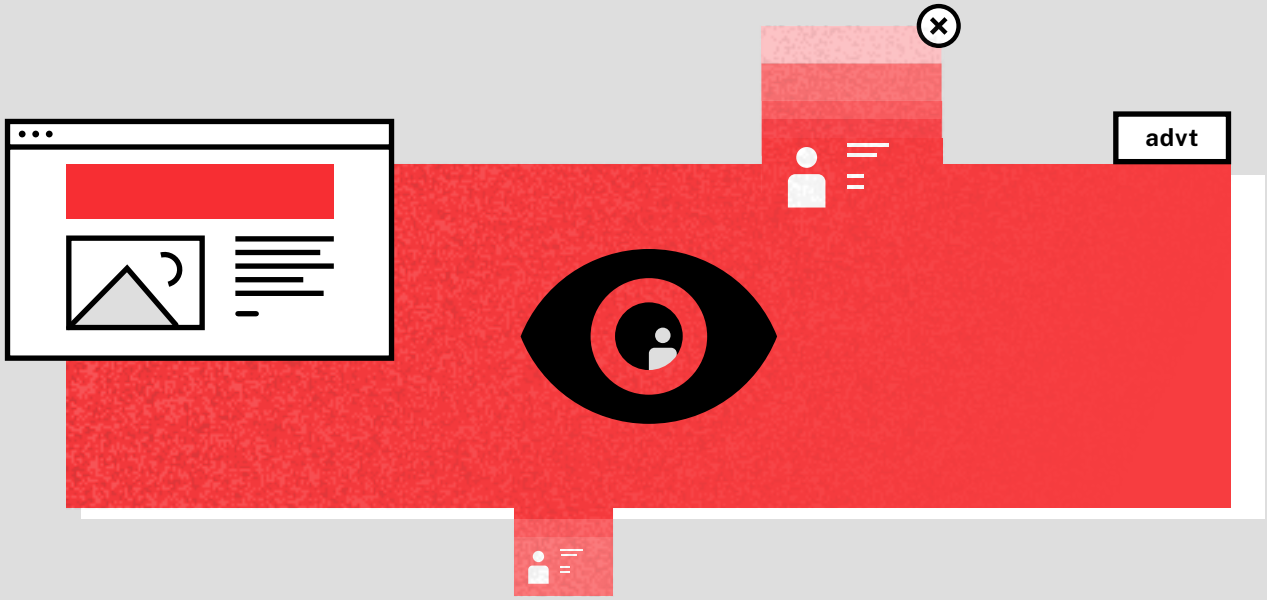


# TO TRACK OR NOT TO TRACK?

## Towards privacy-friendly and sustainable online advertising

**Karolina Iwańska**
Panoptykon Foundation

PANOPTYKON FOUNDATION

Supported by
**Mozilla Foundation**



November 2020

# INTRODUCTION

The dispute over privacy online boils down to which model of monetisation of online content we are ready to support and whether we can imagine it without behavioural advertising. Most online publishers and advertisers can't (for reasons that are not as simple as their supposed malevolence), so they often submit users to an unfair choice: if you want to access our content you have to accept being tracked and profiled.

There is no doubt that online publishers have the right and the need to earn money from the content they provide. It's important that good quality journalism has the appropriate resources to thrive. Rather, the question is: should online content be monetised on the basis of intransparent and unethical practices, such as tracking people without their knowledge or meaningful consent? Even if the answer is "no", actually making it a reality proves to be difficult.

## (Rotten) business as usual

Over the last 25 years, since the first online ad banner appeared in 1994, advertising based on surveillance has become ubiquitous. It's been fast-tracked by the development of the automated way of selling ads, in particular "real-time bidding" first introduced in the United States in 2009. Bob Hoffman, an advertising veteran and a fervent critic of the current state of the ad industry, calls the decade of real-time bidding "the advertising's decade of delusion".[1]

It goes without saying that this way of selling ads has not been designed to put privacy first. 2009 was nine years before data protection made headlines (also in the US) with the introduction of the GDPR. While this law did not significantly modify basic data protection principles that have been in place in Europe for the last 20 years, it carried the promise (or the threat) of stronger enforcement and – importantly – applied not only to companies registered in the EU, but to all those that process personal data of EU citizens.

In theory, on 25 May 2018 when the GDPR went into effect, we all should have seen a tremendous change in how online advertising works. But it didn't happen. Some cookie pop-ups have been adjusted, some internal policies edited, "we care about your privacy" emails sent, but overall – business carried on as usual.

## The ad tech industry's game of influence

The reason for this is that the ad tech industry – companies which broker user profiles and attention, fill empty ad slots on publishers' websites and provide the audience for campaigns run by advertisers – have over the years promoted a win-win narrative around behavioural advertising and made sure that no alternatives can flourish. In effect, policymakers and data protection authorities are facing the challenge of reconciling the protection of the fundamental right to privacy with the

[1] http://adcontrarian.blogspot.com/2019/10/advertisings-decade-of-delusion.html

concern that enforcement of the law would lead to depriving publishers of a key source of revenue. Publishers themselves are often siding with the advertising industry and lobbying against privacy laws, for example the upcoming ePrivacy Regulation which – according to the industry – would be "like a bad movie"[2] for the Internet.

And then we have Google and Facebook – advertising hegemons who use their dominant positions to sweep advertisers' money by offering them the detailed profiles and the eyeballs of over 2 billion users globally. Publishers are forced to compete in this uneven race for clicks, sacrificing their readers' or viewers' privacy and – increasingly – the quality of the content they produce.

## What's the alternative?

At the end of 2020, we seem to be closer to the tipping point. The advertising industry is under investigation by a number of data protection authorities and key decisions are likely to be made as soon as in early 2021. Concerns are voiced more and more loudly even within the industry. But half-measures or cosmetic fixes that they propose will address the symptoms but will not suffice to deal with the negative consequences of the current ad model that go beyond data protection: increasingly stronger financial dependence of online publishers on ad tech middlemen, the rise of clickbait and sensationalist content, the deterioration of public discourse, and an influx of bots that launder advertisers' money for organised crime.

Therefore, the urgent question is: what is the alternative? How can the European Union pave the way to an online future which is a real win-win for publishers and citizens, and not solely for advertising intermediaries?

This brief aims to contribute to this debate by answering the following questions:

- How does open web advertising in its current form violate the principles enshrined in the GDPR?
- What are the economic and societal consequences of the current advertising model and how do they undermine the win-win narrative promoted by the advertising industry?
- What are the existing alternatives?
- Why does surveillance-driven advertising persist?
- Finally, what regulatory interventions are needed to promote the uptake of privacy-friendly and sustainable alternatives?

Answering these questions has been made possible by over two years of investigation, discussions with tech experts, publishers (big and small), advertising industry insiders, and regulators. The last year of this work, including the preparation of this brief, has been supported by the Mozilla Foundation's fellowship programme[3].

[2] https://www.likeabadmovie.eu/

[3] https://foundation.mozilla.org/en/fellowships/

# PART 1

# ONLINE ADVERTISING AND ITS DISCONTENTS

# 1. ONLINE ADVERTISING AND ITS DISCONTENTS

## 1.1. Spotlight on real-time bidding

Online advertising is a catch-all term that includes many forms of ads, different technical methods for delivering them, and different actors who take part in this process. Unpacking this general term is essential for accurately identifying problems and their sources, as well as formulating effective responses.

This brief does not cover all forms of online advertising. As its purpose is to examine negative consequences of surveillance-driven funding models for publishers and to propose a way forward, it zooms in on the so-called **open display advertising –** advertising served on publishers' websites or in their mobile apps. On this market, ads can essentially be sold in two ways: through direct sales or on ad auctions.

**Direct sales,** as the name suggests, involve direct, not intermediated contact between an advertiser (or an agency) and a publisher. Ads sold in this way can either be sold by humans in traditionally negotiated deals or by computers using programmatic technology which automates the process.

**Online ad auctions** on the other hand enable multiple advertisers to simultaneously compete for the possibility of showing an ad to the person currently visiting the website. Ad auctions are automated and normally involve technical intermediaries ("ad tech"). Private auctions, known as **private marketplaces (PMPs)**, are usually operated by the publisher, a group of publishers or an external gatekeeper who selects, verifies and approves advertisers taking part in the auction. As a result the number of advertisers competing for the possibility to show an ad is limited. **Open ad auctions** on the other hand allow any company to participate, provided that it meets technical requirements and promises to adhere to relevant terms and conditions.

The winner of online auctions is selected within milliseconds which is where the term of "real-time bidding" comes from. Although there are thousands of companies which take part in auctions, in practice two organisations define technical and organisational standards for **open real-time bidding systems:** Google and the Interactive Advertising Bureau (IAB), responsible for systems respectively known as Authorized Buyers and OpenRTB.

Both systems have recently come under increased scrutiny of privacy advocates and data protection authorities. In 2018 and 2019 GDPR complaints have been filed against Google and the European branch of the IAB in 17 European jurisdictions. Lead supervisory

authorities have initiated investigations: the Irish Data Protection Commission into Google and the Belgian Data Protection Authority into IAB Europe. As of November 2020, the investigations have not yet been concluded, but the Belgian DPA has issued a preliminary report that is damning for the IAB. Please refer to Part 3 for a more detailed analysis of enforcement issues.

This brief **focuses on real-time bidding** as a method of selling ads that is based on large-scale data sharing between a large number of intermediaries. As such, it poses a massive threat to privacy and makes it impossible for users to effectively control their own data. This brief argues that **the very way in which this system is designed brings about acute negative consequences not only for individuals, but also for publishers, advertisers, and the society at large.**

## 1.2. How real-time bidding works

### STEP 1: Tracking and profiling

As a person navigates the web, small text files (cookies) are installed in their browsers. These cookies collect and save information about the user's device, browser settings, IP address, location, or what they've read or watched. Some cookies are **necessary for technical purposes:** information about the device might enable site owners to identify what adaptations it should make to fit a mobile screen, or suggest the right language version. Other cookies are **needed for user authentication,** remembering items added to an online basket, or to help autofill forms. Finally, **some cookies are used to track and analyse users' behaviour for statistical or advertising purposes:** what articles they read and when, how much time they spend on the website, or which ads they click on.

Cookies can be read only by the website that installed them. Google.com cannot read cookies installed by bbc.com. It either has to ask the BBC to share information related to that cookie (for which the BBC would have to obtain users' consent), or place its own cookie which directly monitors users when they visit the BBC's website. The second scenario is what usually happens: third-party cookies are set by advertising companies along first-party cookies installed by the website itself. **Some websites allow as many as 400 third-party advertising partners to set cookies on their "premises".**

Data that would directly identify the user, such as their name or address, is not very relevant in the context of online advertising. **What matters are characteristics which make the user less or more likely to react to a particular ad.** This is why advertising companies use unique IDs linked to a dynamic marketing profile, which includes interests, purchase power, state of health, information on an important moment in life etc. The assumptions that form a profile can grow more and more sophisticated and accurate, as data linked to cookies accumulates over time and companies record users' activities across multiple websites and devices.

**Three layers of an online profile**

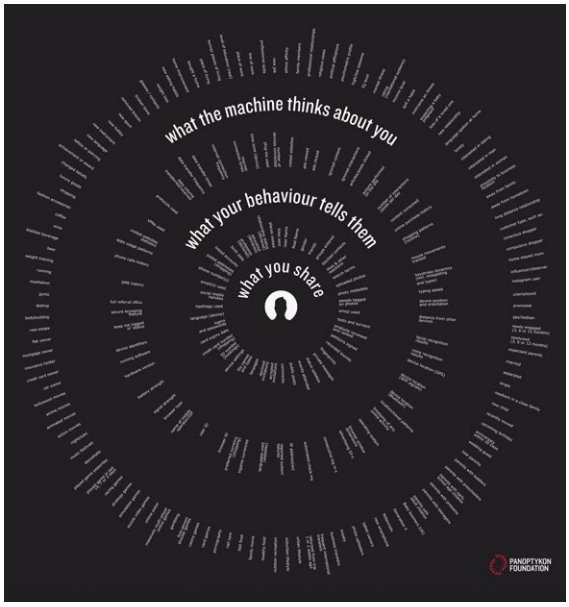


Users of smartphones and other mobile devices are known to various intermediaries under standardised **advertising IDs**, e.g. Google Advertising ID for Android devices (AAID) and Identifier for Advertising for iOS devices (IDFA). Although these IDs can be reset by the user researchers established[4] that thousands of apps available for Android devices violate this policy by using tracking IDs other than the AAID which are non-resettable. Apple, since the introduction of iOS 14 in 2020, requires mobile apps to ask for users' explicit consent to use IDFA for tracking purposes. Previously, limiting tracking was available based on the user's opt out, rather than consent.

In the world of laptops and PCs the situation is different – there are as many different IDs as there are different intermediaries. Making sense of who the user is is possible via **cookie syncing** – a process which enables various ad tech companies to exchange information on IDs they have assigned to users and profiles linked to them. Identifiers and profiles are stored by specialised **data management platforms** (DMPs).

## STEP 2: The auction

A standard real-time bidding transaction begins when the browser makes a connection to the website that the user is trying to load. The **website (publisher)** verifies whether there is a cookie already available for this user (this is normally the case when the user visited the website in the past). If the user is not known to the publisher, a new cookie is installed and a new ID is assigned. It is then sent to the ad server which connects to the so-called **supply-side platform (SSP) –** a specialised piece of software which enables real-time communication with other players in the auction. The SSP then sends a **bid request** to the **ad exchange** where a virtual auction takes place, and awaits offers. The bid request contains information on the available ad space and pricing as well as information about the user (see more in the box below). Bid requests are often sent to several ad exchanges at the same time.

### WHAT GOES INTO THE BID REQUEST?

The real-time bidding process is highly standardised - all companies that would like to take part in it must conform to detailed technical guidelines developed by Google

[4] https://blog.appcensus.io/2019/02/14/ad-ids-behaving-badly/

and the IAB. These standards dictate what types of information about users are sent by the SSP to the ad exchange and - as a consequence - which data is broadcast to hundreds of bidders.

**A standard bid request contains:**

- a user ID set by the SSP,
- so-called full referral URL, meaning the link to the website where the ad is supposed to appear, a phrase or the link from and a category assigned to the website which - although relating to the content of the website - can reveal features of people visiting it and be highly sensitive (e.g. support for victims of abuse),
- year of birth,
- gender,
- location,
- IP address (some systems truncate it),
- interests or segments previously assigned to the user,
- other information the SSP might hold.

The ad exchange then broadcasts this information to tens if not hundreds of companies working for advertisers, known as **demand-side platforms (DSP)**, programmed to find users matching a predefined profile relevant to one of the campaigns they are currently running on behalf of advertisers. These profiles are usually established by advertising agencies. Their clients – brands – normally issue a standard instruction: "reach people that will be interested in what I have to offer". The role of the agency is to determine which features this person might have and where they can be found online. The basic knowledge about potential clients comes from data about past purchases provided directly by advertisers. In order to generate it, most companies maintain Customer Relationship Management systems, loyalty programmes, or their own online shops. This is the starting point for agencies which in the next step try to determine which features and behaviours are characteristic for people who have not yet bought the product but might do so in the future. In marketing jargon the profile of a hypothetical customer is called a **lookalike**.

**How data is broadcast** →

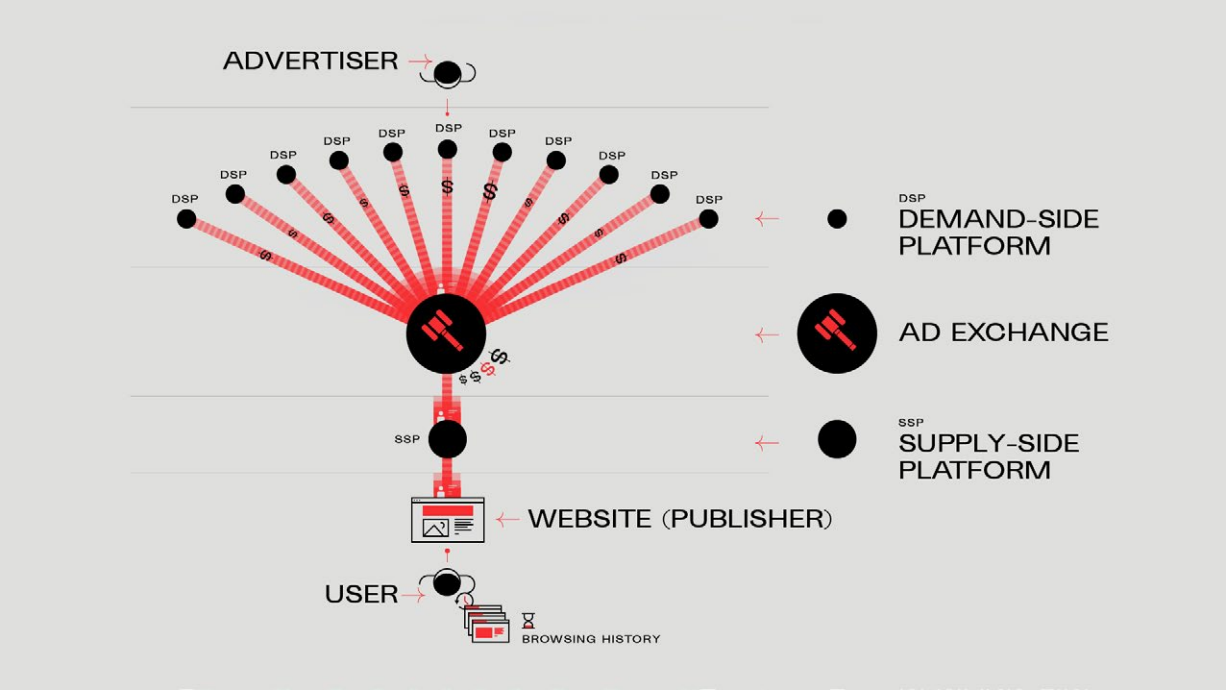


In milliseconds DSPs match the user ID they received from the ad exchange against their database, looking for any additional information about the user that might make this assessment more precise. What is known about the user's previous transactions? Which websites have they just visited? Were they comparing prices? What are they looking for and how much are they ready to pay? On the basis of data they have from observing the user across the web, they decide on whether to place a bid (just like in an old-fashioned auction) and at what price, **hoping to win the possibility to show the right ad to the right person at the right time.**

The ad exchange selects the winning bid (usually the second-highest, to support competition and avoid outrageous bids) and transmits the good news, together with the content of the ad and the tracking code of the DSP, back to the SSP. **Only then does the user see the ad on the website.** According to terms and conditions of major ad exchanges, only the winning DSP gets to keep the ID and the data sent by the SSP in the bid request to enrich the profile of the user for the purposes of future auctions.

**All of this is fully automated and run by algorithms** – from creating a user's profile to determining the bidding price and the winner of the auction. From the perspective of the user it remains completely unnoticeable – **the entire transaction lasts 200 ms** ($^1/_5$ of a second). It's less than a blink of an eye for which we need on average as much as 300 ms. This obscurity is why very few people are interested in what's happening behind the scenes of the commercial web. As a result we end up with a dangerous information asymmetry exploited by professional players to influence the behaviours of unaware users.

→

# 1.3. A system broken by design and by default

**The way online advertising works creates problems that are not just side effects that can be fixed with a few cosmetic tweaks. Rather, they stem from the very design of real-time bidding systems.**

The ad tech industry claims that behavioural ads make it possible for users to browse the internet for free while at the same time rewarding publishers for creating content, and enabling advertisers to promote their products or services. Which sounds like a win-win situation. In reality, however, **there is but one winner: online advertising intermediaries.** Users, publishers, and even advertisers are all, to a smaller or larger extent, losing in this game.

For users, online advertising in its current form creates huge risks to privacy and individual autonomy. Publishers are kept in a prisoner's dilemma where they are facing the choice between having to hand over up to 70% of their profits to advertising middlemen or not being able to fund themselves at all. Advertisers have to rely on unreliable (if not completely useless) metrics, manage an increasing irritation with ads best illustrated by the growing number of adblocks, and deal with ad fraud (bots fabricating clicks on ads) which funds organised crime. Finally, societies have to bear the consequences of attention-driven behavioural advertising by observing the deterioration of media quality and an increase of sensationalist content, which is not without effect to the quality of public debate and democracy.

### PRIVACY: Real-time bidding is by design incompatible with the GDPR

Open web advertising in its current form violates the principles enshrined in the GDPR in the following way:

- **Uncontrolled sharing of data is the heart of the problem: violation of the principle of integrity and confidentiality**

Once the publisher ("represented" by the SSP) sends data about a user to a couple of ad exchanges, users lose control over how their data is then used by a potentially unlimited number of other actors. Intimate details about people are broadcast to hundreds, if not thousands of companies taking part in online ad auctions, billions of times per day[5]. These companies use the data they receive for their own purposes, e.g. to build user profiles.

Publishers, despite being  the only company with whom users have a direct relationship, do not have the means to make sure that it doesn't happen. There are no real safeguards that would keep data secure, as companies tend to over-rely on contracts as "guarantees" of security. For example, Google's advertising guidelines[6] say that only companies that win a given auction may keep data to enrich user profiles, but these are just contractual, not

[5] https://www.iccl.ie/human-rights/info-privacy/rtb-data-breach-2-years-on/
[6] https://www.google.com/intl/en/doubleclick/adxbuyer/guidelines.html

technical, measures. In fact, sources within the industry claim that some companies take part in ad auctions only to get access to people's data, and without the intention to win. As the CEO of Tapad, a company that focuses on cross-device tracking, puts it: "In the ecosystem there is a general understanding that it's in everyone's advantage to share data"[7].

Google and IAB's RTB systems have in fact been **constructed in a way that makes it impossible to control data once it is shared on the ad exchange –** neither by the publisher who shared it, much less by the user. This is a data breach that affects basically everyone that has ever used the Internet. The violations of other GDPR principles stem from this feature inherent to real-time bidding systems.

- **Consent is illusory: violation of the principle of lawfulness and fairness**

The ad tech industry claims that users freely give their consent to share data for advertising purposes because they want to have "a personalised experience" online. But what these companies call "consent" cannot be qualified as consent under the GDPR.

First, **consent is not freely given.** The pop-ups which users must act on when they visit new websites are designed to be deceiving by nudging users towards consenting and clicking the huge green button saying "yes, I agree", as opposed to the "no" option hidden a couple of clicks away, behind layers of legal jargon. What's more, consent fatigue causes people to close the consent notice just to access content they're interested in as fast as possible. In most cases simply closing the banner is also – counterintuitively – interpreted as "explicit" consent.

**Dark patterns in consent notices** →

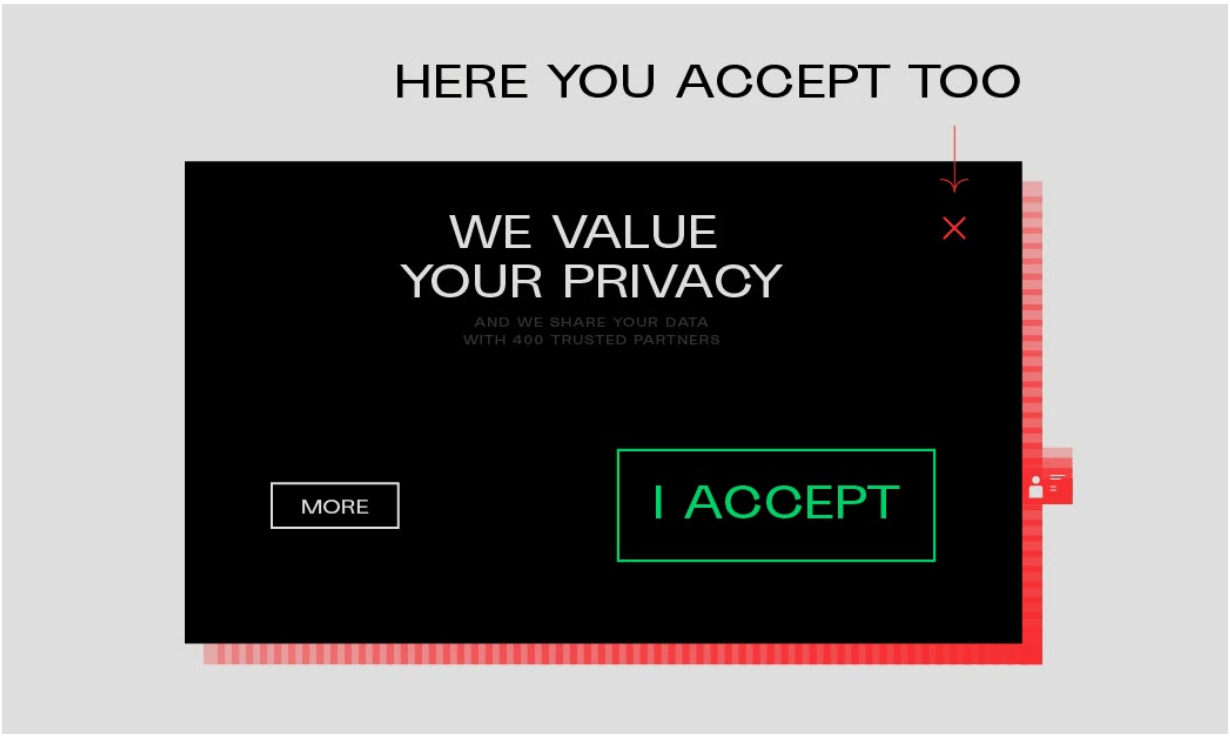


[7] https://adexchanger.com/data-exchanges/tapad-ceo-on-cross-device-graphs-and-where-its-data-comes-from-hint-not-telenor/

→

This phenomenon continues even despite guidelines and decisions from data protection authorities that clarify requirements for freely given consent. According to them, the user must not feel compelled to consent and must be offered real choice[8]. In terms of design, both options (yes and no) should be given the same prominence: refusing consent should not be "one click away", in another layer of the notice, but should require the same amount of steps to select as giving consent[9]. However, it is no surprise that the industry chooses to use dark patterns to impede the rejection of tracking. The experience of a Dutch public broadcaster NPO shows that **only 10% of users agree to tracking cookies if they are given genuine choice**[10].

Second, **consent is not informed**. The requirement for informed consent means that users should understand the consequences of giving consent in order to make an informed decision, including how their data is processed, by whom, and for what purposes. Despite this requirement, companies don't clearly explain how ad targeting works. They share a list of hundreds of "trusted partners" but fail to explain their role and reveal in clear terms that these partners are not just technical intermediaries but companies whose job is to track and profile people across the web. Given how insecure and opaque real-time bidding is, companies that collect users' consent don't have full control over what's going to happen to people's data and who will eventually have access to it. For this reason, it's possible to argue that "consent" to this type of processing will never be sufficiently informed.

Third, **consent does not make everything legal**. Consent does not legitimise targeting which is disproportionate or unfair[11]. Even if consent boxes were designed in a way that seems to offer a genuine yes/no choice, the nature of real-time bidding, where users cannot effectively control their data, means that collecting and processing data for real-time bidding should be considered unlawful.

- **Intrusive profiling goes beyond what is necessary or expected: violation of the principle of data minimisation**

The ad tech industry has developed all sorts of methods that make it possible to track people across the web. **Cross-site tracking** enables ad companies to compile data on what people are reading or watching on a variety of websites. **Cross-device tracking** makes it possible to paint an even more comprehensive picture of a person and their many social

[8] As interpreted by the European Data Protection Board in Guidelines on consent: https://edpb.europa.eu/sites/edpb/files/files/file1/edpb_guidelines_202005_consent_en.pdf

[9] As interpreted by the Danish Data Protection Authority:

https://www.datatilsynet.dk/tilsyn-og-afgoerelser/afgoerelser/2020/feb/dmis-behandling-af-personoplysninger-om-hjemmesidebesoegende/

[10] https://www.ster.nl/media/h5ehvtx3/ster_a-future-without-advertising-cookies.pdf

[11] As interpreted by the European Data Protection Board in Guidelines on consent and Guidelines on targeting social media users:

https://edpb.europa.eu/sites/edpb/files/consultation/edpb_guidelines_202008_onthetargetingofsocialmediausers_en.pdf

roles, for example by combining data from a work computer and a personal smartphone. New invisible ways of tracking, such as browser fingerprinting, are aimed at circumventing limited availability of cookies. Data sharing is a norm – after an ad auction is finished, companies who took part in it exchange identifiers with each other in a process known as "cookie syncing". This means that the next time, they will know that a user who bbc.com recognizes as "User ABC" is in fact the same person as "User 123" in their database.

The idea behind all of these techniques is to gather any and all data that can be used to build detailed online profiles of individual users[12]. These online profiles contain everything the user has ever read online, their IP addresses, precise locations, device information, interactions with content, mouse movements – basically everything that can be collected about someone's online activity. A lot of this information reveals sensitive data, unlawful to be processed without explicit consent. Rather than collecting data that is strictly necessary, ad companies gather insights "just in case" because machine learning algorithms might find an unexpected correlation in large volumes of historical data. As such, **the driving logic of the advertising industry ("more profiling and more targeting") is contrary to the principle of data minimisation enshrined in the GDPR**. This violation has also been confirmed by the British Information Commissioner's Office investigating the ad tech industry: according to this agency data collection for the purposes of real-time bidding is **excessive and mindless**[13].

- **Users are kept in the dark about their marketing profiles: violation of the principle of transparency**

Transparency is one of the most important principles enshrined in the GDPR. If users can't verify their own data that goes into building their profiles, they can't effectively control it and exercise other GDPR rights, such as the right to erasure or the right to complain to a supervisory authority.

However, it is **extremely difficult if not impossible for people to obtain access to their own data collected by online trackers**. Users can't verify inferences made about them that determine their online profiles and "labels" they are assigned by advertisers. To start with, even tech-savvy users will not be able to identify all companies who have access to their data. Some of them set cookies on users' devices and can – with some effort – be identified by their domain names, but others obtain users' data "second-hand" which makes it technically impossible for even the most skilled user to trace them.

What's more, key identifiers used by ad tech companies to single out users and target ads are not even revealed to the people they concern. Because companies use different identifiers for the same user, they can and – as experience shows – do reply "I don't know who you are" when faced with an access request containing an ID that they replaced with another. It is a "catch 22" situation that cannot be reconciled with GDPR requirements.

[12] https://en.panoptykon.org/articles/three-layers-your-digital-profile
[13] https://ico.org.uk/media/about-the-ico/documents/2615156/adtech-real-time-bidding-report-201906.pdf

Even when a user manages to find the right company and the right identifier, experience shows that **a variety of excuses will follow**[14]: either that cookies identify a device and not a user and are therefore not personal data (which is an absurd argument in the age of personal devices), or that data has been anonymised. More often than not companies mistakenly use the term "anonymisation" to describe a process which can in fact be reversed and which still allows them to single out users and deliver a personalised ad. If data had been effectively anonymised, targeting a specific message to a specific user would not have been technically possible. Finally, if a user puts up a fight and manages to obtain their data, datasets are complicated and presented in a way that an average person would not understand, let alone use it to identify potential abuses[15].

- **Real-time bidding is broken by design and by default: violation of data protection by design and by default**

The GDPR introduced the requirements for privacy by design and by default. Data controllers should take privacy into account when designing, implementing and operating any technology which processes personal data. High privacy standards should be offered to users by default, which means that they do not have to do anything (such as change settings) to be protected to the highest possible degree. Targeted advertising couldn't be further from these principles – in fact it is **broken by design and by default**.

For example, companies do not respect the Do Not Track signal[16], despite it being explicitly communicated by the user's browser; Google's advertising settings for publishers send user data to all third parties by default; companies that sell contextual ads cannot opt out of receiving people's personal data through bid requests; and there is no easy way for users to withdraw consent or access data that was collected about them.

### BEYOND DATA PROTECTION: real-time bidding's potential for discrimination

Targeting systems are only as good as the data which fuels them, algorithms are only as good as the people who design them[17] – and inequalities existing in the world are reflected in human and data biases. The nature of behavioural targeting means that ads are "optimised" for those most likely to want to see them or react to them. This is sensible when applied to buyers of power tools for example, but fails when placing job applications in a job market which under-indexes on certain groups, like women, migrants, or people of colour. And it becomes incredibly problematic when certain characteristics such as race or location (even if these criteria are not explicitly labelled as sensitive by targeting

[14] More details: Jef Ausloos, René Mahieu and Michael Veale, "Getting Data Subject Rights Right: A Submission to the European Data Protection Board from International Data Rights Academics, to Inform Regulatory Guidance": https://osf.io/preprints/lawarxiv/e2thg/

[15] See for example: https://privacyinternational.org/long-read/2433/i-asked-online-tracking-company-all-my-data-and-heres-what-i-found

[16] https://gizmodo.com/do-not-track-the-privacy-tool-used-by-millions-of-peop-1828868324

[17] https://medium.com/@szymielewicz/black-boxed-politics-cebc0d5a54ad

algorithms) are used to exclude people from certain services. People will not even know they are discriminated against because they don't have access to ads that they haven't seen.

## ECONOMIC COSTS: real-time bidding exploits publishers and advertisers

- **Publishers give up as much as 70% of their revenue to ad tech companies, while they could be making more money from contextual advertising**

The ad tech industry's ultimate argument to discourage regulators from imposing stricter privacy rules is that advertising supports publishers, allowing them to pay for creating content, and makes it possible for everyone, rich or poor, to access news. However, it has become clear that **only a small fraction of the money that is spent on behaviourally targeted ads goes to publishers**. In what is known as the "ad tech tax", advertising intermediaries capture from 55% (according to industry data[18]) to 70% (as demonstrated by The Guardian[19]) of every dollar spent on an ad.

Some publishers might even be deprived of ad revenue at all because a desire to ensure that advertising appears in "brand safe" environments renders certain types of content "unmonetisable". Advertisers create blacklists of sites and keywords they'd rather not appear next to. For example, a plane company might block keywords such as "terror" or "crash". However, keyword lists also frequently block articles featuring words such as "lesbians" or "muslims", leading to major funding problems for publications for the LGBT and muslim communities. In fact, a British advertising body Outvertising estimates that 73% of safe LGBT content is excluded from funding in this way[20]. While this problem may persist also in contextual advertising systems, it seems to be amplified by the fact that in real-time bidding auctions advertisers do not have a direct relationship with publishers and can't directly negotiate the terms.

It is also **not true that behavioural ads bring significantly more revenue than contextual ads**. An academic case study into an American media conglomerate showed that behaviourally targeted ads translate to only 4% more revenue for publishers[21]. And that does not take into account the additional costs related to serving behavioural ads, e.g. GDPR compliance, maintaining the technical infrastructure or the negative financial consequences of ad fraud and audience arbitrage. Google's estimate that news publishers might lose as much as 62% of their revenue when ads are not behaviourally targeted was criticised by the publishing industry as taken out of context and used as a weapon to further Google's own cause[22].

[18] https://www.adweek.com/digital/3-benefits-resulting-from-ad-tech-tax-cuts/

[19] https://mediatel.co.uk/newsline/2016/10/04/where-did-the-money-go-guardian-buys-its-own-ad-inventory/

[20] https://www.linkedin.com/pulse/save-digital-advertising-world-togetherwecan-jerry-daykin/

[21] https://techcrunch.com/2019/05/31/targeted-ads-offer-little-extra-value-for-online-publishers-study-suggests/

[22] https://digiday.com/media/math-wrong-publishers-grumble-googles-ad-targeting-research/

**Real-life cases demonstrate that contextual advertising can even be more profitable for publishers.** In January 2020 a Dutch public broadcaster NPO switched off all third-party tracking and opted for contextual advertising. As a result, their revenue increased by 61% in comparison with January 2019 and grew to 76% in February 2020 in comparison with February the previous year. Even during the pandemic, when publishers' revenues rapidly dropped[23], NPO's revenue from contextual ads not only did not fall, but continued to be higher in March-May 2020 than in the same period in 2019. According to Ster, NPO's exclusive advertising house, they have also sold out all available ad slots in record time.

Similarly promising results were obtained by Kobler, a Norwegian contextual advertising platform, which found that advertisers buying contextual adverts were prepared to pay 3.4 times more than the average price that Norwegian publishers receive from sales of behavioural ads, with a **final price of 2.3 times more than the average**[24]. Over the 6 months that Kobler analysed the average price advertisers were ready to pay increased by 25% while the ad spend through Kobler's platform quadrupled.

**These examples show that it might no longer be true that if publishers chose to reveal less information about their users, advertisers would simply spend their money elsewhere.**

- **Advertisers fall prey to ad fraud and unreliable metrics**

Advertisers are often portrayed as those who profit the most from targeted advertising by being able to reach customers who will be interested in their products or services. However, we have recently seen an upsurge of evidence challenging this assumption.

One of the biggest worries of advertisers is ad fraud – bots fabricating views, clicks, and engagement for which advertisers have to pay. The total cost of ad fraud is widely debated. The World Federation of Advertisers predicts that by 2025, at $50bn, ad fraud will be the second highest source of income for organised crime after drug traffic. The US Association of National Advertisers estimated that in 2017 it was only $6.5bn. Other sources vary from $34bn to $66bn[25]. One thing is clear: no one really knows what the real number is. In fact, the ISBA – an organisation representing leading UK advertisers – found that **15% of the money spent on ads simply vanished and couldn't be traced back to anyone in the advertising supply chain** – not publishers, not agencies, not ad tech intermediaries[26].

Another worry for advertisers is the rising number of **adblocks**: around 30% of internet users globally installed them[27]. Advertisers are also exploited by ad agencies which engage in so-called **inventory arbitrage**: they buy ad space from publishers and sell it at a higher price to their own clients.

[23] https://www.iab.com/wp-content/uploads/2020/04/IAB-Coronavirus-Impact-Buy-vs.-Sell_4.15.20FINAL.pdf
[24] https://kobler.no/contextual-insights/
[25] https://www.businessofapps.com/ads/ad-fraud/research/ad-fraud-statistics/
[26] https://www.isba.org.uk/media/2424/executive-summary-programmatic-supply-chain-transparency-study.pdf
[27] https://www.socialmediatoday.com/news/global-ad-blocking-behavior-2019-infographic/551716/

Finally, **there is very little evidence that behavioural ads are actually effective**. Researcher found[28] that the accuracy of that targeting is often extremely poor. More so, a 2019 article from The Correspondent claims that **the effectiveness of online ads simply cannot be measured**[29]. The article exposed that benchmarks used to measure the number of clicks and purchases that occur after an ad is viewed are bloated and unreliable. In particular, these benchmarks do not make a distinction between the selection effect (the fact that people were already interested in a particular product, even without the ad) and the advertising effect (the ad is the direct reason why people click or buy). The advertising industry chooses to ignore this in order to stay afloat.

But even if we accept to rely on standard metrics in the industry, it will turn out that **in terms of performance there is no added value for behavioural ads, as opposed to contextual ads**. In fact, multiple studies indicate that contextual ads can perform better: increase the intent to purchase the advertised product by 63%[30], double the number of visits to the advertisers' website[31], and improve the general perception of the brand[32].

### SOCIETAL COSTS: deterioration of the quality of public debate and increasing carbon footprint

Advertising is attention-driven. The more people can see or click on an ad, the more advertisers are willing to pay. **This dynamic is largely responsible for the rise of clickbait and sensationalist content.** The possibility to create blacklists can deprive "controversial" publications of revenue, ultimately deteriorating the diversity and quality of media and public discourse. In-depth reporting on issues considered too negative or difficult may also suffer from lack of funding. Recently the term "coronavirus" has become indesirable for many advertisers which deprived news websites of advertising money despite record numbers of visitors[33].

Fully automated bidding systems are opaque also for their end users – publishers and advertisers. This creates strong incentives for hoax publishers to participate in ad auctions. As a result, **money spent on advertising by the world's biggest brands ends up supporting extremist and fake news content**[34].

[28] https://pubsonline.informs.org/doi/pdf/10.1287/mksc.2019.1188

[29] https://thecorrespondent.com/100/the-new-dot-com-bubble-is-here-its-called-online-advertising/13228924500-22d5fd24

[30] https://www.businesswire.com/news/home/20160316005448/en/New-Study-Reveals-Effective-Brands-Connect-Consumers

[31] https://www.ster.nl/media/h5ehvtx3/ster_a-future-without-advertising-cookies.pdf

[32] https://www.exchangewire.com/blog/2020/04/02/illuma-technologys-contextual-targeting-outperforming-traditional-brand-uplift/

[33] https://www.businessinsider.com/integral-ad-science-doubleverify-help-brands-avoid-coronavirus-news-2020-3?IR=T

[34] https://www.thetimes.co.uk/edition/news/big-brands-fund-terror-knnxfgb98 and https://www.tandfonline.com/doi/full/10.1080/21670811.2018.1556314

Another social cost of online advertising that is likely to be ever more relevant in the coming years is its **impact on climate change**. Vast amounts of (unnecessary) data need to be processed and stored, and that requires a lot of energy. A 2016 study evaluated the carbon footprint of online advertising at 60 mln metric tons[35], which constitutes 10% of total Internet infrastructure emissions, or as much as 60 mln flights between London and New York[36]. That number has inevitably grown since then and will continue to grow, particularly with the increasing use of machine learning, which is hugely energy-intensive. In the times of climate emergency, this aspect of online advertising simply cannot be ignored, especially when there seems to be no justification for such a mass collection and processing of data.

### PART OF A SYSTEMIC PROBLEM: surveillance capitalism and platform domination

This picture would be incomplete without looking more broadly at the online information landscape and data economy. Real-time bidding **upholds business models based on surveillance**, built on the premise that people and intimate information about them can be treated like a commodity. The language used by the ad tech industry is the best illustration of that. Here people's attention ("impressions") is sold at auctions, people are categorised into "segments" which – as one company boasts – can be moulded with their "plasticine-like 92 mln cookies", and "infinite data from infinite devices" stands for intimate details about people's lives.

Even though there are thousands of adtech companies, online advertising remains a **duopoly of Google and Facebook**. By some estimates, the two advertising giants control 84% of the global digital ad market. These ads are mostly sold on their platforms (such as Instagram or YouTube), but Google additionally regulates and operates ad exchanges on external websites. And it doesn't hesitate to use its dominant position for its own benefit. For example, if publishers want to sell ads with the use of Google's infrastructure, they have to accept that Google will use information about visitors of independent websites to make Google's own profiles more detailed and better target ads on Google's own services, such as YouTube[37].

Even when third-party trackers or ad auctions are no longer legal, Google and Facebook will survive thanks to the amount of data they gather about people through their own platforms. This doesn't mean that we should turn a blind eye to privacy violations happening on ad exchanges but makes it inevitable to embed this conversation in a broader discussion on the future of the Internet.

[35] https://www.sciencedirect.com/science/article/pii/S0195925517303505

[36] https://www.theguardian.com/environment/ng-interactive/2019/jul/19/carbon-calculator-how-taking-one-flight-emits-as-much-as-many-people-do-in-a-year

[37] More on Google's unfair practices on ad exchanges: Dina Srinivasan "Why Google Dominates Advertising Markets": https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3500919

# PART 2

# WHAT ARE THE ALTERNATIVES?

# 2. WHAT ARE THE ALTERNATIVES?

> *[...] redesign of online advertising practices for the 21st century will likely extend [...] to privacy-focused user level targeting, and contextual targeting. [...] A system that enables brands to promote the result of their work, one that funds a strong and independent publishing industry, and one that allows users to benefit from the rich personalization that only online media can offer — on their own terms.*

It may come as a surprise that these words were not uttered by a privacy activist or an EU official. Instead, they appeared on the website of Criteo, a large French advertising company that focuses on personalised retargeting. Even the advertising industry, often accused of systemic infringements of privacy laws, has come to terms with imminent change.

Despite these promising declarations from the industry, many discussions about alternatives to the dominant funding model of online publishing still tend to come down to an oversimplified either/or scenario. **Either we stick to behavioural advertising in its current form or people will have to pay for content.** This narrative is to some extent responsible for maintaining a situation in which mass-scale online tracking and profiling is "excused" as part of a rotten compromise between privacy and access to information online.

While people should by all means be encouraged to pay for quality content, success stories that describe increasing revenues from subscriptions are unfortunately mostly limited to large, well-established publishers with a wealthy readership and consistent paywall strategies[38]. Subscription models contribute to diversifying sources of income, which is essential for any business strategy, but in reality a vast majority of publishers' revenues still comes from various forms of advertising[39]. Even micropayments, despite initial enthusiasm, have so far not proved to be a sustainable alternative[40]. **Perhaps it's time to swallow the hard truth: most internet users have been used to paying for content with their attention rather than money for so long that changing their behaviour would require deep, structural changes into the whole information ecosystem.**

These changes are not impossible and should by all means be supported. However, the assumption for this section is that **in the foreseeable future advertising will remain one of the key sources of funding for online publishing**. But it does not mean that publishers are stuck with sharing as much as 70% of their revenues with the ad tech industry, and their readers – with intrusive tracking and profiling. There are a number of advertising models – some already in operation, some for now functioning only as concepts – that are designed to address (at least some of) the underlying problems of real-time bidding.

[38] https://www.nytimes.com/2020/05/06/business/media/new-york-times-earnings-subscriptions-coronavirus.html,

[39] https://whatsnewinpublishing.com/publishers-still-rely-on-traditional-revenue-streams-research-shows/

[40] https://www.wired.com/story/shouldnt-we-all-have-seamless-micropayments-by-now/

This part presents an overview and assessment of the most important types of existing alternatives to the real-time bidding model. While this overview is surely not exhaustive, it is designed to help privacy and human rights advocates, policymakers, smaller publishers, and other organisations that are not at the heart of industry discussions in navigating recent developments. The last section proposes general requirements for privacy-friendly advertising systems and identifies which of the described alternatives meet these criteria. What follows in Part 3, is a discussion on why surveillance-based advertising continues to dominate online spaces and what regulatory interventions are needed for privacy-friendly alternatives to become a viable source of funding for publishers.

ALTERNATIVE ADVERTISING MODELS CAN BE DIVIDED INTO FOUR BROAD CATEGORIES:

1. **security-oriented targeting,**
2. **first-party targeting,**
3. **publishers' collaborations,**
4. **contextual targeting.**

## 2.1 Security-oriented targeting

Sharing personal data with a potentially unlimited number of companies is an inherent feature of open real-time bidding. The industry is booming with products and services that aim to increase the security of various stages of the targeting process, e.g. identity management. However, these tools are intentionally not covered in this section because they only serve as cosmetic interventions that do not fix the systemic problems with the current system. Instead, it's worth taking a look at initiatives that encompass the whole targeting process and propose substantial changes to how ads are delivered today.

A joint premise of the initiatives described below is that ads are targeted to cohorts, rather than to individual users with unique features. A central gatekeeper (e.g. a browser) controls data and administers these cohorts. "Raw" personal data is not shared with advertisers. Ad auctions that normally run between multiple servers (thus sharing data with many companies) are either designed to run locally on a device or on a single server.

- **Google Privacy Sandbox**

In August 2019 Google announced a proposal for reforming the standard of online advertising called the Privacy Sandbox[41]. The flagship idea – so called FLoC (Federated

[41] https://www.blog.google/products/chrome/building-a-more-private-web/

Learning of Cohorts) – is to **target ads not to individuals but to interest groups that users belong to**. The browser would be responsible for monitoring and analysing patterns of users' behaviours across different websites, and creating cohorts of similar users ("flocks"). The flock that the user belongs to, rather than "raw" personal data, would then be shared with other companies.

Privacy Sandbox is not yet in operation but Google – as an advertising company that happens to own a browser used by 60% of all Internet users – is in the right position to dictate the rules of the game for the entire advertising industry. **In fact, in January 2020 the company announced that by 2022 Chrome would phase out support for all third-party cookies[42], looking to make Privacy Sandbox the new standard for the web.**

A similar logic would be used for retargeting[43]. According to Google's proposal originally called TURTLEDOVE[44] and then updated and renamed as Dovekey[45], website operators would add users to a specific interest group based on their activity. A single gatekeeper (a browser or another third party[46]) would then combine and store interests from different websites. Again, it would not give advertisers access to users' personal data – instead, the gatekeeper would send two separate requests to advertisers: one request that contains the context (the URL) of the page where the ad will appear, and another one which contains information about which interest groups the user belongs to. In response to each of these requests, advertisers would send two separate sets of ads that match – respectively – the context and the interests. In the final step, the gatekeeper would run an ad auction and decide which of the pre-uploaded ads is shown.

Some companies critically engage with this framework and propose their own modifications. Three of such initiatives are described below:

- **SPARROW by Criteo**

Criteo, one of Google's main competitors, maintained the bird-related wording by proposing a framework called SPARROW[47], designed to address TURTLEDOVE's shortcomings. According to Criteo, Google's idea fails to give advertisers and publishers the control and transparency they need in the targeting process. This situation could be improved by:

1. **appointing an independent gatekeeper, rather than the browser, to execute ad auctions:** Criteo argues that there's no justification why only browsers could act as gatekeepers in the new model. According to the company, cloud

[42] https://blog.chromium.org/2020/01/building-more-private-web-path-towards.html

[43] Retargeting is a form of targeted advertising by which ads are targeted to consumers based on their previous actions, e.g. visits to an online store.

[44] https://github.com/WICG/turtledove

[45] https://github.com/google/ads-privacy/tree/master/proposals/dovekey

[46] In the original proposal the browser was supposed to play the role of the gatekeeper. After criticism from the industry and a proposal developed by Criteo in response (see below), this idea was updated to enable a different third-party to perform ad auctions.

[47] https://www.criteo.com/blog/sparrow-why-birds-may-play-a-key-role-in-the-future-of-advertising/

service providers or supply side platforms (SSPs) already have the technical infrastructure needed to play the role of the gatekeeper. Criteo surely wants to avoid a situation in which Google – which happens to own the dominant browser on the market (Google Chrome) – becomes even more powerful. Google later adopted this suggestion in the updated version of TURTLEDOVE called Dovekey.

2. enabling lookalike targeting: SPARROW would also allow advertisers to target not only users who have already expressed interest in a particular brand or product (and therefore have been assigned a particular predefined interest group), but also their lookalikes.

- **PARROT and TOUCAN**

The PARRROT proposal[48] was developed by Magnite with the aim to maintain cohort-based targeting proposed in TURTLEDOVE but at the same time allow publishers and SSPs to make key decisions on auctions. These decisions include a number of factors necessary to determine auction outcome, such as ad quality filtering or blocks on certain domains, creatives or advertisers. Technically, it is based on header bidding, i.e. a programmatic technique where publishers offer their ad inventory to multiple ad exchanges at one time before deciding who they will sell the ad to.

TOUCAN is an example of yet another iteration of the retargeting frameworks proposed by Google and Criteo, developed by Adcessible.io as an experiment[49]. The key modification is that it would give users more control over who can add them to interest groups and which ads they see. The basic premise is that users, within their browsers, carry around an "ad profile" composed of a set of advertising preferences, interest groups and an archive of ads served to them. This information would be collected – based on users' explicit consent – by an ad profile administrator directly from publishers' own first-party data (i.e. data that they collected directly from the person visiting their website). This would be technically possible thanks to a solution developed by Apple (Storage Access API). Users could see and control all of their choices (such as adding and removing predefined interest groups and selecting websites that can show personalised ads) via a central dashboard. They could also mandate a representative to do that on their behalf. As such, this proposal is designed to make it impossible for data brokers and other ad tech companies to obscure themselves in long lists of advertising partners displayed in the cookie consent box.

- **Brave Rewards by Brave**

The operator of the Brave browser introduced a model that combines security-oriented advertising and micropayments[50]. In Brave Rewards users see behaviourally targeted ads only after giving their explicit consent. As a "reward" for their attention, they receive 70%

[48] https://github.com/prebid/identity-gatekeeper/blob/master/proposals/PARRROT.md
[49] https://github.com/Greg-Asquith/toucan
[50] https://brave.com/brave-rewards/

of the ad revenue share which they can auto-contribute to publishers that they wish to support. Technically this system allows the browser to infer users' interests based on the content of the sites they visit. Similarly to other proposals in this section, users' data is not shared with advertisers – profiling and targeting is operated in the browser and ad auctions run locally on the user's device.

### ASSESSMENT: FEWER INTERMEDIARIES DOES NOT MEAN A FAIRER AD SYSTEM

Alternatives described above could potentially limit the number of companies that have access to users' personal data. However, technical experts have voiced concerns about the actual security of Google and Criteo's systems which remain unanswered. But even if these proposals do indeed improve security, understood as access to "raw" personal data, it will not be sufficient to address all problems of real-time bidding and ensure a fair and privacy-friendly ad system.

First, these alternatives **do not address concerns that go beyond security**: the potential for discrimination, unfairness, or lack of transparency. Even if users are technically part of a larger cohort (interest group), the effects of targeting (i.e. exposure to a particular targeted message) are **experienced individually and can have a significant individual impact**. In this context it's irrelevant if an ad for an exceptionally high-interest loan has been shown to hundreds of people classified as "in financial trouble" if it affects each and every one of them individually.

None of the proposals goes into detail (at least in publicly available materials) as to **how interest groups are created and how sensitive they can be**. It's unclear whether cohorts would be limited to "interests", such as winter boots, fridges, or sport, or if they could also reveal "personal characteristics", such as high income earners, people with college degrees or new parents. Browsing patterns often reveal sensitive information, e.g. political affiliations or health conditions. Google expressed the intention to remove such information from the mechanism of creating cohorts but admitted that it's not easy to set clear boundaries for what is sensitive and what is not[51]. The risk of exploiting vulnerabilities could perhaps be managed with introducing a minimum number of users that may be part of a specific cohort, but it is likely that it will never be fully eradicated.

Although some of the initiatives offer control tools through which users can manage their choices and interests attributed to them, it is unclear what other insights would be available.

For example, none of the proposals mentions if it would explain to users how

[51] See also: https://www.eff.org/deeplinks/2019/08/dont-play-googles-privacy-sandbox-1

particular interest groups are generated and what specific data was responsible for the user being assigned to a given cohort. These concerns are not unfounded – **they arise from experiences with existing data management tools offered by existing gatekeepers**, Facebook and Google. Companies often choose to reveal only uncontroversial interests that result from data directly provided by users (e.g. via pages or videos they liked), rather than observed, behavioural data (e.g. behavioural patterns) that are interpreted by algorithms to infer characteristics that people did not explicitly reveal.

Last but not least, mandating a single gatekeeper to be responsible for the whole process of ad delivery **poses a threat of developing new intransparent and unaccountable closed ecosystems (walled gardens) or enhancing existing ones.** The latter in particular is a very likely scenario, given Google's plans to phase out third-party cookies. Unless stopped by policymakers, Chrome is likely to become the dominant gatekeeper in the new model of cohort targeting, thus **extending Google's walled garden to cover the entire open web.**

## 2.2 Advertising in the post-cookie world: first-party target ing

Piling privacy problems of real-time bidding and stronger privacy laws have over the last couple of years led browsers and operating systems providers to increase default protections from tracking (e.g. Safari Intelligent Tracking Protection or Firefox Enhanced Tracking Protection). These developments have inspired the search for new ways of monetising content that are not so heavily reliant on third-party tracking. A real panic, however, was caused by Google Chrome's announcement of its plans to **phase out support for all third-party cookies** (the basic tracking technology) by 2022. For advertisers this announcement meant that they would no longer be able to reach as much as 60% of all Internet users and publishers would be even more reliant on Google. **As a result, targeting based on first-party data has become a realistic, widely discussed alternative**.

Online publishers are in a position to directly collect a number of data about their users and their behaviour resulting from interactions with the website or the app, such as the articles they read, liked, or commented on, their login and subscription details, or metadata of their device (e.g. language, browser version, operating system, IP address, location). **Publishers already use their first-party data to sell ads directly to advertisers** (i.e. without the use of ad exchanges). However, a common practice among publishers is to **enhance data about their readers with additional information purchased elsewhere** (e.g. from data brokers), in order to make user profiles richer and more attractive for

advertisers. The New York Times was one of the first (if not the first) of large publishers to recently announce that they would no longer use third-party data for this purpose[52].

Although most publishers are still very much reliant on cross-site tracking, third-party data and intermediaries, these revelations show a direction where the industry is inevitably headed. Some publishers have already built the infrastructure that automates and facilitates first-party targeting. Just to name two examples:

- **Zeus Insights**[53] is a first-party targeting platform developed by the Washington Post. Zeus monitors the content that users are reading or watching and records their activities. In the next step, the publisher combines these insights with existing data to infer what the user might be interested in. The Washington Post plans to license Zeus to other publishers.
- **Forte**[54] is a first-party targeting platform built by Vox Media. The company predicts that in the next two years third-party targeting will become obsolete and advertisers will purchase ads from platforms which rely on data collected directly by publishers, with users' awareness and consent.

### How "shallow" first-party targeting could work

As mentioned above, when publishers engage in direct sales they often rely on additionally purchased data. An interesting proposal for how non-enhanced first-party targeting could work was developed by a data scientist working for one of the largest global publishers. His name is kept confidential because this is not (at least yet) the position of his employer.

This proposal is driven by the idea that we should establish norms on what kind of personalisation is acceptable in the online context. In this model, **the publisher could target ads based only on their first-party data which would in addition be limited to:**

- coarse location (in order to avoid showing a geographically irrelevant advert),
- coarse demographics (gender, age range, income range, job industry),
- shallow behavioural data (user interests based on content he or she has visited at the publisher's website collected over the last 10 days).

More detailed information could be used only if the user actively expressed their consent upon his or her own initiative, preferably through browser settings or other interfaces that limit pop up fatigue.

Ads could be delivered via **a small number of intermediaries** but their role would be

[52] This change applies only to the NYT mobile ecosystem. https://www.axios.com/new-york-times-advertising-792b3cd6-4bdb-47c3-9817-36601211a79d.html

[53] https://www.zeustechnology.com/

[54] https://www.axios.com/vox-media-ad-targeting-platform-forte-b44146cd-de35-4b97-94f7-e341768c792a.html

strictly limited to technical processing: they would not be allowed to enrich data or use it for their own purposes. Companies involved should be obliged to share all identifiers with the controller (publisher) so that the publisher has full knowledge and control over data flows.

Publishers would be required to provide **farm-to-table accountability** for every advert the user sees. This means that users should not have to navigate complicated external settings but simply click on a button within the ad to see all relevant information about targeting. It should include the list of all intermediaries involved in serving this particular ad and targeting parameters used by advertisers. Users could exercise their GDPR rights by direct interaction with a given advert, e.g. by clicking on a button to erase all data that led to targeting them with this ad.

According to this proposal, **publishers could rely on their legitimate interests to target ads.** This suggestion might seem controversial in the light of the GDPR and the existing text of the ePrivacy directive. However, the author argues that the legitimate interest test (required by the GDPR) is passed because the publisher has full control over data flows, targeting is limited to what users may reasonably expect, and meaningful transparency and control tools are available.

### ASSESSMENT: PRIVACY-FRIENDLY ALTERNATIVE THAT REQUIRES A REGULATORY PUSH

"Shallow" first-party targeting ticks a lot of boxes as far as privacy and data protection is concerned, **provided that it is not enhanced by data collected cross-site**. This model builds on the direct relationship between users and publishers. Other companies potentially involved in delivering ads cannot use the data they receive for another purpose and have to share all identifiers with the publisher. This means that the publisher serves as a "one stop shop" for users willing to exercise their GDPR rights even when technical intermediaries are involved. If these rules are not merely contractual but can also be enforced technically, **"shallow" first-party targeting has the potential to eradicate data "broadcasting" and profiling based on cross-site tracking.**

Insights collected from visits to a single publisher's website are likely to be way less intrusive than profiles built by combining browsing data from across the web. If the scope of first-party data is limited to "shallow" behavioural insights collected over a short period of time, the risk of discrimination and exploitation of vulnerabilites is much smaller. If these conditions are clearly communicated to users and if users can easily opt out of behavioural ads and still access the service, shallow first-party targeting is a valid privacy-friendly alternative to real-time bidding.

Financial sustainability of this model remains an open question. The answer to

it **depends on whether the online ad industry will continue to be driven by the logic of accumulating more and more data for targeting**. If so, switching to first-party targeting might only be an option available for large and established publishers who have a significant base of committed readers, varied content that allows publishers to infer users' interests, financial capacity to analyse data (or commission it directly from another entity), and a level of prestige that attracts advertisers. Small or local publishers will not have the same competitive advantage and will most likely not have the means to make this system technically operable. **Part 3 further explains this dynamic and proposes a set of measures that would make first-party targeting more sustainable.**

## 2.3 Breaking the platform dominance: publishers' collaborations

In online advertising **the rules of the game for publishers are dictated by bigger players:** online platforms and the ad tech industry. The former – with their market dominance, reach, and detailed user profiles – attract over 80% of advertisers' money. The latter impose fees that leave publishers with as little as 30% of the original ad budget. This lack of level playing field and an advertising market which promotes mass data collection and detailed targeting, makes even large publishers struggle to attract enough advertisers. As a result, publishers have begun to form partnerships and collaborate with each other. These collaborations are designed to make use of publishers' biggest assets – quality content and direct relationship with users – to create a counterbalance for platform dominance.

> *(...) The only way to effect change is for publishers to work together. When 75-80% of the market is dominated by no more than five companies then you have to collaborate in order to be competitive.*
>
> Damon Reeve, CEO of the Ozone Project[55]

### Joint advertising platforms

One form of collaboration between publishers that has emerged in the last couple of years are joint advertising platforms. These platforms are designed to give advertisers a single entry point to the inventory of all participating publishers. Some initiatives go as far as to share and combine data held by individual publishers, thus enabling advertisers to reach a much wider and more diverse audience across different websites.

[55] https://whatsnewinpublishing.com/damon-reeve-speaks-about-the-ozone-project-the-uks-leading-digital-advertising-marketplace/

Joint advertising platforms usually operate as private marketplaces (PMPs) – automated ad auctions that are open only for specifically invited advertisers, rather than to any company that meets the technical requirements and agrees to the terms of the ad exchange (as is the case for open auctions). Over the last couple of years PMPs have become more attractive for publishers. It is estimated that in 2020 ad spending on private marketplaces will surpass that on open auctions[56].

The most notable examples of joint advertising platforms include:

- **TrustX**[57] – a cooperative private marketplace operator created in 2016 by thirty premium news, sports and entertainment publishers. It is owned by Digital Content Next – a US non-profit publishers' trade association. It connects premium publishers with trusted advertisers. The aim of this initiative is to rebuild trust between publishers, brands and consumers, to erase opacity of the current ad tech practices, as well as to limit Google's and Facebook's dominance in the advertising revenue streams.

- **The Ozone project**[58] – a platform which combines first-party data from over 90 of the UK's premium publishers, such as The Guardian or The Telegraph, into a closed ecosystem. This gives advertisers access to 99% of the UK internet users and allows publishers to compete with leading online platforms.

Data sharing collaborations and joint advertising platforms are also emerging in other countries, such as France (La Place Media[59]), Czech Republic (CPEx[60]) and Denmark (DPN[61]).

### Joint identity management

Joint advertising platforms have so far proved to be more attractive for publishers, but a form of collaboration that goes even further is joint identity management. It aims to create unified standards for how publishers manage their audiences in order to facilitate data sharing and cooperation in other areas (e.g. managing subscriptions). For the purpose of this overview it's worth highlighting two examples:

- **ITEGA**[62] (The Information Trust Exchange Governing Association) aspires to create and govern a new ecosystem for publishers. Developers of ITEGA often compare its role to that of ICANN in governing domain names. It introduces a new infrastructure that would create unified data collection standards for publishers. Part of this infrastructure is a custom new architecture and protocols that would manage identity and facilitate data exchanges between member publishers (such as a universal, cross-site user ID and payment authentifications, managing data rights, and subscriptions).

[56] https://www.emarketer.com/content/private-marketplace-ad-spending-to-surpass-open-exchange-in-2020

[57] https://trustx.org/

[58] https://www.ozoneproject.com/

[59] https://www.adexchanger.com/publishers/how-french-publishers-reclaimed-programmatic-by-creating-la-place-media/

[60] https://www.cpex.cz/en/press-center/cpex-launching/

[61] https://www.exchangewire.com/blog/2013/02/18/top-danish-publishers-launch-dpn-publisher-exchange-to-take-on-fbx-and-adx-in-the-local-market/

[62] https://itega.org/

ITEGA declares that transparency and control tools for users would default to the highest possible privacy standard.

- **DigiTrust** was a proposal for a unique advertising ID from the Interactive Advertising Bureau - an industry representative body composed mainly of advertising companies. However, after a couple of years in operation it was discontinued in July 2020 because the costs of technical support for the system were too high for publishers.

## ASSESSMENT: MORE COMPETITION DOES NOT NECESSARILY ADDRESS THE PATHOLOGIES OF BEHAVIOURAL ADVERTISING

Joint advertising platforms allow publishers to build closer and more lasting relationships with advertisers. As opposed to open auction systems, they give participating publishers control over who is placing bids on their ad inventory. It also allows them to earn more because ad impressions are usually valued at higher prices and fees for intermediaries are smaller. Advertisers who are concerned with brand safety do not face the risk that their ads will appear in hoax or untrustworthy outlets. The risk of ad fraud is also reduced. But a couple of concerns still remain.

First, existing platforms have proved to be an option reserved for premium publishers who can offer valuable audiences and are attractive for brands to advertise on. Small or newly-established publishers might not be invited to join such collaborations.

Second, increased transparency in the advertising supply chain does not automatically translate to positive effects for users' privacy. **Combining data from different publishers could in practice amount to cross-site tracking**. Profiles built this way may in reality be very detailed and potentially reveal users' vulnerabilities or sensitive characteristics. This may still be the case even when individual publishers limit their tracking to "shallow" behavioural data collected over a specified, short period. In the context of existing joint advertising platforms it is not clear whether users explicitly agree to having their data combined and whether they can verify their profiles and control them. For similar reasons, joint identity management systems will also be problematic in the context of privacy and data protection.

In summary, **simply creating more competition will not address the pathologies of behavioural targeting.** Quite the contrary – it may further exacerbate them. At the same time, collaborations which do not amount to cross-site tracking but which leverage the assets of publishers may still be very effective in helping publishers create a counterbalance for platform dominance and should by all means be supported.

# 2.4 Contextual targeting

Contextual targeting is unique against the background of other alternatives because it is based on a completely different logic of targeting. As opposed to behavioural advertising, **ads are tailored not to the user, but to the content of the website.** This alternative is not new at all: it was widely used before cookies and real-time bidding stole the spotlight. Even now, contextual advertising is the foundation of Google Search where ads are targeted to the keywords used in the search query rather than the characteristics of the user. In contextual systems, some personal data may still be collected for the purposes of measuring the effectiveness of ad campaigns, but this data is not used for profiling or targeting. Lack of personal data in the targeting system and – as a result – **no risk of violating the GDPR is exactly why contextual targeting is increasingly gaining popularity again.**

Technically, contextual ads can be sold and delivered with the use of all existing methods, i.e. through direct sales or in an automated, programmatic way on private and open ad auctions. Contextual advertising does not exclude any intermediaries from the process either: SSPs can still be involved in managing the publishers' ad inventory, DSPs can conclude deals on behalf of advertisers and ad exchanges can facilitate the whole process. **The difference is that none of these intermediaries handles personal data anymore.**

While behavioural advertising is mostly concerned with who is visiting a website, contextual advertising focuses on what is the context of the website being visited. **Over the years, the analysis of the context to which ads can be targeted has become very sophisticated.** Machine learning and natural language processing algorithms are applied to find keywords and determine the topic of a website. Ads can also be contextually targeted to video content as specialised software can be used to generate subtitles in order to determine the exact topics that are being discussed.

### ASSESSMENT: PRIVACY-FRIENDLY, EFFECTIVE AND POTENTIALLY MORE PROFITABLE

Contextual targeting seems to be an actual win-win for all groups who should benefit from online advertising: users, publishers and advertisers.

First, contextual targeting is **genuinely privacy-friendly**. Contextual ads do not rely on the behaviour or characteristics of users, so tracking and profiling mechanisms become obsolete. Contextual advertising eradicates the risk of discrimination and manipulation through data. This doesn't mean that the content of the ad itself cannot be unfair or discriminatory. However, the fact that such ads are visible at a specific subpage regardless of who visits it, means that it should be easier for publishers to monitor them and for users to report them. Moreover, with contextual ads users'

online experience can be more satisfying, as they are not bombarded with requests for consent and hundreds of trackers do not slow down their browsers.

Second, because publishers and advertisers do not have to comply with the GDPR, contextual targeting is **much less burdensome**. This may directly translate to improved cost-efficiency as it is no longer necessary to maintain the technical infrastructure (e.g. consent management tools). There is also a high chance that intermediary fees would be reduced – in the current system they are bloated mainly due to the need to compensate ad tech companies for collecting and analysing personal data. In effect, **contextual ads can bring publishers more profit.** And – as discussed in Part 1 – there is real-life evidence to support this claim:

- after opting for contextual advertising the Dutch public broadcaster saw a **68% monthly average revenue increase** compared to the corresponding period the previous year and consistent revenue increase even after the advertising market was severely hit by the Covid-19 pandemic;
- publishers working with a Norwegian contextual advertising platform were **paid on average 2.3 times more for contextual ads** than for behavioural ads.

Third, there is no evidence suggesting that contextual targeting is less effective than behavioural advertising. Quite the contrary – there are more and more examples proving that **contextual ad campaigns perform better than those driven by personal data**.

This is not to say that contextual advertising is bullet-proof. Because no data about users is collected it can pose challenges in terms of frequency capping (i.e. avoiding showing the same user the same ad multiple times[63]). Some content may also be difficult to contextualise. However, the contextual advertising industry has already shown that it is perfectly able to innovate in this regard – it's enough to see the "case studies" sections on the websites of come contextual ad platforms, such as GumGum, Zefr or Illuma.

To sum up, **contextual targeting is a promising alternative to behavioural advertising not only because it is privacy-oriented.** It can also support publishers in a more sustainable way, while at the same time delivering the same, if not better, results for advertisers. However, the dominance of behavioural advertising driven by the interests of ad tech intermediaries and online platforms makes it difficult for contextual advertising platforms to scale up and for publishers to transform their monetisation strategies. Because of this, regulatory intervention is necessary. Please refer to Part 3 for details.

[63] Frequency capping can work if the user allows local storage for non-personal data: https://www.linkedin.com/pulse/frequency-capping-ad-campaign-measurement-under-gdpr-sean-blanchfield/

# 2.5 Bird's eye view: what requirements for a privacy-friendly ad system?

The overview above, even if not exhaustive, demonstrates that the online advertising market is dynamic. Many ideas are circulating, some are already being implemented, and more are certainly to come. In this context, it's worth taking a step back from specific proposals in order to answer the question: **what features should a privacy-friendly or "acceptable" advertising model have?**

Some might argue that there is no such thing as "acceptable" advertising in the first place because advertising – often described as designed to drive consumption and create (artificial) needs – is by definition manipulative and harmful. This brief does neither challenge nor confirm this conviction. Instead, it acknowledges that advertising is there to stay and will continue to be one of the main sources of funding for online publishers. But it doesn't mean that policymakers can't define specific limits or requirements for online advertising models.

**"Acceptable" online advertising should:**

- fully respect fundamental rights, in particular the right to privacy, the protection of personal data, and the right to non-discrimination,
- support quality online publishing by offering fair remuneration for publishers,
- not incentivise creating content that is solely designed to provoke clicks (e.g. by clickbait titles or sensationalist topics/wording),
- reduce the potential of ad fraud,
- minimise necessary computing power.

**In this context there are three essential conditions for alternative advertising models that have the highest potential to contribute to achieving these goals:**

- cross-site tracking and profiling should not be allowed,
- first-party data collection should be fully transparent, subject to users' control and limited in scope and in time,
- publishers and advertisers should have as direct a relationship as possible.

Advertising models incorporating these conditions, if implemented at scale, have the potential to:

- strengthen the position of publishers,
- eliminate intermediaries invisible and unaccountable to users (e.g. data brokers),

- contribute to rebuilding trust between publishers and their users which could translate to increased financial support (e.g. via subscriptions),
- end the invasion of privacy (by disabling mass data collection) and ensure users' information autonomy,
- bring about positive effects for advertisers (e.g. direct relationships contribute to ads appearing in brand safe environments while improved accountability in the supply chain can effectively limit ad fraud).

Three of the alternatives described above meet these conditions and deserve regulatory support:

- contextual advertising,
- "shallow" first-party targeting,
- joint advertising platforms (provided that first-party data is not combined)

The next part takes a closer look at why these advertising models are not implemented at scale and discusses what regulatory interventions are needed to support the uptake of privacy-friendly alternatives.

# PART 3

# TOWARDS PRIVACY-FRIENDLY ALTERNATIVES: RECOMMENDATIONS FOR EUROPEAN POLICYMAKERS

# 3. TOWARDS PRIVACY-FRIENDLY ALTERNATIVES: RECOMMENDATIONS FOR EUROPEAN POLICYMAKERS

## 3.1. Why surveillance-driven advertising persists

Despite the fact that advertising models based on contextual and first-party targeting hold the promise of better privacy protection and more revenues for publishers, **the online advertising market does not currently provide many options for publishers**, particularly small or local ones. Most contextual advertising platforms are still small and it is not as easy to plug into them as it is to embed Google's or other big ad platforms' trackers. Why is that?

For 20 years the advertising industry and leading online platforms have engaged in convincing regulators, users, and publishers that behavioural advertising is the only sustainable way to fund the Internet. The fact that in the 1990s and early 2000s most publishers were still more concerned with prestigious print editions that the web facilitated the expansion of online advertising middlemen. Over the last 8 years alone, the ad tech industry grew from 150 to over 7,000 companies[64]. In the meantime, Silicon Valley startups, unrestrained by regulations, have expanded to become new online empires and main sources of information online, forcing publishers to compete with them for people's attention and advertisers' money. For internet users paying for access to online services with attention instead of money has become normalised to the extent that it is often difficult to imagine a different scenario.

**Given the dominance of this business model and the industry narrative which presents it as the only sustainable option, it is not surprising that regulators and policymakers don't want to worsen the condition of publishers or "break the Internet".**

As a result, the GDPR – which should be a powerful tool in fighting unlimited profiling and targeting – is not properly enforced. Despite a number of complaints having been filed in 2018 and early 2019 in 17 European jurisdictions against Google and IAB – two standard-setters for the online behavioural advertising market – no decisions have been issued yet. **However, the lack of GDPR enforcement, instead of helping publishers, further cements their dependence on the ad tech industry and throttles the development of privacy-friendly alternatives.**

In the best-case scenario, the implementation of privacy-friendly alternatives would bring evidence proving that they constitute profitable and viable options for publishers. As a result, regulators would have no reason to be concerned about the negative side effects of strict law enforcement. Further tightening of data protection rules and the prospect of

[64] https://chiefmartec.com/2019/04/marketing-technology-landscape-supergraphic-2019/

immediate and high fines would create even stronger incentives to innovate and develop alternative technical solutions that do not rely on the exploitation of people's personal data, thus fueling enforcement and policy initiatives with even more evidence.

In reality, however, **we are stuck in a vicious circle**. The failure of data protection authorities to send a regulatory push prevents the mass uptake of alternative solutions. When alternatives are not implemented at scale, gathering enough evidence to undermine the ad tech industry's lobbying is a massive challenge. This further consolidates the market driven by the logic of more data and more personalisation.

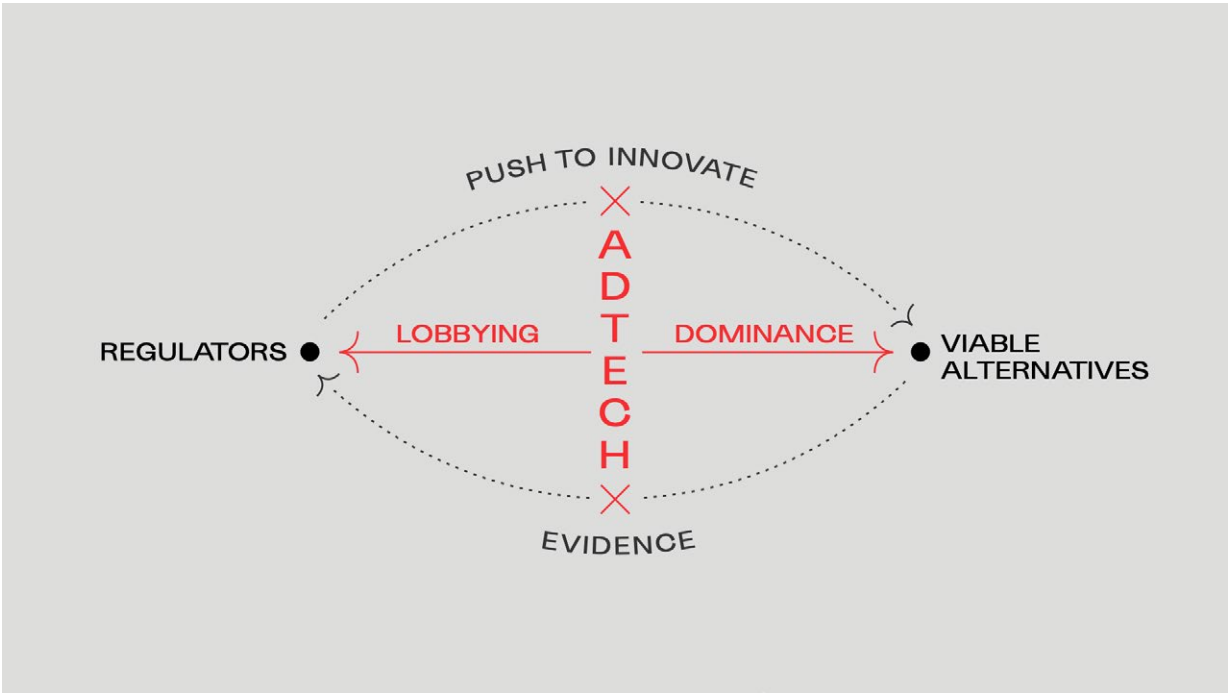


## 3.2. Recommendations

The answer to the title question of this brief – to track or not to track – is not straightforward, given the complex dynamics described in the previous sections. A simple call on publishers to adopt privacy-friendly alternatives ignores the fact that in the current market landscape this move constitutes an understandable business risk or might not even be available at all. At the same time, the pathologies of behavioural advertising cannot be eliminated by encouraging competition on the data market. In fact, the very nature of this market is in collision with the values that the EU should promote: respect for fundamental rights and fostering a healthy digital sphere.

**Any effective solution should therefore aim to fix the cause, not the symptoms**. In this context, only the enforcement of existing laws, paired with new rules which create strong regulatory incentives for the adoption of privacy-friendly alternatives, will eliminate the negative individual and societal consequences of advertising based on surveillance.

**Therefore, on the general level, the EU should:**

- ensure that the GDPR is consistently enforced in all Member States;
- create a coherent and systemic regulatory framework designed to support the uptake of privacy-friendly advertising models which includes both:
  1. effective limits on cross-site tracking (ePrivacy regulation), and
  2. effective restriction of the power of online platforms (the Digital Services Act package);
- make use of financial incentives and other soft measures to promote privacy-friendly advertising models.

The next sections explain these recommendations in more detail.

### 3.2.1. Plugging the GDPR enforcement gap

Strong and fast enforcement of the GDPR is necessary to create room for privacy-friendly alternatives to flourish. The European Union should ensure that national Data Protection Authorities have the tools and resources necessary for that. However, experience with GDPR enforcement so far shows that:

- the law is not consistently applied in all Member States: despite a number of decisions questioning practices of individual ad tech companies[65] or dark patterns in consent notices[66], these practices prevail in other countries,
- national data protection authorities do not engage in joint operations to investigate data processing practices which affect all EU citizens,
- authorities do not have the capacity to investigate big tech, both in terms of financial resources (half of all national DPAs have budgets below 5 mln euros) and in terms of technical expertise needed to investigate technical cases[67].

The ad tech industry engages in lobbying with regulators investigating its data practices. As an illustration, the Polish branch of the Interactive Advertising Bureau provided a workshop on the technical aspects of real-time bidding for the Polish Data Protection Authority. The workshop was held after the Polish DPA had already received Panoptykon Foundation's complaint against the IAB's practices related to the very subject matter that was being discussed. In response to Panoptykon's FOI request, the authority refused to share any

[65] https://iapp.org/news/a/the-vectaury-decision-is-not-an-obituary-for-digital-advertising/

[66] https://www.datatilsynet.dk/presse-og-nyheder/nyhedsarkiv/2020/feb/nye-retningslinjer-om-behandling-af-personoplysninger-om-hjemmesidebesoegende/

[67] According to data obtained by a browser operator Brave, only six of Europe's 28 DPAs have more than 10 tech specialists on board:
https://brave.com/dpa-report-2020/

materials produced as part of the workshop and claimed it did not have plans to hold a similar workshop with representatives of civil society. Having the IAB "teach" DPAs about real-time bidding is like asking Facebook to provide DPAs with a workshop about political microtargeting in the midst of the Cambridge Analytica scandal.

Taking these problems into account, the European Commission should ensure that:

- Member States provide enough funding to the Data Protection Authorities so that they are able to investigate complex technical matters and multinational tech corporations,
- national DPAs can directly apply for additional targeted funding administered by the European Commission,
- the European Data Protection Board provides independent technical training to national DPAs,
- procedural law facilitates effective collaboration and joint operations.

### 3.2.2. Creating incentives for privacy-friendly advertising in the ePrivacy regulation

Three years since the adoption of the ePrivacy regulation by the European Parliament, the Council has still (as of November 2020) not reached the general approach needed to start the trialogue negotiations. In the meantime, the advertising industry engages in aggressive lobbying against the regulation. According to Corporate Europe Observatory, ePrivacy is one of the most lobbied against files and almost all lobbying efforts come from the industry[68]. Concerns about the regulatory capture of the Council have been expressed even by the European Commission.

It's understandable that European policymakers would like to find a happy medium in which both online privacy and publishers' interests are preserved. But as this brief demonstrates, the existing behavioural advertising model is not such a happy medium – it is detrimental not only to individual privacy, but also to the long-term financial sustainability of publishers. **The ePrivacy regulation, instead of maintaining the status quo, should create strong incentives for the uptake of privacy-friendly and sustainable alternatives: contextual advertising and first-party targeting.**

**To this end, the ePrivacy regulation should introduce the following rules:**

- prohibit cross-site tracking and targeting,
- require publishers to obtain users' consent for first-party targeting,
- prohibit cookie walls,
- enable Internet users to express their preferences through automated means,
- use future-proof wording which encompasses all tracking technologies, not only cookies.

[68] https://corporateeurope.org/en/power-lobbies/2018/06/shutting-down-eprivacy-lobby-bandwagon-targets-council

The definition of cross-site tracking and targeting should be constructed in a way that would in practice lead to **outlawing all forms of combining users' activity from different websites, apps, services or devices for advertising purposes, regardless of what technology is used**. The goal of this prohibition should not only be to eliminate third-party trackers, but to make it impossible for any single company or a group of companies to collect and combine insights about a person across the web. Otherwise, taking into account recent developments on the market (esp. Google's announcement of the Privacy Sandbox[69]), there is a risk that gatekeepers would abuse their positions and build unaccountable walled gardens. It's important to note that this measure would also **limit the targeting power of major online platforms**, such as Facebook and Google, because they would no longer be able to track people on external websites and combine this data with insights gathered on their own platforms (e.g. Instagram, YouTube). A similar limit to cross-contextual targeting is likely to be introduced in the new California Privacy Rights Act[70].

In practical terms, **first-party targeting would become the only allowed form of behavioural advertising**. However, publishers should still obtain users' consent and should not be able to implement cookie walls. These two measures would create a strong incentive for publishers to use first-party targeting in a transparent and accountable manner in order to gain users' trust necessary for obtaining consent. This requirement is also likely to contribute to the uptake of the most privacy-friendly form of advertising: contextual targeting. This was in fact the case for the Dutch public broadcaster NPO which observed that only 1 in 10 users consented to behavioural advertising if given genuine choice. The company decided to implement contextual targeting and actually recorded a significant increase in revenues[71].

In order to **address users' pop-up fatigue** and support meaningful consent ePrivacy should provide for **users to express their preferences through automated means** (e.g. browser settings) **and require all publishers to respect this technical signal.** In this context the situation in California where the CCPA gave consumers the right to opt out of any business selling their data can serve as an inspiration. Regulations interpreting the law, issued by the attorney general, concluded that businesses have to respect universal opt-outs sent by a browser or by the consumer's device. This regulatory push alone led to the development of the technical standard called Global Privacy Control which is likely to finally make "do not track" a reality[72]. In the European context, we should draw lessons from the challenges with the application of Article 21(5) of the GDPR, which introduced a similar measure for opting out of direct marketing but did not specify the technical standard that should be used. Therefore, **it's important to ensure that there is a central, EU-level institution responsible for issuing or accepting appropriate technical specifications.** This role could be performed

[69] See Part 2.1.

[70] https://www.linkedin.com/pulse/california-privacy-rights-act-define-limit-behavioral-johnny-ryan/?trackingId=9xVbA3iMe1ADiy8nz672Ww%3D%3D

[71] See more in Parts 1.3 and 2.4.

[72] https://www.wired.com/story/global-privacy-control-launches-do-not-track-is-back/

by the European Commission or a specialised unit within an appropriate EU agency, e.g. the European Data Protection Board.

The European Parliament – in the context of the DSA discussions – expressed support for an even stricter measure: an overall ban on behavioural advertising. This suggests that all options might still be on the table. However, if it is not politically feasible to introduce topics which have not been considered in the previous versions of the ePrivacy text, such as a ban on cross-site tracking, **the second best option** that would create stronger incentives for the adoption of privacy-friendly advertising systems would be for the law to:

- enable online newspapers (not all online publishers) to rely on their legitimate interest for the purposes of shallow first-party targeting,
- require them to respect automated exercise of the right to object.

In any case, an exception for news publishers in which tracking for advertising purposes is considered necessary for the provision of the service should be a clear red line[73].

### 3.2.3. Limiting platform power: Digital Services Act package

Although this brief deals mostly with open web advertising, it's impossible to ignore the fact that the market logic which strips publishers of advertising revenue and internet users of their privacy is dictated by big online platforms, such as Google and Facebook. Publishers – forced to compete for users' attention with online giants which monitor, analyse and compare intimate details of the lives of 2 billion people – are fighting a war for survival, not for prosperity. As such, **any regulatory response should treat online advertising as a system of interconnected vessels**. More specifically, this means that the ePrivacy regulation, even if it introduces very strict measures, will not be sufficient to bring about long-lasting, sustainable change in the condition of publishers and EU citizens. It needs to be complemented with strict GDPR enforcement towards online platforms and regulations which shift the power balance in the online environment.

In this context, the Digital Services Act package, especially the envisioned Digital Markets Act, should aim to significantly limit the power of online platforms to collect, combine and exploit data for the purposes of delivering targeted advertising as well as for developing and improving platforms' services. In particular, the DSA package should introduce specific obligations and prohibitions for large online platforms, e.g. impose a clear standard of interoperability and prohibit practices and behaviours that impede access to the market for competitors. More detailed ideas on how to limit platform power are available in Panoptykon Foundation's response to the consultation on the Digital Services Act[74].

[73] Such a proposal was included in Recital 21 of the Finnish presidency text and in the early versions of the German presidency text: "In some cases the use of processing and storage capabilities of terminal equipment and the collection of information from end-users' terminal equipment may also be necessary for providing an information society service, requested by the end-user, such as services provided to safeguard freedom of expression and information including for journalistic purposes, such as online newspaper or other press publications as defined in Article 2(4) of Directive (EU) 2019/790, that is wholly or mainly financed by advertising provided that, in addition, the end-user has been provided with clear, precise and user-friendly information about the purposes of cookies or similar techniques and has accepted such use".

[74] https://panoptykon.org/sites/default/files/stanowiska/panoptykon_dsa_consultation_submission_08.09.2020_final.pdf

### 3.2.4. Promoting the uptake of alternatives by soft measures

Apart from binding regulation, European policymakers should also consider using soft measures to promote the uptake of privacy-friendly alternatives. In particular the EU should:

- fund research and innovation projects that aim to develop, test or implement privacy-friendly advertising and content monetisation models for publishers,
- monitor and research the societal harms of behavioural advertising,
- require campaigns or initiatives benefiting from public funding to use contextual and not behavioural advertising for promotion,
- set up a grant programme for small EU-based publishers to support their transition from real-time bidding to contextual or first-party targeting or other monetisation models,
- develop guidelines or best practices for publishers and foster publishers' collaborations,
- promote digital and media literacy aimed at inspiring citizens to financially support quality publications.

# ABOUT

**Karolina Iwańska** is a lawyer and policy analyst at Panoptykon Foundation. Since 2018 she has been leading Panoptykon's investigation into the online advertising industry's practices related to data protection and privacy. In 2019/20 Karolina was a Mozilla EU Tech Policy Fellow examining alternatives to the dominant behavioural advertising model and developing policy recommendations aimed at promoting privacy-friendly funding models for online publishers. This publication presents the results of this work.

@ka_iwanska

**Panoptykon Foundation** is a Polish watchdog NGO with a mission to protect fundamental rights in the context of growing surveillance and fast-changing information technologies. We believe in "watching the watchers" and consider data a source of power. Therefore we keep an eye on entities that collect and use personal data in order to influence people (public authorities, intelligence agencies, business corporations). On the legal front we keep track of new legislation, develop alternative regulatory solutions and intervene to protect human rights. In our advocacy we address both policymakers and business lobbies. Through our research and investigations we expose risks related to commercial and public surveillance in order to raise public awareness. We visualize collected data and engage in artistic collaborations in order to reach broader audiences. Since 2010 we have been an active member of European Digital Rights (EDRi).

@panoptykon
www.panoptykon.org

This publication was supported by Mozilla Foundation. 

For enquiries please contact:
fundacja [at] panoptykon.org or karolina.iwanska [at] panoptykon.org.

Graphic design by Michał Małolepszy (www.littlebetter.pl)