

The after party: Cynical resignation in Adtech's pivot to privacy

Big Data & Society
 July–December: 1–14
 © The Author(s) 2023
 Article reuse guidelines:
sagepub.com/journals-permissions
 DOI: 10.1177/20539517231203665
journals.sagepub.com/home/bds



Lee McGuigan¹ , Sarah Myers West², Ido Sivan-Sevilla³ 
 and Patrick Parham³

Abstract

Digital advertising and technology companies are resigned to a new privacy imperative. They are bracing for a world where third-party tracking will be restricted by design or by law. Digital resignation typically refers to how companies cultivate a sense of powerlessness about privacy among internet users. Our paper looks through this optic from the other end of the lens: How is the digital advertising industry coping with the *increasing salience* of privacy? Recent developments have forced companies to implement “privacy-preserving” designs—or at least promise some semblance of privacy. Yet, the industry remains dependent on flows of data and means of identification to enable still-desired targeting, measurement, and optimization. Our paper analyzes this contradiction by looking at systems that aim to replicate existing functionalities while protecting user “privacy.” We call this a form of “cynical resignation” and characterize its key maneuvers as follows: (a) *sanitizing surveillance*; (b) *party-hopping*; and (c) *sabotage*. We argue that this “cynical resignation” to a privacy imperative represents a policy failure. In the absence of decisive interventions into the underlying business models of data capitalism, companies offer techno-solutionism and self-regulations that seem to conform to new laws and norms while reinforcing commitments to data-driven personalization. This may benefit the largest tech companies, since their privileged access to first-party data will make more companies reliant on them, and their computational power will be even more valuable in a world where modeling is used to compensate for the loss of third-party data and traditional methods of personal identification.

Keywords

Surveillance, Adtech, privacy, resignation, platforms, data capitalism

This article is a part of special theme on Digital Resignation and Privacy Cynicism. To see a full list of all articles in this special theme, please click here: <https://journals.sagepub.com/page/bds/digitalresignationandprivacycynicism>

Introduction

It looks like last-call at adtech's big-data bonanza. The digital economy is bracing for a world where third-party tracking may be restricted by design or by law (Keller, 2022; Veale and Borgesius, 2022). Platform companies like Google and Meta spent the last decade turning behavioral data into historic advertising fortunes (Crain, 2019; West, 2019). Today, they are trying to protect those fortunes in the face of regulatory headwinds, infrastructural changes, cultural pressures, and economic turmoil

¹Hussman School of Journalism and Media, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

²AI Now Institute, New York, NY, USA

³College of Information Studies, University of Maryland, College Park, MD, USA

Corresponding author:

Lee McGuigan, Hussman School of Journalism and Media, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA.

Email: leemcg@unc.edu



(Graham, 2022; Lomas, 2022). As marketers lose easy access to cheap data from third-party cookies and cross-app conversion tracking, industry leaders are publicizing a range of “privacy-preserving” technological solutions centered around first-party data and inferential machine-learning systems (Apple, 2021; Schiff, 2023). These solutions promise to replicate existing capabilities without flouting users’ expectations or data governance laws. It appears that the ad-supported tech sector is becoming resigned to a new privacy imperative (Bindra, 2021; IAB Tech Lab, 2022; Mudd, 2021).

To parse the appearance from the realities, we analyze this “pivot to privacy” using critical concepts of resignation and cynicism. Digital resignation typically refers to how companies cultivate a sense of powerlessness about privacy among individuals, leaving data subjects exposed to unwanted tracking and exploitation (Draper, 2017; Draper and Turow, 2019). It gestures to the normalization of the belief that people should not expect privacy online. Our paper looks through this optic from the other end of the lens: How are digital tech companies and the advertising industry writ large coping with the *increasing salience* of privacy as a policy and public-relations issue? Recent developments have forced companies that monetize data to begin implementing “privacy-preserving” designs—or to at least declare respect for user privacy. At the same time, the digital economy remains dependent on flows of data and means of identification to enable the ad targeting, measurement, and optimization that market actors still demand. How are adtech and platform companies reconciling this contradiction?

Our paper shows how these companies are performing “privacy” without meaningfully changing the dynamics of data capitalism. Examining some emergent technological and self-regulatory solutions, including first-party tracking, cookie-less identification, and “clean room” data laundering, we argue that these solutions, while including some good-faith efforts, ultimately amount to a form of “privacy cynicism” (Hoffmann et al., 2016). They proceed from narrow definitions of privacy and sidestep related concerns about discrimination, power asymmetries, and the institutional structure of media systems. Furthermore, companies with large incumbencies or strategic positions along advertising and data supply chains are using privacy as a pretense to take steps that will make the whole ecosystem more dependent on firms that own first-party data, control infrastructural bottlenecks, or market well-branded “artificial intelligence” services.

“Privacy-preserving” adtech thus represents a cynical resignation. Companies concede that they must react to new laws and cultural norms, but their solutions aim to reproduce the status quo or even justify anticompetitive behavior. A main contribution in our paper is to characterize some key strategic maneuvers being executed by means of these privacy solutions. We call these maneuvers: (a)

sanitizing surveillance (using computational techniques to obfuscate data flows); (b) *party-hopping* (pursuing more invasive first-party tracking, or turning third-party data into first-party data via partnerships and acquisitions); and (c) *sabotage* (restricting data access to disadvantage rivals and increase market power). Priorities and tactics vary across firms and industrial sectors, but, overall, we observe efforts to contain and exploit the force of today’s privacy imperative.

These theatrics have real implications. As tech companies are becoming resigned to legal, technological, and cultural regimes wherein gestures toward privacy are necessary and perhaps lucrative, these same companies are moving to define the boundaries of what privacy will mean and how it will work. Their maneuvers are normalizing a paradigm that purports to “solve” privacy using “privacy-enhancing technologies” (PETs). This technosolutionism is rooted in an information-security mindset that views the world in terms of adversaries and vulnerabilities, not in terms of social relations and power structures. Many critical theories regard privacy as a profoundly *public* matter (Cohen, 2013; Nissenbaum, 2009). By contrast, the pivot to first-party surveillance and proprietary PETs represents a *privatized* way of doing privacy. Divorced from a political-economic critique of surveillance business models, this approach advances an abstract view of “privacy without power” (Marwick, 2022), offering technical fixes for what are actually social problems. Even as these fixes remedy some real abuses, they reinforce marketing and technology companies’ cynical disposition toward data governance—that despite promises about empowering sovereign consumers, “privacy lies outside the purview of democracy, as do most of the important decisions about the structure and values of our communications infrastructures” (Crain, 2021). Without intervention to change the status quo, the companies best-positioned to leverage first-party datasets, control infrastructural junctions, and exercise massive computing power will determine what privacy means for end users, and how competitors and clients will have to conform.

The following sections analyze this cynical pivot to privacy. Our arguments are based on close readings of public-facing texts, including technical documentation, trade reportage, and publicity. We examined company websites, blogs, and press releases, followed coverage of relevant industry developments in popular newspapers and trade publications (e.g. *AdExchanger* and *Digiday*), and synthesized these findings with insights from an extensive review of critical literature. We focus on materials from the last five years and from advanced capitalist economies in the Global North because legal, technical, and cultural developments have brought concerns about privacy in adtech to the forefront of public and business discourses in these places over that timeframe. Of course, adtech spans the globe, and human rights to privacy are universal, so these issues resonate across jurisdictions, with local specificities. Further research

should extend, qualify, or challenge our claims in light of contexts beyond this paper's scope.

Our research approach led us to study recent privacy initiatives by adtech companies and to inductively develop a set of critical categories to characterize them. The empirical examples we present are not drawn from a representative sample, but they do illustrate key themes that emerged from our close attention to industrial trends. The initiatives we examine are, in some cases, prospective or experimental. Nevertheless, as companies articulate what "privacy-preserving" means to them, we get a chance to scrutinize corporate values and push for more rigorous responses. To be clear, some of the techniques discussed below may curb certain harms, and we believe that many system designers working on these solutions are well-intentioned. The space examined here is not monolithic. But the patterns we detail clearly illustrate that "privacy-preserving adtech" is largely about preserving adtech profits.

The new privacy imperative and the old data dependencies

Major sectors of contemporary capitalism center around the extraction and processing of personal data (Birch, 2020; West, 2019). In the technoscientific culture of modern business, organizations exercise what Fourcade and Healy (2017) call a "data imperative," collecting any available data assets in the hope that they can be exploited now or later. Platform enclosure of this information and the infrastructures for circulating, storing, and computing actionable insights from it has concentrated political-economic power among large companies functioning like for-profit knowledge and logistical utilities (Plantin et al., 2018; Viljoen et al., 2021; Wu and Taneja, 2021). Interlocking dependencies among data processors and the customers of data-derived services have extended platform institutions throughout online media industries (Nieborg and Poell, 2018; van der Vlist and Helmond, 2021).

Advertising has been one of the most reliable ways to convert data assets into revenue (Mellet and Beauvisage, 2020), and marketing today is completely intertwined with surveillance and discrimination (Darmody and Zwick, 2020; Gandy, 2021; Wachter, 2020). Chasing explosive growth in the scale and speed of personalization, the digital advertising industry became acutely dependent on forms of user identification and monitoring that facilitate core operations, such as audience segmentation, the valuation of impressions, and the analytics that power probabilistic bets and optimization of decisions about whom to target, how to bid at auction, and so on. This business model helped tracking mechanisms—and partners who facilitate and profit from them—become pervasive on websites and mobile apps. Meanwhile, ad-supported social media thrived on a continuous supply of consumer data

(Crain and Cohen, 2023), so much so that Jasmine McNealy (2022: 1678) describes platforms like Facebook as "large-scale phishing operations designed to collect information about users deceptively and surreptitiously." These dynamics helped nearly-ubiquitous surveillance and scoring become part of everyday experience in "the society of algorithms" (Burrell and Fourcade, 2021). Automating the conversion of almost any online behavior into investment opportunities, adtech processes were integrated into the capital circuits of the commercial internet (Crain, 2019; McGuigan, 2023). As Tanya Kant (2021) suggests, "It is not an overstatement to propose that user targeting underpins the online economy as we know it."

After a long binge, however, a hangover is setting in. The data collection and processing at the heart of these business models have slammed into the counterforce of an emergent legal regime. That regime is exemplified by existing laws, such as the General Data Protection Regulation (GDPR) and the California Privacy Rights Act (CPRA), as well as proposed rulemaking by the Federal Trade Commission and the U.S. Congress, and a growing number of regulatory judgements and court proceedings that challenge aspects of corporate conduct (Swant, 2023). Even though regulatory actions are rooted in different privacy ideologies across jurisdictions (Whitman, 2004), data governance requirements are pushing towards global changes in adtech. Data flows that were enabled by default are now subject to legal frictions that portend existential problems for market configurations premised on circulation, speed, and automation. E.U. regulators, for example, recently fined Meta \$414 million in a ruling that effectively bars Facebook and Instagram from making users accept personalized advertising as part of their terms of service (Satariano, 2023). Other data protection authorities have ruled against the legality of pixel tracking, industry consent frameworks, and the Real-Time Bidding protocol central to programmatic advertising (Lomas, 2023). Despite uneven GDPR enforcement (Mizarhi-Borohovich et al., 2023; Sivan-Sevilla, 2022) and corporate efforts to influence public and legal opinions (Corporate Europe Observatory, 2022), it is clear that business-as-usual runs afoul of the law (Veale and Borgesius, 2022).

Corporate policies, encoded in technical infrastructures, are also interrupting data conduits. Two major developments stand out: (a) Apple's App Tracking Transparency (ATT) protocol requires apps and ad intermediaries to obtain consent before tracking users across apps; and (b) the impending elimination of third-party cookies on web browsers threatens to break many targeting and measurement functionalities in web advertising. We discuss each in turn.

ATT was rolled out alongside iOS 14.5, marking a significant change in the way that apps collect and share data. User-level tracking on mobile apps has been accomplished using device identifiers that marketers and apps could access by default. ATT flips that, requiring opt-in

permission from users. If users do not opt-in, the buyers and sellers of ad inventory will not have access to device IDs and the associated data. This “signal loss” interrupts the routine practice of deciding how to bid on ad inventory based on calculations of individuals’ probable value. For example, developers who advertise their apps want to recognize the population of users who spend money on in-app purchases—and then to attribute the revenue those users generate to the advertising event(s) that motivated them to download the app. ATT has been catastrophic for major brokers in this marketplace. Analysts reckon that it may have cost Meta \$10 billion in revenue in 2022 (Bobrowsky, 2022). As we discuss below, however, Apple’s budding advertising business stands to benefit from its position as an authorized data processor in the iOS ecosystem (Haggin, 2021). Apple can now provide targeting and measurement capabilities that are no longer available to companies unable to obtain user permission—an obligation from which Apple exempts itself, since the company does not define its first-party data collection as a privacy violation (Kollnig et al., 2022).

Meanwhile, web advertising is preparing for the so-called “deprecation” of third-party cookies, which have been the main mechanisms for identifying, tracking, and profiling web users and for evaluating advertising opportunities (Turow, 2011). Following its peers, Google has pledged to no longer support third-party cookies in the Chrome browser. Although the company has delayed this phasing out twice (Lawler, 2022), Google is preparing for post-cookie adtech with a suite of proposals collectively termed the “Privacy Sandbox.” This comprises a range of services to facilitate advertising use cases, such as retargeting, audience segmentation, and measurement/attribution, based on data that is collected, processed, and stored on the “client side” (via the browser). Its first high-profile solution, which sorted users into “cohorts” of similar people, crashed under critical blowback—both in terms of its privacy weaknesses (Wodinsky, 2021a) and its likelihood of making other companies even more reliant on Google’s data and targeting abilities (Wodinsky, 2021b). Despite these hiccups, Google is proceeding with a corporate-led realignment of data governance in adtech.

The writing on the wall appears portentous. Circuits of feedback, analysis, and optimization that the digital advertising industry packaged into revolutionary promises of progress and growth are no longer guaranteed. Analysts at McKinsey recently identified “privacy protection” as “one of the megatrends shaping the evolution of the web” (Ahuja et al., 2022). Another report called the demise of third-party tracking a “reckoning” that will “threaten the US digital advertising industry—and compel its transformation” (Brodherson et al., 2021). Yet, despite the bold rhetoric, adtech firms have pivoted to privacy with a mix of hesitation and finesse. Their path forward is hedged by thickets of deep-rooted mechanisms and mindsets.

Adtech is a vast sociotechnical assemblage, difficult to reform. The technologies and partnerships that have allowed many companies to leverage or monetize personal data are infrastructural to digital advertising markets and professional routines (Alaimo, 2022; Mellet and Beauvisage, 2020). They are also entangled in the plumbing that distributes media content and facilitates access to mediated experiences (Braun, 2019; van der Vlist and Helmond, 2021). Control over these logistical and informational infrastructures is highly stratified. Google, for example, dominates markets for adtech services, and its position in audience-aggregation via search, YouTube, and beyond enables it to bundle products and steer clients’ money toward its own properties (Srinivasan, 2020). The company operates like a for-profit utility, absent regulatory constraints, and it appears disinclined to relinquish its status (Hagey et al., 2022). Large adtech platforms, in short, manage whole ecosystems, leveraging data across sites and services and setting rules that affect many other businesses.

Adtech’s maintenance also draws on a staunch ideological defense. Many professionals, academics, and policymakers treat it like a law of nature that advertising should be the primary mechanism for financing content and services online. Industry spokespeople insist that advertising—especially targeted digital advertising—not only sustains the “free” internet, but actually anchors a load-bearing pillar of the entire US economy (IAB, 2022). It follows that as long as certain security risks can be mitigated with technical designs and industry self-regulation, putting data to work for personalization and optimization yields unassailable market efficiencies and increases in consumer welfare (Deighton and Kornfeld, 2020). Overall, the adtech catechism holds that robust privacy rules threaten to disfigure the internet, sending disastrous shockwaves through adjoining industries and whole ways of life (Winslow, 2022; see Baik, 2020). While the Ads Privacy Lead at Google says “moving to a more private model isn’t just an option—it’s a necessity,” she maintains that personalized advertising is required to “pay for the web everyone wants” (Norburn, 2022). Thus, even as the ad-supported technology sector will admit that privacy is a “fundamental human right” (Apple, 2022), a massive inertia weighs against any structural transformation.

These observations reveal a different privacy paradox, wherein companies enriched by commercial surveillance at once promise a privacy-first future and yet also affirm their commitment to data-driven advertising as the essential economic foundation of digital citizenship. We argue, then, that this pivot to privacy can be characterized as a cynical resignation. Cynicism, in our usage, differs somewhat from the typical formulation of privacy cynicism. The latter describes individuals’ subjective responses to their perceived lack of agency (Hoffmann et al., 2016; Lutz et al., 2020); by contrast, we observe companies acting within the constraints of law and public opinion but also

asserting their agency to define or protect operating conditions that favor their interests. While these companies are compelled to react to a new privacy imperative, they still feel empowered to exercise control over the direction and intensity of their reactions—and over the shaping of privacy’s practical contours. Advertising companies perform resignation in public, admitting, for example, that “the web ecosystem is at risk if privacy practices do not keep up with changing expectations” (Bindra, 2021). And, yet, they continue to lobby for legal frameworks that would preserve existing business models and/or impose compliance costs which they can bear but competitors cannot. On the other hand, adtech’s pivot does mirror the usual definition of privacy cynicism, albeit from the perspective of data controllers rather than data subjects, in that we find high privacy concerns co-existing with only modest countermeasures. Keller (2022: 11–12) summarizes the tension playing out internationally: “the advertising sector is making changes to accommodate demands from regulators and courts for better data protection. Yet, it is fair to say that the online commercial sector in the UK, EU and the United States remains largely committed to the idea that reformed behavioural advertising is a legitimate and necessary feature of the digital economy.” As an article in the *New York Times* puts it, the practice of “gathering people’s online data for targeted advertising is not going away” (Chen and Wakabayashi, 2022). Indeed, some cookie replacements may be *more* privacy-invasive (Kemp, 2022). Michael Veale (2022) sums things up neatly, suggesting that as Google, Apple, and others reconstruct adtech, they “might redefine privacy more than they reform profiling.”

This is the continuation of a well-established thrust and parry. Critical scholars raised concerns decades ago about how addressable marketing manifests “rational” discrimination (Gandy and Simmons, 1986), while administrative researchers countered with assurances that the pro-economic boons would outweigh any anti-social drawbacks (Blattberg and Deighton, 1991). When public sentiments tilt toward the critical concerns, marketing and technology companies have a track record of using privacy rhetoric and self-regulations to make commercial surveillance more palatable (Regan, 1995). These maneuvers have tended to manage perceptions more than to confront underlying problems. Privacy policies, for example, were introduced in the 1990s, offering a pretense of transparency and user empowerment; and yet many people have misperceived these policies as assurances that companies will protect privacy, rather than notices of the data collection and processing allowances to which individuals consent (Draper and Turow, 2019). These and other transparency mechanisms secured, rather than challenged, market arrangements rooted in the commodification of personal information (Crain, 2018), and they individualized privacy as a matter of informed consumer choice (Draper, 2017). Consumer sovereignty was illusory, however, since privacy policies harbor a

“transparency paradox” (Nissenbaum, 2011): a fulsome disclosure of commercial surveillance practices would be too complex for reasonable persons to digest, while a human-legible privacy policy cannot explain everything it authorizes.

The recent developments described above have raised the curtain on a pivotal new act in this ongoing drama. The next section details some high-profile solutions that are being marketed to resolve the current contradiction between social and legal demands for privacy, on the one hand, and industry demands for data-driven optimization, on the other.

Industry response to the new privacy imperative

The shift toward “privacy-preserving” solutions impacts the entire adtech sector—from publishers and intermediaries on the “sell” side, who generate revenue through providing access to audiences, to advertisers and intermediaries on the “buy” side, who try to locate high-value users within those audiences. We demonstrate how the industry is reacting to privacy pressures by examining a set of technological interventions across each sequence in the “activation” of data-driven advertising. That sequence includes: (a) identification of consumers; (b) targeting consumers with ads; (c) measurement and attribution of ad campaigns; and (d) optimization of investments in advertising. The “privacy-preserving” solutions under study propose to replicate and maintain existing capacities across this sequence of distinct steps.

Identification lets sellers and buyers recognize users and associate their characteristics and predicted value with unique identifiers. This informs granular targeting and allows advertisers to recycle observed conversions into future targeting criteria. Attribution and optimization processes enable efficiency comparisons across targeting tactics based on key performance indicators, allowing advertisers to further refine their activation efforts. This sequence has depended on third-party cookies, conversion pixels, and mobile ad IDs. Going forward, these activities will rely on abilities to associate user identities with personally identifiable information (PII), to pair first-party data with data from other sources, to create anonymous attribution reports, and to use machine learning to infer insights from platform-enclosed first-party tracking. Given the status companies like Google and Apple hold in the advertising ecosystem, they are positioned to benefit most from these shifts (Kollnig et al., 2022). They (and a few others) are rushing to deepen their stake in adtech’s circuits of value and data, while companies without the same platform power are scrambling to protect themselves from obsolescence.

“Encrypted” identification solutions

Third-party cookie IDs have been a “key prerequisite in personalized/targeted advertising since all user-centric data is

associated with [them]” (IAB Europe, 2020). The deprecation of third-party cookies is leading industry actors to seek “Universal” or “Alternative” IDs that identify users based on PII and first-party cookies and require those who identify consumers to follow certain norms that limit their ability to use collected data. In contrast to current third-party cookie identification, “Universal ID” solutions offer explicit opt-out mechanisms and potentially promise more transparency on the data being collected. The practicality of those alleged improvements, however, remains questionable.

The three primary identification architectures suggested by the industry are as follows: The Trade Desk’s Unified ID 2.0, that is expected to reach 250 million people; LiveRamp’s RampID, that is expected to cover more than 250 million people; and the Secure Web Addressability Network (SWAN), that is expected to include top-20 publishers and advertisers (Asim, 2021).

Trade Desk’s Unified ID (UID) 2.0 aims to let sellers and buyers of online ads match encrypted user IDs in the bid-stream that circulates information about impressions being sold in programmatic auctions. The proposed architecture requires users to login to participating websites via email or other PII. The login email is then encrypted and converted to a “UID2 token” that is refreshed periodically by the system. Those obfuscation mechanisms ensure that the email address cannot be traced back to the user, but do not prevent the matching of IDs to the same user by those who hold access to the raw data. Such a UID 2.0 identifier is likely to provide wide and persistent visibility of the user by the governing bodies (“UID2 Operators”) and participating bodies, like registered Demand-Side Platforms (DSPs) that match user IDs from the seller to the buyer side. Users can “opt-out” of the UID2 framework, but publishers and service providers may limit access to their products for users who decline to log in. Practically, the onus is on publishers to explain why they need an email address and how personalized advertising powers the “free and open web that users appreciate” (Titone, 2021). This “opt out” mechanism will only prevent publishers from creating UID2 tokens and block the ability of DSPs to bid on ads to show to those users.

For the SWAN solution, an individual is identified via a first-party cookie when they first visit a publisher site that has adopted the architecture (Thomson and Rescorla, 2021). Upon loading the publisher page, the user is presented with a pop-up asking for consent to show personalized advertising on the current site and other sites that have adopted the solution. Here, again, the burden is on the publishers to communicate the value exchange to the users. As part of the pop-up, the user also has the option to share their email address, which can serve as an identifier. Regardless of whether they share their email address, a first-party cookie is placed by the initially visited publisher site within the SWAN network that creates a pseudonymous identifier stored on the user’s browser. The

identifier is potentially based on PII, making it more deterministic than a third-party cookie identifier, and enables persistent identification of the user across all participating websites.

The LiveRamp RampID architecture is distinct from the previous two, in that it is interoperable with other ID solutions (Asim, 2021), making the identification instrument even more persistent since it can identify the user across all three ID solutions. The RampID uses user email addresses, upon login to publishers’ sites in exchange for content. Instead of placing a first-party cookie, LiveRamp matches IDs across partners (Asim, 2021). Once a RampID is created, LiveRamp provides cookie syncing across industry actors to match user IDs in DSPs, for example matching to The Trade Desk’s cookies. Once a publisher’s and a marketer’s Ramp IDs are linked, the marketer can buy the user impression. The ability of this solution to further identify users in the ecosystem is also attached to other offline PII (phone number, address history), based on information that advertisers can match with first-, second-, and third-party data.

Each of these solutions creates an identifier that enables advertisers to track users across sites, replicating that capability enabled by third-party cookies. The cookie-less IDs rely on first-party data—users’ PII—passed to a centralized governing actor or through various page redirections (in the case of SWAN), that is then distributed to “registered actors” who can identify the user in the bid stream. Cross-site tracking and profiling is therefore enabled for “authorized partners,” who get users’ consent and conform to the contractual terms of these ID architectures. In that regard, little has changed—advertisers and ad servers are still gathering large amounts of user browsing history data for targeting purposes. Those proposed identification architectures are in fact designs of new tracking systems that come with self-regulatory controls of how the data should be handled and how the real user identity should be masked. Alarming, the solutions mostly use PIIs as identifiers, making user IDs more deterministic and potentially more persistent than third-party cookie IDs. Importantly, however, the new identification architectures do pose one tracking limitation: only those who are registered to use each solution will have cross-site visibility on users. By contrast, third-party cookies let any party involved in the ad auction gain visibility on individuals’ online behavior.

“Tidying” targeting through clean room solutions

There is a growing trend in adtech towards investing in technology that can exploit the first-party data that organizations own. As companies lose the ability to pair that data with other sources via third-party cookies, the digital advertising industry is turning to data clean rooms as a “privacy-preserving” method for maintaining profiling capabilities.

At a high-level, a clean room is a software environment that allows an organization to pair its first-party data with other sources, including other organizations' first-party data and purchased third-party data from data brokers. The name highlights the objective of sanitizing cross-partner data pooling through mutually agreed upon privacy practices (IAB Technology Laboratory, 2023a). A data clean room does not allow participating parties to view joined data sets, preventing reidentification of users or the exposure of PII. Deterministic PII such as email address or a universal ID solution often serves as the identifier that can pair datasets together (Haggin, 2022; Hercher, 2023a).

The primary uses of data clean rooms take place from the buyer side and include: the maintenance of audience insights and segmentation for targeting, campaign measurement and attribution, and audience modeling and activation capabilities that can be used internally within an organization and exported for targeting across activation channels, such as display, search, or social media. These capabilities would allow for continued insight into behavioral and contextual signals for targeting, the creation of look-a-like audiences, and optimization based on user behaviors (IAB and Ipsos, 2023).

While many solutions fall under the label of data clean rooms, the different types of developers often have various goals and particular values to contribute to their customers' digital advertising capabilities. Clean rooms developed by major platforms, such as Google's Ads Data Hub and Amazon's Marketing Cloud, are characterized as walled gardens, which allow platforms to leverage the significant amount of first-party data kept within the platform (Hercher, 2023a). These clean rooms may be bundled with these platforms' other ad and cloud infrastructure offerings, and can be tooled in order to offer direct activation of audience segments and ad campaigns. These provide efficiencies and a powerful incentive for users to select them over alternatives. Platforms have come to define this technology as "privacy compliant" based on control over the flows of first-party data and ability to track user behavior across various locations within its own suite of products (Hercher, 2023b). Alternatively, clean rooms developed independently by parties that are not primary sites of activation allow for advertisers to store and match their data in a centralized location and export it for use elsewhere (Hercher, 2023b). This matching process carried out by non-platform parties is primarily done for the purpose of creating segments for targeting that advertisers can push to activation channels. Independently developed clean rooms highlight that it would be almost impossible to coordinate compliance across all parties involved in creating an interoperable output and have transparency into the various data flowing in and out of the clean rooms (Hercher, 2023a).

Since the practices in this type of privacy-preserving environment derive from different definitions of privacy and limitations of data usage, the employment of a data

clean room does not guarantee compliance with laws or industry norms. Consequently, the IAB has been attempting to standardize elements of clean room services. It has introduced "Open Private Join and Activation" guidelines for data enrichment that occurs through the pairing of different first-party data sets (IAB Technology Laboratory, 2023b). This acknowledges that the solution is being used to match data sets using personally identifiable information, with visibility from the parties involved, to continue granular level targeting. Furthermore, this technology does not allow outside parties to see if the paired data contains information related to sensitive categories that violate privacy norms.

"Anonymous" ad attribution solutions

Attribution refers to the capability of linking ad views or clicks—called "source events"—to subsequent consumer behaviors, like a purchase or app download—called "target events." Third-party cookies and tracking pixels currently enable this through persistent monitoring of identifiable users. These linkages let advertisers measure the effectiveness of ad campaigns and optimize accordingly to maximize ROI. In preparation for the post-third-party tracking era, Google, Meta, and Apple have publicized new attribution solutions that, according to company documents, "improve" user privacy by inserting sophisticated technologies into the mediation of these data flows (Apple, 2021; Nalpas and White, 2021; Taubeneck et al., 2022).

All three solutions promise to mask user identity during the generation and distribution of aggregated reports that link source and target events. Their designs are complex and some implementation details are vague. But, in general, they aim to reduce the ability of unauthorized actors to identify, track, and profile individual users, while still joining the records of events that are generated by a user's behaviors. They accomplish this via a suite of technical mechanisms. The Interoperable Private Attribution (IPA) system proposed by Meta and Mozilla, for instance, uses unique identifiers called "write-only match keys," which are stored on and only readable by the local device, to record source and target events from the same user; the IPA system then generates reports by way of a multi-party computation method that disperses personal information and encryption keys across trusted servers, so that useful insights can be gleaned without allowing a single party to access complete records of user behavior; those reports aggregate data and use differential privacy techniques to obfuscate individuals within a crowd; finally, privacy budgets, which restrict actors' abilities to query the servers, prevent adversaries from deanonymizing attribution data with brute force (Taubeneck et al., 2022). Google and Apple use some but not all of these same techniques, and their approaches lean heavily on special permissions granted to Chrome and Safari browsers as well as Apple iOS. For example, Apple's SKAdNetwork, which lets registered ad

networks measure the impact of advertising on app downloads and engagement, is considered private by design and thus exempted from Apple's ATT mechanism (Apple, 2023). They also add delays to the production of attribution reports to further obfuscate users' identities. Essentially, these companies claim that as long as the identifiers assigned to users are not accessible to third parties, and as long as the anonymous information associated with individuals is stored locally (on a device or browser) or cryptographically protected, then the data collection and analysis required to provide attribution services is privacy-compliant.

These solutions aim to maintain accurate reporting on the performance of ad campaigns while promising to conceal user identity and limit the data shared in attribution reports. Observing and influencing users, however, is still the priority; the goal remains to induce higher value consumer behaviors at lower advertising costs. While they improve certain information-security functions, like preventing data leakage, these "private" attribution solutions continue to service corporate demands for online surveillance. They use clever technological designs both to guard against adversarial abuses, and to protect the premise that advertisers, intermediaries, and platforms are all entitled to determine how effectively digital ads are performing.

"Privacy-safe" optimization

The last set of solutions examined here takes us into the black-box magic of optimization. Meta and Google are both steering advertisers toward adtech products that use machine learning to identify valuable consumer targets and optimize marketing performance, exercising near-total algorithmic discretion about where to serve ads. These products, branded respectively as Meta's Advantage+¹ and Google's Performance Max, are not strictly "privacy-preserving," but they are definite responses to the challenges and opportunities encountered in adtech's pivot to privacy. "PMax isn't a privacy product," a trade reporter explains. "But it's part of a future digital ad model without third-party tracking" (Hercher, 2022a). Meta, on the other hand, explicitly markets Advantage+ functions as "privacy-safe" means for scaling the delivery of personalized advertisements (Meta, 2022). In particular, Meta is trying to work around the difficulties posed by Apple's ATT (Murphy and Criddle, 2023). "Having read the room," as one observer puts it, "Meta is investing in AI to develop and deploy privacy-enhancing technologies (PETs) to underpin its ad platform" (Schiff, 2023).

Meta and Google position these products in the familiar registers of automation and optimization. In Advantage+ and Performance Max, clients specify conversion targets, the price they want to pay for those conversions, and budget caps, and these systems work to maximize each

campaign's return on advertising spend. These systems are designed to compensate for the new limitations on third-party measurement and data sharing. Take Meta for example. Rather than relying just on "its conversion pixel and [software development kit] to connect impressions to conversions," Meta ingests advertisers' sales records by integrating the latter's databases into its servers (Hercher, 2022b). This shifts the processing of conversion records from a third-party relationship, wherein Meta observes events on advertisers' websites or apps, to something that looks like a second-party relationship, where advertisers share the first-party sales data they collect. While ATT stops Meta from monitoring the cross-app chain of events from ad exposure to purchase behavior, "Apple can't block a company from piping its CRM, CDP or data warehouse to Meta" (Hercher, 2022b). And where conversion data are missing, these optimization systems fill in the gaps with probabilistic modeling. By shifting responsibilities for measurement and data ownership among parties, and because Meta and Google do not report granular details to advertisers about how these black boxes target or influence individual users, this whole arrangement has a veneer of privacy.

These solutions have garnered attention for being opaque but effective (Murphy and Criddle, 2023). They also demonstrate the broader transition away from situating "ground truth" in advertising around widely accessible behavioral signals and toward the use of machine-learning models that rely on massive first-party data assets and computing resources. Meta, for instance, has reportedly "invested in dramatically expanding its computing power in order to train these more complex AI models on larger data sets" (Murphy and Criddle, 2023). Large platform companies are, in these cases, promising to replicate the optimization capabilities that marketers are used to by absorbing more adtech and data-analytic functionalities and by leaning more on inferential profiling and evaluation. Through these products, Meta and Google can tell a compelling story about why advertisers should continue to see their platforms as the best options for maximizing ROI; and they can, at an operational level, exert additional control over the distribution of ad revenue, perhaps favoring their own properties. Google and Meta may steer customers toward these products by making them prerequisites for accessing other spending channels. For example, advertisers who want to promote themselves on Google Maps now must use Performance Max (Hercher, 2022a). This not only pumps revenue through these pipes, but it also builds the training data necessary for optimization. And since Meta and Google only tell ad clients how much money was spent and how much revenue was likely produced, all the information about what placements worked and what users converted stays with the platforms. This could drive adoption through a network effect: as advertisers feed data into these systems, and as the models learn

to recognize valuable prospects for certain products, those models become more attractive to other advertisers targeting the same people. In sum, platform companies are using these products to exploit and intensify information asymmetries, while presenting them as solutions that will comply with restrictions on third-party tracking.

Few details about the inner workings of these products are publicly known. What is clear, however, is that they represent a waypoint for adtech's pivot to privacy, wherein optimization is powered by inferential machine learning, third-party tracking is replaced with first-party tracking and data-sharing partnerships, and platforms internalize more computing and logistical operations within walled gardens, increasing dependency on those core adtech companies. These maneuvers are emblematic of an approach to "privacy" that we call cynical resignation.

Cynical resignation

Analysis of industry responses to the new privacy imperative demonstrates how the recent wave of privacy regulations—including the GDPR, CPRA, and increased momentum behind the passage of a US federal privacy bill—have left tech firms cynically resigned to a world in which they will have to comply with new privacy mandates, even as they continue to lobby against them and resist enforcement. This is a reversal of conventional understandings of privacy resignation in that it is *companies* that are now resigned to the persistence of privacy regimes, rather than consumers being resigned to routine privacy invasions via commercial surveillance (Draper and Turow, 2019; Lutz et al., 2020). Corporate resignation to privacy mandates is, unlike that of many consumers, cynical rather than apathetic: while they "perform" privacy, companies are simultaneously attempting to get ahead of these mandates in order to ensure that core elements of their business model will remain untouched (Hagey et al., 2022).

Under current policy approaches, this appears to be a largely successful maneuver: policy proposals concentrate on third-party data tracking and data security, offering significant leeway to firms to both develop technical interventions that enable them to sidestep meaningful privacy measures and pursue market strategies that go even further, enabling a few of the biggest firms to profit and consolidate power. In this section we characterize these maneuvers in three forms, mapping the industry playbook for cynical privacy resignation.

Sanitizing surveillance

One of the means through which firms evade privacy mandates is through the development of computational techniques that enable them to obfuscate data flows and sanitize

profiling activities. While these technological systems may be technically compliant with privacy regimes, they do not make any meaningful underlying change to the practices that animated privacy concerns in the first place—particularly concerns that are sensitive to issues of power. For example, significant investments in PETs enable firms to comply with requirements not to retain data that reflects individual-level tracking, while still allowing them to profit from the underlying insights that can be derived from this data. Such methods include the use of multi-party computation, which enables multiple parties to match and query data from multiple business partners without revealing individual-level information, and the use of "clean rooms" that enable ad targeting to take place within a secure environment operated by a third party, without the targeting firm gaining direct access to individual insights.

These create growth opportunities for firms that maintain a cloud infrastructure business, because they are able to sell features such as clean room services as add-on compliance features for their potential clients. This enables their clients to do such things as retain anonymized audience sets within a clean room over time for persistent tracking, and combine data from across many subsidiary platforms (Hercher, 2023c). What they do not do is meaningfully address the role of data as a key source of tech firms' power. Indeed, some of the regulatory moves happening around PETs are being used to avert more bright line data protection measures (Veale, 2022).

Another example is the turn to inferences as the key locus of insights (Solow-Niederman, 2021). Google in 2021 announced plans to shift to using "data-driven attribution," which substitutes industry standard "last-click" attribution methods with algorithmic models designed to infer the connection between ad impressions and user activity in the absence of the data that cookies or other identifiers provide (Srinivasan, 2020). As one trade publication describes it, "Data-driven attribution may not be more privacy-compliant than last-click [attribution], in and of itself. But in a privacy-forward environment where connecting ad impressions to online user activity is often prohibited, it will be the reliable methodology" (Hercher, 2021). Though this data is probabilistic, and does not contain personal identifiable information, it nevertheless retains the information asymmetries that characterize other models in behavioral advertising, in many ways concentrating these asymmetries even further by positioning Google's algorithmic models as the ground truth on which publishers and advertisers can evaluate the effectiveness of their placements.

Party-hopping

Maximizing information asymmetries is at the heart of the second technique that characterizes cynical privacy resignation, what we're calling "party-hopping": pursuing

more invasive first-party data practices in lieu of disreputable third-party practices. It is a policy dodge enabled by failures to enforce a meaningful definition of privacy. Party-hopping works on the premise that data collection and processing are unproblematic so long as the number of companies accessing or owning that data can be restricted. It both dismisses data minimization and encourages consolidation. While many companies are able to engage in first-party tracking, big tech firms—which can draw on their network effects and combine streams of data drawn from subsidiaries within the platform ecosystems they control—benefit most from this transition, and several emerging industry techniques illustrate how they intend to maximize this data advantage.

One example of party-hopping is through the development of new technologies that enhance companies' first-party tracking capabilities. The clean rooms described above offer one illustration of what this looks like in practice. While many clean rooms are designed for interoperability, those offered by Google, Amazon, and Meta function as walled gardens, as one recent industry article bemoaned: "The privacy magic isn't in encryption or even the degree to which unbeknownst users are tracked online, it's whether the data is held in the hands of one company" (Hercher, 2023b). By drawing clients into the use of their clean room technology and integrating it with their other service offerings, cloud providers are able to keep more users operating within their walled gardens (Hercher, 2023b). Meta's development of its Meta Pixel to replace its standard conversion pixel is an additional example. The former triggers the website hosting it to set a first-party cookie on the user's browser and then share data with Meta. This acts as a workaround to get past Apple's intelligent tracking prevention, which blocks third-party cookies, enabling Meta to replicate the data flows established prior to Apple's "privacy enhancing" measures (Zawadzinski and Wlosik, 2022).

Universal ID solutions work along a similar principle. By cloaking persistent tracking in the legitimating cover of a first-party consent agreement, unique identifiers and profile data can once again be activated across sites by various advertisers and intermediaries who signed up for the same global ID solutions. SWAN, for example, effectively turns every publisher within the network into a "first party," making users eligible to be targeted by advertisers across sites of other SWAN partners, despite encryption of user-level information by site. As discussed further below, Apple's ATT also hinges on the pretense that, by definition, first-party tracking is legitimate and third-party tracking is not.

Across these examples are a set of techniques that large tech firms are best placed to exploit in environments where third-party tracking is eroded. By controlling bottlenecks, such as browsers, devices, cloud computing, and marketplaces, they have the capacity to technologically innovate

around regulatory environments to preserve their data advantage. And where this fails, they have the capital to buy their way into data advantages.

Sabotage

A third model of cynical resignation takes the form of what we describe as sabotage—using "privacy" as a means of disadvantaging rivals and increasing market power. This strategy not only seeks to retain core elements of data tracking in advertising, it reinforces platforms' dominance in the market by directly undermining external competitors' abilities to gain insights from third-party data. Our terminology here draws on Shapiro's (2023) Veblen-influenced definition of platform sabotage as the strategic withdrawal of efficiency from other market actors. This is quite applicable to digital advertising, whose quintessential value proposition is marketing efficiency, and wherein access to personal data and addressable targeting capabilities are considered the fundamental sources of that efficiency (Turow, 2011). We call this sabotage because an offending company does not improve its products so much as withhold perceived efficiencies from the rest of the marketplace.

Apple's announcement of its App Tracking Transparency privacy tools is a case in point. ATT mandates that when apps seek to collect user information to share it with third parties, a window must show up on the Apple device asking for permission, and if the user declines the app must stop tracking the user's data. The announcement of this measure led Meta to take out full-page newspaper ads denouncing the feature as harmful to small businesses. It also led Germany's competition regulator to initiate a proceeding against Apple to review ATT's impact on other companies, specifically examining whether Apple's ability to unilaterally set rules for its app store had anticompetitive effects on the market.

This maneuver should be viewed in light of Apple's growing advertising business. As mentioned above, ATT operationalizes privacy in a self-serving way; its definition of "tracking" prohibits third-party data collection (absent opt-in consent), while Apple's own first-party data collection and personalized advertising services are exempted from such prohibitions and permissions requirements (Kollnig et al., 2022). That means advertisers can get more detailed data about the performance of their ads if they buy from Apple rather than other ad-sellers, and that apps financed by personalized advertising may become more dependent on Apple. As one analyst told the *Wall Street Journal*, "it's not really clear why they would withhold [performance data]...unless they were just really trying to privilege their own ad networks" (Haggin, 2021). This party-based privacy policing generates a windfall for the company. In the first six months after the launch of ATT, Apple's advertising business reportedly tripled its market share (McGee, 2021).

Similarly, Google's intention to block third-party cookies on Chrome and its subsequent implementation of Privacy Sandbox foreshadow significant downstream effects for other companies (Wakabayashi, 2021). This led to complaints by publishers and the UK Competition and Markets Authority, both of whom worried that Privacy Sandbox would have negative impacts on competition in the digital ads market. The result was a negotiated set of concessions made by Google in exchange for the CMA's agreement not to pursue an investigation. The agreement is heavily reliant on a monitoring trustee to track and evaluate whether Google has adequately complied with these commitments (Competition and Markets Authority, 2022). This is not the first time Google has used "privacy" to justify terms of service and data sharing (or withholding) that make its own adtech products more valuable to advertisers and publishers, relative to competitors (Srinivasan, 2020).

FTC Commissioner Rebecca Kelly Slaughter (2021: 21) remarked on these sorts of practices in a recent speech: "Shutting off the data spigot for others while filling your own well is the kind of anticompetitive innovation that we're bound to see more of if this space remains unregulated." Where privacy regulations fail to take into account these kinds of industry practices, they are likely to result in further concentration of data advantages in the hands of large tech firms who are positioned to benefit from platform sabotage.

Conclusion: the fight to define meaningful change

The digital advertising industry, and the tech sector more broadly, is forging ahead with a paradigm that may define and encode the meaning and function of privacy for the foreseeable future. Now that they've accepted a world in which privacy remains salient, their actions indicate firms are seeking to bring privacy within their own control: parading behind a welcome rhetoric about privacy-preserving technology, while actively working to reproduce their existing profit centers. The largest among these firms are exploiting privacy through self-regulatory technical measures that cement their power and make small firms more dependent on them. In such a system, only the largest firms have the data, scale/reach, and computing capacities to solidify the status quo while simultaneously performing "compliance" with laws and norms. Consider, for example, the black-box optimization products that Meta and Google are positioning as catch-all adtech solutions for a "privacy-first" future. These combine all of the tactics described above: they enrich the platforms' own first-party data with first-party sales data that advertising clients integrate into ad-targeting databases; they use machine-learning models to sanitize surveillance and discrimination; and they let these companies claim further control over the allocation of advertisers' money and hoard data and computing resources.

Adtech's apparent pivot thus takes advantage of a policy gap: in the absence of privacy mandates that are sensitive to political-economic and sociotechnical power, these firms can perform privacy without making fundamental changes to the business model. Each of the maneuvers described above—sanitizing surveillance, party-hopping, and sabotage—work in concert to solidify adtech companies' ability both to retain the status quo and to avert regulatory interventions that would invoke structural changes to the digital economy in favor of protecting users' privacy. Given this, cynical resignation may be interpreted as the product of a policy failure: a failure to introduce strong curbs on corporate data collection and profiling, and a failure to anticipate tensions between privacy and competition policy that has enabled these companies to take steps to ensure the broader public gets neither.

To make privacy meaningful, and forestall these self-interested moves by companies, policy measures have to tackle the structural harms at the root. For example, data minimization, collection limitation, and purpose limitation measures can be one effective means through which to cut off firms' access to data (Center for Democracy & Technology, 2022). Other regulatory proposals could ban certain types of business models wholesale, or institute strong curbs on profiling activities (US Congress, 2022). And stronger competition enforcement, particularly enforcement that treats privacy harms as an example of toxic competition (Stucke, 2022), can avert firms' use of acquisitions to both build a stronger foothold in the advertising ecosystem and expand their access to data flows, ensuring that we don't face a race to the bottom as firms seek to leverage privacy rhetoric to their own gain. The industry's newfound commitment to privacy opens a political horizon to define what it will mean to honor that commitment.

Acknowledgements

The authors would like to thank the ComplianceNet community for their great feedback during the 2023 annual conference at American University.


Declaration of conflicting interests


The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Lee McGuigan  <https://orcid.org/0000-0003-3157-2944>

Ido Sivan-Sevilla  <https://orcid.org/0000-0003-2194-5006>

Note

1. Advantage+ is branded onto a suite of automated optimization products; for simplicity, we use the label as a general umbrella.

References

- Ahuja K, Bauer T, Meder C, et al. (2022) As the Cookie Crumbles, Three Strategies for Advertisers to Thrive. *McKinsey & Company* (blog), April 6. Available at: <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/as-the-cookie-crumbles-three-strategies-for-advertisers-to-thrive> (accessed 9 January 2023).
- Alaimo C (2022) From people to objects: The digital transformation of fields. *Organization Studies* 43(7): 1091–1114.
- Asim A (2021) A Comprehensive Guide to Third-Party Cookie Alternatives. *Digiday*. November 4, <https://digiday.com/media/a-digiday-media-guide-to-third-party-cookie-alternatives/>.
- Apple (2021) Meet Privacy-Preserving Ad Attribution. *Apple Developer*. Available at: <https://developer.apple.com/videos/play/wwdc2021/10033/> (accessed 2 February 2023).
- Apple (2022) What's New with SKAdNetwork. *Apple Developer*. Available at: <https://developer.apple.com/videos/play/wwdc2022/10038/> (accessed 3 February 2023).
- Apple (2023) Attributing Ads with SKAdNetwork and Private Click Measurement. *Apple Developer*. Available at: <https://developer.apple.com/app-store/ad-attribution/> (accessed 2 February 2023).
- Baik JS (2020) Data privacy against innovation or against discrimination?: The case of the California Consumer Privacy Act (CCPA). *Telematics and Informatics* 52: 101431.
- Bindra C (2021) Building a Privacy-First Future for Web Advertising. *Google Ads* (blog), January 25. Available at: <https://blog.google/products/ads-commerce/2021-01-privacy-sandbox/> (accessed 2 February 2023).
- Birch K (2020) Automated neoliberalism? The digital organisation of markets in technoscientific capitalism. *New Formations* 100: 10–27.
- Blattberg RC and Deighton J (1991) Interactive marketing: Exploiting the age of addressability. *Sloan Management Review* 33(1): 5–14.
- Bobrowsky M (2022) Facebook Feels \$10 Billion Sting from Apple's Privacy Push. *Wall Street Journal*, February 3. Available at: <https://www.wsj.com/articles/facebook-feels-10-billion-sting-from-apples-privacy-push-11643898139>.
- Braun J (2019) The Devil in the Details: User Tracking is Hurting More than our Privacy, It's Doing Serious Damage to Public-Interest Media, Too. *Flow*, February 22. Available at: <https://www.flowjournal.org/2019/02/the-devil-in-the-details/>.
- Brodherson M, Broitman A, Macdonald C, et al. (2021) The Demise of Third-Party Cookies and Identifiers. *McKinsey & Company* (blog), April 12. Available at: <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-demise-of-third-party-cookies-and-identifiers> (accessed 10 January 2023).
- Burrell J and Fourcade M (2021) The society of algorithms. *Annual Review of Sociology* 47: 213–237.
- Center for Democracy & Technology (2022) Comments on Commercial Surveillance ANPR. <https://cdt.org/wp-content/uploads/2022/11/CDT-Comments-to-FTC-on-ANPR-R111004.pdf>.
- Chen B and Wakabayashi D (2022) You're Still Being Tracked on the Internet, Just in a Different Way. *New York Times*, April 6. Available at: <https://www.nytimes.com/2022/04/06/technology/online-tracking-privacy.html>.
- CMA (2022) Investigation into Google's "Privacy Sandbox." <https://www.gov.uk/cma-cases/investigation-into-googles-privacy-sandbox-browser-changes>.
- Cohen JE (2013) What privacy is for. *Harvard Law Review* 126: 1904–1933.
- Corporate Europe Observatory (2022) How corporate lobbying undermined the EU's push to ban surveillance ads. Available at: <https://corporateeurope.org/en/2022/01/how-corporate-lobbying-undermined-eus-push-ban-surveillance-ads> (accessed January 18).
- Crain M (2018) The limits of transparency: Data brokers and commodification. *New Media & Society* 20(1): 88–104.
- Crain M (2019) A critical political economy of web advertising history. In: Brügger N and Milligan I (eds) *The SAGE Handbook of Web History*. London: Sage, 330–343.
- Crain M (2021) Power Play: Big Tech's Feud Over Mobile App Tracking. *The Reboot*, June 10. Available at: <https://web.archive.org/web/20210610172629/https://thereboot.com/power-play-big-techs-feud-over-mobile-app-tracking/>.
- Crain M and Cohen NS (2023) Social media and audience commodification: Toward an applied theory. In: West E and McAllister M (eds) *The Routledge Companion to Advertising and Promotional Culture*, 2nd edition. New York: Routledge, 115–125.
- Darmody A and Zwick D (2020) Manipulate to empower: Hyper-relevance and the contradictions of marketing in the age of surveillance capitalism. *Big Data & Society* 7(1): 1–12. DOI: 10.1177/2053951720904112.
- Deighton J and Kornfeld L (2020) *The Socioeconomic Impact of Internet Tracking*. New York: Interactive Advertising Bureau.
- Draper NA (2017) From privacy pragmatist to privacy resigned: Challenging narratives of rational choice in digital privacy debates. *Policy & Internet* 9(2): 232–251.
- Draper NA and Turow J (2019) The corporate cultivation of digital resignation. *New Media & Society* 21(8): 1824–1839.
- Fourcade M and Healy K (2017) Seeing like a market. *Socioeconomic Review* 15(1): 9–29.
- Gandy OH (2021) *The Panoptic Sort: A Political Economy of Personal Information*, 2nd edition. New York: Oxford University Press.
- Gandy OH and Simmons CE (1986) Technology, privacy and the democratic process. *Critical Studies in Media Communication* 3(2): 155–168.
- Graham M (2022) More Changes Loom for Online Marketers. *Wall Street Journal*, January 25. Available at: <https://www.wsj.com/articles/more-changes-loom-for-online-marketers-11643150679>.
- Hagey K, Schechner S and Kupa M (2022) Why Google Plays Down Its Ad-Tech Business but Is Determined to Keep It. *Wall Street Journal*, November 1. Available at: <https://www.wsj.com/articles/why-google-plays-down-its-ad-tech-business-but-is-determined-to-keep-it-11667292084>.
- Haggin P (2021) Apple's Privacy Changes Are Poised to Boost Its Ad Products. *Wall Street Journal*, June 9. Available at: <https://web.archive.org/web/20210609195842/https://www.wsj.com/>

- articles/apples-privacy-changes-are-poised-to-boost-its-ad-products-11619485863.
- Haggin P (2022) Advertisers Turn to ‘Clean Rooms’ to Keep Consumer Data Private. *Wall Street Journal*, October 19. Available at: <https://www.wsj.com/articles/advertisers-data-clean-rooms-11666038545>.
- Hercher J (2021) Goodbye, Last Click Attribution. *AdExchanger*, September 27. Available at: <https://www.adexchanger.com/online-advertising/goodbye-last-click-attribution-google-ads-changes-default-to-data-modeling/>.
- Hercher J (2023b) Acxiom Takes Its Data Broker Biz to the Cloud with Snowflake’s Help. *AdExchanger*, February 7. Available at: <https://www.adexchanger.com/ad-exchange-news/acxiom-takes-its-data-broker-biz-to-the-cloud-with-snowflakes-help/>.
- Hercher J (2022a) Meet Performance Max, the Blackest Black Box of All Google Ad Products. *AdExchanger*, December 14. Available at: <https://www.adexchanger.com/commerce/meet-performance-max-the-blackest-black-box-of-all-google-ad-products/>.
- Hercher J (2022b) More Performance, Less Transparency: Inside Meta’s Advantage+ Shopping Black Box. *AdExchanger*, December 19. Available at: <https://www.adexchanger.com/commerce/more-performance-less-transparency-inside-metas-advantage-shopping-black-box/>.
- Hercher J (2023a) Why 2023 Is a Pivotal Year for Indie Data Clean Rooms. *AdExchanger*, January 18. Available at: <https://www.adexchanger.com/data-exchanges/why-2023-is-a-pivotal-year-for-indie-data-clean-rooms/>.
- Hercher J (2023c) Will the Walled Gardens Turn Clean Room Tech into Yet Another Platform Plaything? *AdExchanger*, January 25. Available at: <https://www.adexchanger.com/platforms/will-the-walled-gardens-turn-clean-room-tech-into-yet-another-platform-plaything/>.
- Hoffmann CP, Lutz C and Ranzini G (2016) Privacy cynicism: A new approach to the privacy paradox. *Cyberpsychology* 10(4). DOI: 10.5817/CP2016-4-7.
- IAB (2022) IAB slams bill that would eliminate data-driven advertising. <https://www.iab.com/news/iab-slams-bill-that-would-eliminate-data-driven-advertising/>.
- IAB and Ipsos (2023) *State of Data 2023: Data Clean Rooms & the Democratization of Data in the Privacy-Centric Ecosystem*. https://www.iab.com/wp-content/uploads/2023/01/IAB_State_of_Data_2023.pdf.
- IAB Europe (2020) A guide to the post third-party cookie era. <https://iab europe.eu/knowledge-hub/iab-europe-guide-to-the-post-third-party-cookie-era/>.
- IAB Tech Lab (2022) Time to build for privacy. <https://iabtechlab.com/blog/time-to-build-for-privacy/>.
- IAB Tech Lab (2023a) *Data Clean Rooms: Guidance and Recommended Practices*. <https://iabtechlab.com/wp-content/uploads/2023/02/FINAL-DRAFT-PUBLIC-COMMENT-Data-Clean-Room-Guidance-IAB-Tech-Lab.pdf>.
- IAB Tech Lab (2023b) *Open Private Join and Activation: A Data Clean Room Interoperability Specification*. <https://iabtechlab.com/wp-content/uploads/2023/02/FINAL-DRAFT-PUBLIC-COMMENT-Open-Private-Join-Activation-IAB-Tech-Lab.pdf>.
- Kant T (2021) Identity, advertising, and algorithmic targeting: Or how (not) to target your “ideal user”. *MIT Case Studies in Social and Ethical Responsibilities of Computing* (Summer 2021). DOI: 10.21428/2c646de5.929a7db6.
- Keller P (2022) *After Third Party Tracking: Regulating the Harms of Behavioural Advertising Through Data Protection*. King’s College London.
- Kemp K (2022) “A rose by any other unique identifier”: Regulating consumer data tracking and anonymisation claims. *Competition Policy International TechReg Chronicle* (August 2022): 21–29.
- Kollnig K, Shuba A, Van Kleek M, et al. (2022) Goodbye tracking? Impact of iOS app tracking transparency and privacy labels. *FACCT ’22: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (June 2022): 508–520. DOI: 10.1145/3531146.3533116.
- Lawler R (2022) Google Delays Blocking Third-Party Cookies Again, Now Targeting Late 2024. *The Verge*, July 27. Available at: <https://www.theverge.com/2022/7/27/23280905/google-chrome-cookies-privacy-sandbox-advertising>.
- Lomas N (2022) On Meta’s ‘Regulatory Headwinds and Adtech’s Privacy Reckoning. *TechCrunch*, February 4. Available at: <https://techcrunch.com/2022/02/04/on-metas-regulatory-headwinds-and-adtechs-privacy-reckoning/>.
- Lomas N (2023) Use of Meta Tracking Tools Found to Breach EU Rules on Data Transfers. *TechCrunch*, March 16. Available at: <https://techcrunch.com/2023/03/16/meta-tracking-gdpr-data-transfer-breach/>.
- Lutz C, Hoffmann CP and Ranzini G (2020) Data capitalism and the user: An exploration of privacy cynicism in Germany. *New Media & Society* 22(7): 1168–1187.
- Marwick A (2022) Privacy without power: What privacy research can learn from surveillance studies. *Surveillance & Society* 20(4): 397–405.
- McGee P (2021) Apple’s Privacy Changes Create Windfall for Its Own Advertising Business. *Financial Times*, October 17.
- McGuigan L (2023) *Selling the American People: Advertising, Optimization, and the Origins of Adtech*. Cambridge, MA: MIT Press.
- McNealy J (2022) Platforms as phish farms: Deceptive social engineering at scale. *New Media & Society* 24(7): 1677–1694.
- Mellet K and Beauvisage T (2020) Cookie monsters. Anatomy of a digital market infrastructure. *Consumption Markets & Culture* 23(3): 110–129.
- Mizarhi-Borohovich I, Newman A and Sivan-Sevilla I (2023) The civic transformation of data privacy implementation in Europe. *West European Politics*: 1–30. Published online ahead of print 15 March. DOI: 10.1080/01402382.2023.2184108.
- Mudd G (2021) Privacy-Enhancing Technologies and Building for the Future. *Meta for Business* (blog), August 11. Available at: <https://www.facebook.com/business/news/building-for-the-future> (accessed 1 March 2023).
- Murphy H and Criddle C (2023) *Meta’s AI-Driven Advertising System Splits Marketers*. *Financial Times*, February 27.
- Meta (2022) Bring in the Holidays with Advantage+ Shopping Campaigns. *Meta for Business* (blog), October 26. Available at: <https://www.facebook.com/business/news/advantage-plus-holiday-shopping-campaigns>.
- Nalpas M and White A (2021) Attribution Reporting. *Chrome Developers* (blog), May 18. Available at: <https://developer.chrome.com/docs/privacy-sandbox/attribution-reporting/> (accessed 16 February 2023).

- Nieborg DB and Poell T (2018) The platformization of cultural production: Theorizing the contingent cultural commodity. *New Media & Society* 20(11): 4275–4292.
- Nissenbaum H (2009) *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford Law Books.
- Nissenbaum H (2011) A contextual approach to privacy online. *Daedalus* 140(4): 32–48.
- Norburn C (2022) Google's Ads Privacy Lead on Why It's Fighting to Save the Ad-Funded Internet. *The Drum*, September 23. Available at: <https://www.thedrum.com/opinion/2022/09/23/google-s-ads-privacy-lead-why-it-s-fighting-save-the-ad-funded-internet>.
- Plantin JC, Lagoze C, Edwards PN, et al. (2018) Infrastructure studies meet platform studies in the age of Google and Facebook. *New Media & Society* 20(1): 293–310.
- Regan PM (1995) *Legislating Privacy: Technology, Social Values, and Public Policy*. Chapel Hill, NC: University of North Carolina Press.
- Satariano A (2023) Meta Fined \$414 million After Ad Practices Ruled Illegal Under EU Law. *New York Times*, September 23. Available at: <https://www.nytimes.com/2023/01/04/technology/meta-facebook-eu-gdpr.html>.
- Schiff A (2023) Meta is Betting on Three Acronyms: AI, ML and PETs. *AdExchanger*, February 2. Available at: <https://www.adexchanger.com/platforms/meta-is-betting-on-three-acronyms-ai-ml-and-pets/>.
- Shapiro A (2023) Platform sabotage. *Journal of Cultural Economy* 16(2): 203–220. DOI: 10.1080/17530350.2022.2159495.
- Sivan-Sevilla I (2022) Varieties of enforcement strategies post-GDPR: A fuzzy-set qualitative comparative analysis (fsQCA) across data protection authorities. *Journal of European Public Policy*: 1–35. Published online ahead of print 19 November. DOI: 10.1080/13501763.2022.2147578.
- Slaughter RK (2021) *Wait but Why? Rethinking Assumptions About Surveillance Advertising*. Federal Trade Commission. Available at: https://www.ftc.gov/system/files/documents/public_statements/1597998/iapp_psr_2021_102221_final2.pdf (accessed October 22).
- Solow-Niederman A (2021) Information privacy and the inference economy. *Northwestern University Law Review* 117(2): 357–424.
- Srinivasan D (2020) Why Google dominates advertising markets. *Stanford Technology Law Review* 24: 56–175.
- Stucke ME (2022) The relationship between privacy and antitrust. *Notre Dame Law Review Reflection* 97(5): 400–416.
- Swant M (2023) From the FTC to SCOTUS, the Ad Tech World Has Its Hands Full of Privacy and Policy Issues. *Digiday*, February 20. Available at: <https://digiday.com/media-buying/from-the-ftc-to-scotus-the-ad-tech-world-has-its-hands-full-of-privacy-and-policy-issues/>.
- Taubeneck E, Savage B and Thomson M (2022) Interoperable Private Attribution (IPA) Overview. <https://docs.google.com/document/d/1KpdSKD8-Rn0bWPTu4UtK54ks0yv2j22pA5SrAD9av4s/edit>.
- Thomson M and Rescorla E (2021) Comments on SWAN and Unified ID 2.0. https://mozilla.github.io/ppa-docs/swan_uid2_report.pdf.
- Titone T (2021) Unified ID 2.0 Explained. *Ad Tech Explained*, May 3. Available at: <https://adtechexplained.com/unified-id-2-0-explained/>.
- Turow J (2011) *The Daily You*. New Haven: Yale University Press.
- U.S. Congress (2022) *Banning Surveillance Advertising Act*. H.R.6416 <https://www.congress.gov/bill/117th-congress/house-bill/6416>.
- van der Vlist FN and Helmond A (2021) How partners mediate platform power: Mapping business and data partnerships in the social media ecosystem. *Big Data & Society* 8(1): 1–16. DOI: 10.1177/20539517211025061.
- Veale M (2022) Future of Online Advertising: Adtech's New Clothes Might Redefine Privacy More than They Reform Profiling. *netzpolitik*, February 25. Available at: <https://netzpolitik.org/2022/future-of-online-advertising-adtechs-new-clothes-might-redefine-privacy-more-than-they-reform-profiling-cookies-meta-mozilla-apple-google/>.
- Veale M and Borgesius Z (2022) Adtech and real-time bidding under European data protection law. *German Law Journal* 23(2): 226–256.
- Viljoen S, Goldenfein J and McGuigan L (2021) Design choices: Mechanism design and platform capitalism. *Big Data & Society* 8(2): 1–13. DOI: 10.1177/20539517211034312.
- Wachter S (2020) Affinity profiling and discrimination by association in online behavioral advertising. *Berkeley Technology Law Journal* 35: 367–430.
- Wakabayashi D (2021) Google Delays a Privacy Change to Its Chrome Web Browser. *New York Times*, August 10. Available at: <https://www.nytimes.com/2021/06/24/business/google-privacy-chrome.html>.
- West SM (2019) Data capitalism: Redefining the logics of surveillance and privacy. *Business & Society* 58: 20–41.
- Whitman JQ (2004) The two Western cultures of privacy: Dignity versus liberty. *Yale Law Journal* 113(6): 1151–1221.
- Winslow G (2022) IAB: Proposed FTC Rules Could 'Criminalize the Internet'. *tvtech*, November 15. Available at: <https://www.tvtechnology.com/news/iab-proposed-ftc-rules-could-criminalize-the-internet>.
- Wodinsky S (2021a) Google's Quest to Kill the Cookie Is Creating a Privacy Shitshow. *Gizmodo*, June 11. Available at: <https://gizmodo.com/googles-quest-to-kill-the-cookie-is-creating-a-privacy-1847072958>.
- Wodinsky S (2021b) Google's Plan to Quash Cookies Draws Scrutiny from Regulators. *Gizmodo*, January 8. Available at: <https://gizmodo.com/googles-plan-to-quash-cookies-draws-scrutiny-from-regul-1846018255>.
- Wu AX and Taneja H (2021) Platform enclosure of human behavior and its measurement: Using behavioral trace data against platform episteme. *New Media & Society* 23(9): 2650–2667.
- Zawadzinski M and Wlosik M (2022) What Facebook's First-Party Cookie Means for AdTech. *Clearcode*, June 8. Available at: <https://clearcode.cc/blog/facebook-first-party-cookie-adtech/>.