiple examples of financial products including credit and financing being advertised to different groups in seemingly discriminatory ways. A 2020 study found that men were more likely to be shown ads for credit and financing, whereas women were more likely to see ads for debt relief. In the US, women are less likely to be offered finance and credit, and when they are, tend to receive worse terms than men.[73] In 2021, US news site The Markup found several examples of finance companies advertising financial products on

68 Tech Transparency Project, 2021, Pills, Cocktails and Anorexia: Facebook allows harmful ads to target teens

69 Tech Transparency Project, 2024, Meta approves harmful teen ads with images from its own AI tool

70 The Verge, 2024, Meta and Google secretly target minors on YouTube with Instagram ads

71 The Guardian, 2015, Google's ban on payday loan ads recasts debate of morality in media for digital age

72 Facebook.com, 2025, About audiences for housing, employment or financial products and services campaigns

73 Sara Kingsley, Clara Wang, Alexandra Mikhalenko, Proteeti Sinha, and Chinmay Kulkarni, 2020, Auditing Digital Platforms for Discrimination in Economic Opportunity Advertising

Facebook and deliberately excluding younger users: a credit card ad was not shown to under 25s, and a property related financial product was only shown to over 35s.[74]

Although any direct targeting of credit products should now be impossible due to Facebook's new restrictions, it is unlikely that the restrictions will stop targeting by proxy. A number of different proxy indicators could contribute to the targeting of financial products to consumers with lower credit rating, or an urgent need for credit for example:

Poverty is much more prevalent in households renting their home than than those owning their own home[75]

While jobseeking does not automatically mean someone is unemployed it could be an indicator, and unemployment particularly in the long term is associated strongly with poverty.[76]

It is possible to select an audience for Facebook ads for financial products and services using these characteristics , either intentionally through the choices an advertiser makes themselves, or unintentionally as a result of Facebook's own automated ad targeting. In combination, it is likely that an audience with these characteristics will include a high percentage of people on a low income who may be more susceptible to taking out high interest, unsustainable loans or credit.

Evidence from an investigation into Google ads shows that this kind of potentially discriminatory targeting can happen automatically. Google was serving adverts for poor quality credit products to users who searched for terms such as 'quick money now' and 'need money help', strongly suggesting these were people in financial distress.

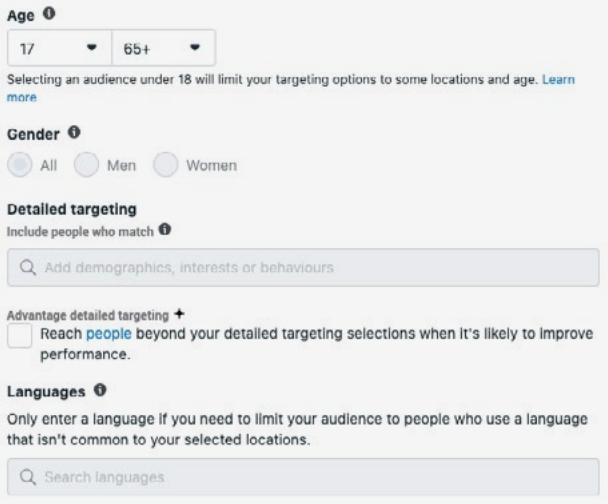

Figure 9: Screengrab from Meta's ad portal showing the targeting options available for an ad in the Financial Products and Services category

The ads promised things like 'no credit check' and money delivered quicker than a takeaway pizza; potentially in breach of the Advertising Standard Authority's rules. The content of some adverts also implied they were being targeted at people with poor credit histories or who were otherwise in financial difficulty. People in financial distress have few options and can find themselves trapped in unaffordable debt by high interest rates.[77]

Even if Meta itself now restricts the direct targeting of financial products, there is evidence that Meta users' data is used to make financial judgements about them. User interests still contain finance-related categories such as 're-financing' for example. A fraud prevention firm explicitly advertises 'social media credit scoring' as a service, promising financial companies "the tools you need to get a complete view of your customers". Credit providers actually boast of their ability to target 'sub-prime' users on Facebook, who have less access to affordable credit and so are more at risk of taking out poor quality finance products which may trap them in extremely high interest rates and debt.

74 Corin Faife, Alfred Ng for The Markup, 2021, Credit card ads were targeted by age, violating Facebooks' anti-discrimination policy

75 Institute for Fiscal Studies, 2023, Housing quality and affordability for lower-income households

76 Joseph Rowntree Foundation, 2024, UK Poverty 2024

77 The Guardian, 2022, Google profiting from 'predatory' loan adverts promising instant cash

HEALTH AND MENTAL HEALTH

The Observer discovered that British charities providing people with advice and support for their mental health, including depression, eating disorders and self-harm, were sharing user data with Meta. Via the Meta Pixel, Meta received details of sites that users visited and which buttons they clicked on. Some of the sites were specifically aimed at under-18s. In many cases this data was linked to the individual's Facebook account, as well as to their IP address, which includes location information. All of the charities that used the Meta Pixel had it on their homepages, so the data was shared before a user even had the opportunity to opt in or out of data sharing. Facebook's access to this data would have enhanced its ability to profile users and therefore target adverts towards them. For people dealing with mental health challenges, this could result in being shown upsetting, intrusive and misleading ads or content.[78]

The Observer also found 20 NHS trusts with the Pixel on their websites, potentially sharing sensitive health information with Facebook without user permission. People looked at information on conditions including HIV, eating disorders, cancer and mental health, and their browsing was then shared with Meta. This opened them up to the potential of being targeted with adverts for products or services related to these conditions. This is intrusive, potentially unwanted, and could have caused harm, for example weight loss adverts, scam treatments or inferences about sexuality or sexual behaviours. Facebook says it has filters to exclude this sort of data from its profiling and ad targeting, but they have also admitted that their systems don't catch everything.

In late 2024, Meta announced plans to limit advertisers' access to data about users' health, as well as about financial services and politics. Meta says this is part of an effort to reduce the amount of sensitive personal data that they collect and hold. It comes on the back of regulatory intervention in the US against medical-related companies that were sharing user data with Meta via the Meta Pixel.[79] Meta faced legal action for receiving sensitive medical information and not taking proper action to stop it from occurring.[80] This change is welcome in light of the limited ability Meta users have to filter out content they do not wish to see in their feeds.

HOUSING

Another lawsuit in the US led to significant changes in Facebook's ad targeting. The Lookalike Audiences feature, which enables advertisers to find new audiences based on the characteristics of their existing customers, was enabling discrimination despite the 2019 ruling. It was therefore possible to exclude certain audiences from seeing adverts for housing, based on protected characteristics including race. Facebook was required to change how they show housing adverts to ensure they are shown to a representative mix of people, not unfairly excluding any age, gender, race or ethnicity.[81]

Facebook also announced it would extend these changes to employment and credit adverts, stopping offering the Lookalike Audiences tool for these adverts, and also introducing a new tool to make the audiences that adverts are served to less biased (the Variance Reduction System), although this tool has only been introduced in the US.[82]

78 The Observer, 2023, UK mental health charities handed sensitive data to Facebook for targeted ads

79 Marketing Brew, 2024, Meta plans crackdown on health related user data

80 The Observer, 2023, NHS data breach: trusts shared patient details with Facebook without consent

81 Martech, 2022, Facebook agrees to revamp adtech over discrimination charges

82 Rachel Griffin, 2022, Tackling Discrimination in Targeted Advertising

EDUCATION

A lawsuit was filed in early 2025 against Meta in the US alleging that Meta's ad targeting disproportionately shows adverts for more expensive, for-profit colleges rather than public not-for-profit universities to Black users. This would be limiting the information that Black users get about educational opportunities, and perpetuating existing inequalities, as there is evidence that for-profit colleges provide a lower quality education, resulting in lower earning potential and higher levels of student debt. The lawsuit suggests that Facebook's ad targeting is over- or under-predicting the likelihood of certain groups of users in seeing the different adverts from the different types of college, and therefore serving up the adverts in such a way as to result in discrimination.[83]

This has been called 'predatory inclusion', meaning the inclusion of already marginalised groups or individuals in a scheme or offer which superficially provides them with benefits, but in reality is extractive. Rather than 'expanding access' as this advertising may appear to be doing, it is actually ending up with Black students in worse debt.[84]

SCAMS AND PREDATORY ADVERTISING

Targeting allows those looking to scam people, sell fraudulent products or services or otherwise defraud consumers to do so more effectively. Because Facebook's ad algorithms are specifically designed to find the users most likely to engage with particular topics, messages and ads, an unscrupulous advertiser has the perfect tools to target vulnerable users who are more likely to fall victim to their scams.

Older Facebook users in the US were tricked out of savings in a scam investment which was specifically advertised to people aged 60 and over who were politically or economically conservative. The ads should have been removed on the basis of their content and the scam that they were attempting to lead people into, but some versions of the ads remained live for nearly two years. Theoretically, Facebook has automated systems to detect scam adverts, but they clearly did not work in this case.[85]

The Tech Transparency Project found thousands of scam adverts on Facebook offering 'free' Donald Trump merchandise. One advert took users through a number of websites to finally end up on a page which seemingly only required payment for postage of the 'free' merchandise. It was in fact signing people up for an $80 / month subscription, in breach of Meta's policies against 'deceptive and misleading practices'. It is easy to see how scams involving political affiliation could be targeted towards the Facebook users who are more likely to engage with them. The investigation also found fake health benefits being promoted, promising users over $1,000 in credit. Again, it is easy to imagine what kinds of user interests these scam adverts could be targeted at.[86]

---

83 Lawyers' Committee for Civil Rights under Law, 2025, New Lawsuit Challenges Big Tech Firm Meta for Discrimination in Advertising Higher Education Opportunities

84 Tressie McMillan Cottom, 2020, Where platform capitalism and racial capitalism meet: The sociology of race and racism in the digital society, Sociology of Race and Ethnicity, Vol 6, Issue 4

85 Jeremy B. Merrill, Hanna Kozlowska, Quartz, 2019, How Facebook fueled a precious-metal scheme targeting older conservatives

86 Tech Transparency Project, 2024, Meta hosts sprawling pro-Trump scam ad network as election nears

WHAT WILL THE IMPACT OF GENERATIVE AI BE?

Meta has been using AI in its advertising offer for many years, with algorithms deciding which ads to show to which users. In the last couple of years it has also introduced Generative AI (GenAI) into its suite of tools for advertisers, and the 'behind the scenes' workings of ad targeting. Meta's GenAI is in part trained using public posts and comments from Meta users, unless they have opted out.[87]

Advertisers can use GenAI tools to create their own adverts, including creating different versions, changing images and text, and fine-tuning content, language, design and the people that appear in adverts to appeal to different audiences.[88] Meta reports that over a million advertisers used GenAI tools in one month in 2024, and that adverts created with their GenAI tools result in 11% higher click-through rates and 7.6% higher conversion rates compared to campaigns that did not use them.[89]

Meta's AI-powered advertising tools are called Advantage+, and are trained on results from ads across Meta platforms, including Facebook and Instagram. An Advantage+ Shopping Campaign (ASC) uses machine learning to 'dynamically serve your ads to the audiences most likely to convert'. It removes most of the manual targeting options, instead using AI to automatically target adverts to the 'best' audiences. i.e. those most likely to click and make a purchase. It promises to 'automatically show ads to people most likely to respond', using machine learning to learn constantly from user interactions with ads and from their other activity, on- and off-site.[90]

Advantage+ also promises to expand ad audiences beyond its established ad targeting methods. It will constantly assess ad performance and if the system finds better 'performance opportunities' outside of the audiences defined by the advertiser it will show the ads to these 'new and unexpected' audiences.[91] This feature is not allowed to be used for characteristics or products/ services which have restrictions on them, such as age or gender.

This next iteration of Meta's ad targeting underlines the centrality of advertising to Meta's business model. Embedding GenAI in both customer-facing and behind the scenes aspects of the ad system promises to improve performance, but also adds another layer of opacity to the whole process. By further automating ad audience selection, both advertisers and consumers will have less chance of truly understanding why certain people are seeing the ads they are shown. It may also make the detection of biased or discriminatory proxies even harder, meaning more of them slip through the net and more users experience discrimination or harm as a result.

This makes it harder for conscientious advertisers to ensure their ads are shown to users fairly and with due care, and for users to understand what data has been used to include or exclude them in an ad audience. It takes human decision-making and control even further out of the equation, reducing opportunities for accountability, explainability and making meaningful change when things go wrong. GenAI tends to amplify bias more than other types of AI, partly because it uses larger datasets than other AI models, making it harder to audit any biases or rebalance the data, so the more it is used in ad targeting the greater the risk of further embedding bias and creating more discriminatory outcomes.

As users we may be subject to ever-more sophisticated manipulation through ad content and creative adjusted by GenAI, as well as more fine-tuned targeting, trained on data we provide ourselves. As things stand, we have little scope to opt out of this if we want to keep using these platforms.

87 Facebook.com, How Meta uses information for generative AI models

88 Forbes, 2023, 5 Amazing ways Meta is using generative AI

89 Meta, 2024, Meta's AI products just got smarter and more useful

90 Meta, no date, About Advantage+ Audience

91 Meta, no date, About Advantage lookalike; Meta, no date, Maximise performance with Meta's suite of AI-powered ad tools

RECOMMENDATIONS

There is a significant gap between our rights on paper to opt out of data profiling and targeted ads and our actual ability to realise these outcomes. Too many consumers are putting up with the uneasy compromise between giving up our privacy, opening ourselves up to the risk of harm and discrimination and using ubiquitous websites and platforms. Decisions about our data are heavily weighted in favour of the platforms, tech companies and advertisers. However, there are practical ways to improve the balance and ensure consumers can fully exercise their data rights.

RESPECT PEOPLE'S RIGHT TO CONSENT TO TARGETED ADVERTISING

Every user of a site or platform which uses profiling and ad targeting should only see targeted ads if they have consented for their data to be processed for this purpose. People should also be able to simply and effectively use their right to opt out at any time. It should not be a paid-for privilege, but a universally available right. Opting users out of data profiling and targeting should be the default for sites like Facebook, with users who prefer targeted ads able to opt in if they wish.

The problem with wanting to opt out of profiling and targeting is not that we have no right to do so, but that our existing rights are not respected by the likes of Meta. Most data rights exist for individual consumers, and these are hard to enforce, as the O'Carroll vs Meta case exemplifies. This means we as users of platforms like Facebook are limited in our ability to make meaningful change. As discussed in section 6.1 we can opt in and out of certain features but if we want to keep using their platforms we have to accept the fact we are being profiled and served targeted ads.

Profiling puts us into categories with other people on the basis of online behaviours and characteristics that we may never know about, and these categories may hold no clear meaning for us. We may be placed in a category which means we are excluded from seeing ads and content that we would actually like to see, or included in an ad audience which shows us something we'd rather not, but this category might be entirely opaque or impossible to understand. We may never know if we're included in a category that shows us in a negative light, considers us as 'risky' or an undesirable customer.[92] This limits our ability to meaningfully opt out, make informed choices or hold the platforms to account.

If UK users were able to fully enforce their rights under the GDPR in theory their data would not be processed to create a profile that could be used in ad targeting. This would mean that users would see a random selection of adverts when using a platform like Facebook, or ads chosen as a result of the content someone is looking at rather than their individual characteristics. It might also mean that non-advertising content such as brand-sponsored posts by influencers would also be shown randomly rather than through targeting. It may not stop the targeting of non-marketing content, but as the rights are currently not enforceable and no-one entirely understands how Meta actually targets users, we cannot be sure. It could also mean that the data belonging to the individual user who enforces their rights would not be processed for ad targeting to other users, for example as part of a lookalike audience.

Beyond being able to exercise our rights, all users should be opted out of profiling and targeting by default, with only those who actively choose it having their data being processed for profiling: Meta argues that many of its users prefer targeted ads, in which case this could be an opt-in option.

92 Mann, M., & Matzner, T., 2019, Challenging algorithmic profiling: The limits of data protection and anti-discrimination in responding to emergent discrimination. Big Data & Society, 6(2)

In Europe, Meta users have the option to 'pay or consent'. There is a paid subscription available to users who wish to restrict the amount of data profiling carried out on them, but this is not foolproof or universally liked. If users continue on the free version of Facebook or Instagram they essentially have no choice but to continue to allow their data to be processed for profiling and ad targeting.

European consumer organisations complained that the first version of pay or consent, introduced in 2023, breached EU law, and Meta rolled out a second version of it in late 2024. Many believe it is still not lawful, based on unfair and unclear practices, and should not be in place.[93]

Pay or consent creates a two tier privacy landscape, where those who can afford to pay the subscription are granted greater digital privacy than those who cannot.[94] In order for consent and pay to be implemented in a lawful manner, anyone choosing not to pay to remove adverts should still be presented with a clear choice between receiving targeted adverts determined by their personal data, or contextual or untargeted adverts. By making opt-out the default, these rights would not be restricted to those who know how and can afford to exercise them: they should be freely available and easily enforceable to anyone who wishes. We explore this in greater detail in our report on advertising business models.[95]

IMPROVE AD TRANSPARENCY

The transparency introduced by the Meta Ad Library should be built on and strengthened, with all ads subject to a stronger minimum level of transparency, and access to the Library should be freely available without logging into a Meta account.

There has been progress in achieving more transparency around Meta advertising: there is now a Meta Ad Library, a searchable directory of ads placed on Meta platforms. It shows who placed an advert and how much they have spent advertising on Meta, as well as information about paid partnerships between content creators and brands. The library can be used to monitor ads and advertisers, which is a good first step to holding them, and Meta, more accountable. Some believe that transparency is a more achievable and effective goal than a complete ban on targeted ads, as this can have unintended consequences like banning charities or community groups from reaching people who might benefit from their support.

There is scope to improve the ad library:

- It should not require a user to be logged in with a Meta account to use it. In order to look up ads and other information on the ad library, individuals and organisations have to be logged in to a Meta account. This puts up barriers to transparency and accountability, and ironically forces people to hand over their personal data to Meta
- All ads should be subject to the greater transparency applied to political ads. The ad library provides much more detail about who is shown political ads than mainstream ads, including a breakdown by gender and age. This information should be available as a minimum requirement for all ads on Meta platforms.

93 BEUC, 2025, Consumer groups red card Meta's latest pay-or-consent policy

94 Derek E Bambauer, 2025, Target(ed) Advertising, UC Davis Law Review Vol. 58

95 Privacy without Paying: Alternatives to Meta's Surveillance Advertising Models, Open Rights Group, May 2025

USER SWITCHING AND INTEROPERABILITY

Interoperability and user switching can reduce advertising harms by incentivising competition to improve the quality of the advertising experience.

Another important way of mitigating the discriminatory and harmful impacts of targeted advertising would be to promote interoperability between social media platforms, advertising systems, and third-party tools. Interoperability would enable users to access social media platforms via alternative services or applications. These could have enhanced privacy, transparency or safety features. We will make the case for social media interoperability in a separate report in mid 2025.

Interoperability already exists as the result of government regulation in Open Banking,[96] which allows interoperable bank information and user decisions including provider switching. Mobile phone switching is a regulatory requirement for mobile phone numbers to be moved to a new operator, so users can change provider. Social media interoperability could similarly allow users to switch provider, from for example Threads.net to BlueSky or Mastodon without losing their contacts. Threads.net, owned by Meta, has in fact promised this kind of platform switching, so that Threads users could move to a Mastodon service, or from Mastodon to Threads, while retaining their network connections. BlueSky also promises similar kinds of switching. This kind of horizontal (external) interoperability would allow users who are not happy with the kind of advertising they receive on a platform to move to a new one. Even a small number of users leaving would create an incentive for better delivery of more trustworthy ads, so has the potential to allow the market to drive up standards.

Governments and regulators already have a number of policy tools, ranging from penalties and fines, tax relief for open systems, procurement rules that can be used to favour interoperability systems, trust-mark systems, funding the development of interoperable protocols such as ActivityPub, easier compliance with online safety regulations, and market merges for interoperable systems. Powers exist in competition law to require interoperability where there is market harm.[97]

Platforms can also offer vertical (ie, internal) interoperability. By encouraging or requiring them to open up key elements of their advertising infrastructure, such as APIs for ad delivery, targeting transparency, and user data portability,regulators could enable users to deploy their own ad filtering tools or to access platforms via intermediaries that offer protection from harmful or discriminatory ads.

Interoperability does not destroy advertising-based business models. Platforms can still deliver ads and earn revenue while allowing greater user control over how ads appear and what kinds of targeting are permissible.

For example, interoperable systems could allow users to opt into ad categories, block sensitive topics, or use trusted third-party filters without entirely removing advertising from the ecosystem. This kind of internal interoperability is already planned at BlueSky for content recommendation engines. The same thinking could be applied to develop capabilities for users to choose or restrict ad delivery.

This approach would create a healthier balance between commercial sustainability and user safety. It would also drive competition based not only on audience size but also on the quality and safety of the advertising experience, rewarding platforms that offer greater transparency and accountability.

96 Retail Banking Market Order 2017 (CMA Order) and The Payment Services Regulations 2017

97 Section 20(3)(e) Digital Markets, Competition and Consumers Act 2024

DEVELOP AND SUPPORT NEW MODELS OF ADTECH

Ad targeting doesn't have to be done through data profiling: contextual advertising can achieve similar results without collecting personal data and violating user privacy. This and other models of privacy-preserving online advertising should be developed and supported by advertisers and platforms.

Some advertisers are looking at different models of ad targeting, in particular contextual advertising. Rather than build a user profile which follows the user around the Internet, contextual advertising places adverts based on the sites and contents they are looking at. If someone is on a website about walking holidays, they will see ads related to that; if they are browsing beauty blogs, the ads they see will be relevant to the content of the blogs, not chosen because of the user's data profile. This means ad targeting is still possible, but does not rely on harvesting and exploiting our personal data.[98]

Contextual advertising has a range of potential benefits for advertisers and consumers:

- It has been shown to be as effective as ad targeting based on data profiling
- It doesn't rely on extensive data collection and analysis, so can be more accessible and cost effective for advertisers
- This also means it doesn't risk violating users' privacy and data rights
- It provides a user experience that many people value and feel more comfortable with, avoiding potential damage to an advertiser's reputation and brand.

Platforms like Facebook would still be able to generate ad revenue from contextual advertising, without breaching our data rights.

The O'Carroll vs Meta case exemplifies the need to significantly shift the balance in power between consumers and platforms, and emphasises the need to allow consumers to continue to use platforms like Facebook without having our personal data processed against our wishes for advertising that causes unlawful discrimination and significant harm.

98 Privacy without Paying: Alternatives to Meta's Surveillance Advertising Models, Open Rights Group, May 2025

openrightsgroup.org

Published and promoted by Open Rights Group, a non-profit company limited by Guarantee, registered in England and Wales no. 05581537. Registered address Space4 113-115 Fonthill Road, London, England, N4 3HH