



Article

Private attributes: The meanings and mechanisms of “privacy-preserving” adtech

new media & society

1–22

© The Author(s) 2023



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/14614448231213267

journals.sagepub.com/home/nms



Lee McGuigan 

The University of North Carolina at Chapel Hill, USA

Ido Sivan-Sevilla

University of Maryland, USA

Patrick Parham

University of Maryland, USA

Yan Shvartzshnaider

York University, Canada

Abstract

This study analyzes the meanings and technical mechanisms of privacy that leading advertising technology (adtech) companies are deploying under the banner of “privacy-preserving” adtech. We analyze this discourse by examining documents wherein Meta, Google, and Apple each propose to provide advertising attribution services—which aim to measure and optimize advertising effectiveness—while “solving” some of the privacy problems associated with online ad attribution. We find that these solutions define privacy primarily as anonymity, as limiting access to individuals’ information, and as the prevention of third-party tracking. We critique these proposals by drawing on the theory of privacy as contextual integrity. Overall, we argue that these attribution solutions not only fail to achieve meaningful privacy but also leverage privacy rhetoric to advance commercial interests.

Corresponding author:

Lee McGuigan, Hussman School of Journalism and Media, University of North Carolina at Chapel Hill, Carroll Hall, CB 3365, Chapel Hill, NC 27599 NC, USA.

Email: leemcg@unc.edu

Keywords

Adtech, contextual integrity, critical data studies, digital advertising, online tracking, platforms, privacy, surveillance

Introduction

Business models based on surveillance and permissive information flows face intensifying scrutiny from regulators, policymakers, and civil society groups (e.g. Federal Trade Commission (FTC), 2022; Mizarhi-Borohovich et al., 2023; Veale and Borgesius, 2022). Platform companies like Google, Meta, and Apple now promise that privacy will be a central design value in the reconstruction of online advertising (Apple, 2021; Bindra, 2021; Mudd, 2021). This sounds like news worth celebrating, turning the page on the personal-data free-for-all that accompanied the rise of advertising technology, or “adtech” (Crain, 2021; Turow, 2011; Zuboff, 2019). We should hold our applause, however, until we know what “privacy” means to these companies, and how those definitions may be inadequate and/or productive of self-advantageous relationships (Greene and Shilton, 2018; Kollnig et al., 2022; Scharlach et al., 2023).

This study assesses adtech’s reformist rhetoric by examining proposals for “privacy-preserving” advertising attribution. Attribution is a process for measuring advertising effects by matching information about users’ media and marketplace activities (Smith, 2019). It requires intermediaries to produce and join records of advertising exposure or engagement, on one hand, and subsequent purchases or other valued actions (e.g. app downloads), on the other. Attribution essentially assigns credit for marketing outcomes to specific advertising efforts; it thereby lets advertisers and their agencies determine and possibly improve their return on investment (ROI), and, in some cases, it allows revenue to be allocated to the publishers, apps, and intermediaries deemed responsible for “causing” certain consumer behaviors. Because its mechanics rely on persistent surveillance, advertising attribution has empowered companies that are well-positioned to monitor users at multiple touchpoints—such as Google, Meta, and, increasingly, Apple—and it has stimulated demand for tracking and analytics services (McGuigan, 2023; Srinivasan, 2020; Van der Vlist and Helmond, 2021).

Attribution in digital advertising has been executed mainly using third-party cookies, conversion pixels, and mobile device identifiers that let marketers track individuals across websites and apps (MacKenzie, 2021). But those identification and measurement instruments are in transition: web browsers are phasing out support for third-party cookies, and Apple now requires app developers and ad networks to get opt-in permission from users to access device IDs and measure behaviors across apps (Graham, 2022). In response to these changes, and under pressure to curb surveillance advertising’s obvious abuses, adtech companies have promised to incorporate privacy-enhancing technologies (PETs) into their attribution services. Google, Meta (and Mozilla), and Apple have each outlined plans for using computational techniques to help advertisers continue measuring and optimizing the effects of their campaigns, while at the same time preventing unauthorized actors from tracking individual consumers or covertly extracting personal data. Each company is using PETs to navigate the tension between public pressure and business interests. A critical comparison of their maneuvers provides a glimpse on how

the adtech sector conceptualizes privacy “problems” in general, and how the specific “solutions” promised by these companies reflect aspects of their reputations, market positions, and infrastructural or platform power.

Our study discerns and compares the meanings and mechanisms of privacy conveyed in these attribution proposals. This sort of clarification is of urgent importance. Regulators and policymakers around the world are seeking to codify privacy in digitally-mediated environments (e.g. European Commission, 2022; FTC, 2022); meanwhile, adtech companies are appropriating the term in public relations and using their dominant positions to encode strategic definitions of privacy into information and market infrastructures (Veale, 2022). New proposals for ad attribution services are political instruments that stake out the legitimate boundaries of privacy, surveillance, datafication, and corporate power. This is a critical moment to clarify the meanings, contradictions, influencing forces, and implications of “privacy-preserving” adtech.

Based on a critical discourse analysis of their attribution proposals, we argue that Google, Meta/Mozilla, and Apple are each promising reforms that leverage (1) long-standing but limited definitions of privacy and (2) elaborate but techno-solutionist computational mechanisms. Addressing multiple audiences in a vaguely technical idiom, these proposals frame a discursive space where each company’s solution can do the work of legitimizing corporate data governance and platform-imposed “privacy.” They make sense by inviting the policymakers and other publics interested in these documents to picture the world in terms of security threat models, individual harms, and the “creepy” indignities associated with furtive tracking and profiling. While these initiatives may make progress on some real problems, they fail to contend with the broader ecosystems of surveillance and data capitalism. They may also further normalize dubious information flows, dismissing the possibility that attribution’s features, to say nothing of its bugs, raise privacy (and other) problems that are not eradicated by technical fixes.

Building on the latter point, we consider how the very notion of privacy-preserving attribution implies an extension of economic priorities and platform power within the mediation of social life. These proposals assume that the use of PETs is sufficient to justify information flows that combine media and market behaviors. We contend, by contrast, that the legitimization of attribution reflects an effort to shift the expectations surrounding ad-supported media: from an arrangement wherein advertisers are entitled to measure audience attention at the site of media exposure, to one wherein advertisers get to measure advertising effects by observing both the site of media exposure *and* the sites of subsequent consumer behavior. Adtech companies may feel compelled to impress with cryptographic techniques and self-regulatory promises because a definition of privacy rooted in social relations could invalidate the entire enterprise of attribution.

“Privacy-preserving” attribution: background and literature review

Adtech, surveillance, data capitalism

Digital economies depend on forms of data processing and analytics that create well-documented tensions with privacy, as well as related concerns about discrimination and

corporate and state power (Binns, 2022; Gandy, 2021; McNealy, 2022; West, 2019). Proponents of data capitalism, by contrast, argue that privacy impedes the social and economic progress promised by friction-free informational flows (Deighton and Kornfeld, 2020). They view market-oriented data governance as a fair trade-off, productively delimiting privacy for the sake of innovation, efficiency, convenience, and wealth accumulation (see Baik, 2020; Cohen, 2013). Adtech fuels some of the most explosive privacy quarrels, since it is both an exemplar of data-driven fortune seeking and a gateway through which almost all Internet users have been enrolled into systems of routine commercial surveillance—systems that are prone to failure and abuse (Maréchal, 2018).

Personal data and consumer profiling are central to digital advertising (Turow, 2011). While the advertising industry has always tried to communicate as exclusively as possible with people whom marketers consider valuable, digital advertising has become increasingly reliant on pervasive surveillance (Crain, 2021). Large platform companies, such as Google, Meta, and Apple, claim an outsized share of revenue in digital advertising markets, due in part to their comparatively greater access to user and marketplace data (Srinivasan, 2020). Smaller market actors worry about their dependency on these platforms, who, in the name of “privacy,” hoard data in ways that make advertising markets less transparent and competitive (Cyphers, 2021). These asymmetries may be compounded by analytics techniques that seem to further improve “privacy” by replacing some of the signals generated through direct behavioral tracking with probabilistic inferences generated by machine learning models (Kak and West, 2023). The latter approaches favor companies with massive computing power and a position in the supply chain that allows them to collect “first-party” data (Kollnig et al., 2022). One of the advertising functionalities most affected by ongoing changes to “privacy” is attribution.

Attribution

Attribution is a process that documents users’ engagement with advertisements and connects those records with observed marketplace outcomes. The purpose is to help advertisers determine—and ultimately lower—the cost of acquiring customers or achieving other objectives. Attribution also facilitates the allocation of revenue to publishers or intermediaries for advertising transactions that are based on user actions (e.g. purchases, downloads). Presently, attribution uses third-party cookies or mobile device identifiers to recognize individuals across sites or apps and conversion pixels to record user behaviors.

Surveillance and identification are critical to attribution since it is, at its root, a claim about advertising effects. Attribution requires detailed accounting of user behavior to confirm the order of causation and to rule out alternative influences. Until recently, attribution claims were usually derived from the “last click” (i.e. the most recent advertising event got credit for the marketplace outcome); increasingly, though, companies are using a “multi-touch” approach, wherein credit is divided across all the events deemed to have contributed to the outcome, often using machine learning (Clark, 2021). Multi-touch attribution thus activates a *more* surveillant process, since it implies a fuller inventory of the possible influences on consumption. Perfect record-keeping is impossible, of course, so attribution claims involve probabilistic modeling and rarely inspire full confidence

among the professionals using these measures. Nevertheless, the goal of attributing consumer behaviors to advertising events has motivated organizational and infrastructural investments in surveillance, data processing, and data sharing (McGuigan, 2023).

It follows that attribution raises privacy concerns. Mozilla even admits that “current attribution practices have terrible privacy properties” (Thomson, 2022). Companies now seek to maintain existing capabilities, which are still in demand, while complying with new rules and norms. We contend that attribution provides an interesting case study for examining the advertising industry’s privacy rhetoric. Attribution is a key functionality provided by adtech vendors, yet it has been understudied in critical literature on marketing, surveillance, and privacy (for an exception, see Smith, 2019). It is also particularly well-suited to an analysis informed by a theory of privacy as “contextual integrity” (CI) (Nissenbaum, 2009), since attribution requires the collection and matching of data generated across multiple sites of user behavior. Attribution’s core function is to join records created when users encounter advertisements embedded in media content, with records created when users make purchases or download apps on other sites. In short, it requires data flows that encompass both media usage and marketplace behavior.

This raises a key dilemma: What definition(s) of privacy can be reconciled with attribution’s basic processes?

Meanings and mechanisms of privacy

The meaning of privacy is subject to ongoing debate (e.g. Citron and Solove, 2022), varying across legal, philosophical, and technical disciplines (Nissenbaum, 2009). Privacy definitions, and the mechanisms for operationalizing them, are situated within political-economic contexts; as such, they both reflect and shape dynamics of power that structure the experiences of designers, workers, and consumers who develop or interact with socio-technical systems (Greene and Shilton, 2018). The privacy discourses circulated through corporate documentation also help companies position themselves in relation to regulators and other stakeholders by aligning with desirable principles (Scharlach et al., 2023). Some principles have been especially influential at defining what privacy will mean in policy and practice (Cohen, 2013; Epstein et al., 2014). We highlight some perspectives identified by Nissenbaum (2009) as key frameworks for theorizing privacy.

Privacy protections predominantly rely on an “informed consent” model, which puts the onus on the user to comprehend the associated benefits and harms and adjust controls around what information to share, with whom, and for what purpose (Solove, 2013). This paradigm understands privacy as *control* over personal information, and its proponents push for greater transparency in disclosing information handling practices. The main focus within this approach is identifying different information categories and purposes, often through recourse to dichotomies such as private versus public, personal versus non-personal, and sensitive versus non-sensitive information.

Another dominant perspective defines privacy as *limiting access* to individuals’ data (Nissenbaum, 2009: 69–71). The basic idea is that privacy increases as the amount of information disclosed about an individual, or the number of parties privy to it, decreases. This notion of privacy is strongly coupled with security mechanisms such as encryption,

multi-party computation (MPC), and differential privacy. Those mechanisms further index concepts that are related to an access-based definition of privacy, including confidentiality and secrecy.

These approaches have merits, but they are ultimately inadequate when it comes to evaluating emergent socio-technical systems (Nissenbaum, 2009). The drawbacks of using dichotomies, such as sensitive/non-sensitive, and the informed consent models to define privacy become evident when we are inundated with information collection practices that challenge established norms. Likewise, evolving means of extracting inferences or predictions from large datasets undercut efforts to solve privacy problems through “anonymity” (Barocas and Nissenbaum, 2014).

Nissenbaum’s (2009) theory of privacy as CI addresses the limitations of these other perspectives, particularly for analyzing new and complex technologies. In contrast to control- or access-based accounts of privacy, CI emphasizes the need to collect and circulate information in accordance with norms that promote the purposes, functions, and values of a given social context. CI’s heuristic framework evaluates the appropriateness and legitimacy of norm-breaching information flows by considering their ethical and social implications. CI thus requires that information flows be justified by more than technical means of confidentiality, anonymity, or consent management. For example, an instance of data processing could be made “more private” according to an access-based definition while still remaining out of alignment with social norms and values. The privacy-preserving attribution solutions examined herein may exemplify this latter situation, and this leads us to the following research questions about the proposals from Google, Meta/Mozilla, and Apple:

RQ1. What do these companies mean when they talk about privacy?

RQ2. How do their solutions differ from each other in terms of privacy?

RQ3. How might each company’s approach to privacy reflect political-economic factors?

These questions lead us to a further judgment about whether we should be satisfied that these attribution solutions “preserve privacy.” We reach our conclusion here by drawing on CI and critical political economy (CPE). These two approaches pair productively. While CI puts norms at the center of its analysis, it does not prescribe the normative content of a given context or social sphere. By contrast, CPE inherits a normative thrust from moral philosophy, with a commitment to praxis in the foreground (Mosco, 2009), and work in this tradition shows how battles over privacy, discrimination, and other concerns are linked to the commodification and privatization of information and its infrastructures (e.g. Crain, 2021). Political-economic critiques of datafication that begin by situating information flows within social relations (e.g. Viljoen, 2021) provide especially useful insights for considering privacy and adtech. CPE also helps show how corporate privacy discourses service lobbying and public relations efforts (McGuigan et al., 2023); these attribution proposals advocate for privatized, techno-centric, and self-regulatory solutions, framing privacy and media financing as problems best solved by platform

companies and smuggling these issues out of the realms of collective action, strong public governance, and political debates about values and power. These conceptual tools help us critically interpret companies' claims about privacy-preserving adtech.

Methods

To analyze adtech's privacy discourse, we assembled a corpus of publicly available materials that describe the purpose and functionality of the attribution solutions designed by Google, Meta/Mozilla, and Apple. These texts address multiple audiences whose priorities and technical expertise vary; our corpus includes documentation written for software systems administrators, as well as publicity produced to explain these initiatives to non-experts. We limited our corpus to materials published by these companies or their employees on corporate websites, owned developer blogs, and, in two cases, GitHub. We collected texts that directly referenced attribution or synonymous functions (e.g. conversion measurement). Our corpus comprises 18 texts (6 representing each company), which ranged in length from 680 words to 8400 words and averaged around 3200 words.

We chose to compare these three companies because, on one hand, they operate large platforms with nearly unrivaled access to consumer and market data, they set commercial terms for partners, customers, and competitors, and they own commanding shares of digital advertising revenue (Kollnig et al., 2022; Nieborg and Poell, 2018; Srinivasan, 2020); and, on the other hand, they are each positioned differently in the adtech industry and in advertising and data supply chains, and they vary in their abilities and means of exercising infrastructural power (Van der Vlist and Helmond, 2021; Veale, 2022). This allows us to consider how these companies operationalize privacy in relation to their different priorities, advantages, and vulnerabilities, while also highlighting common assumptions about privacy and its public relations appeals, which help to narrow its meanings and contain its political force, likely in ways that help maintain or even extend the commercial and data governance relations that these platforms capitalize via their adtech businesses.

We coded explicit or latent meanings of privacy in our corpus using an iterative process. First, we discerned themes inductively; this produced an initial list of "privacy meanings" that reflected the terminologies used in each document. Since the three companies use different terms to refer to similar things, we then consolidated and refined our coding categories to capture core underlying principles. By synthesizing the themes that emerged inductively with the concepts found in the literature reviewed above, we arrived at five categories for classifying "privacy meanings" across these attribution proposals (see Table 1). We then coded statements that (1) articulated or implied a privacy meaning and/or (2) described a method, or "mechanism," for achieving privacy (e.g. encryption). We also noted instances in these documents that reference trade-offs or tensions between privacy and the commercial objectives these tools are designed to achieve. Finally, we contextualized the findings by considering each company's position in the adtech industry, the likely audiences for the documents in the corpus, and each solution's stage of development (e.g. prospective, experimental, implemented).

Table I. Privacy meanings and privacy-supporting mechanisms.

Privacy meanings		Privacy-supporting mechanisms									
	Anonymity	Limiting access	Anti-tracking	Control	CI	Differential privacy	Local (on-device/ on-browser) processing	Multi-party computation	Data aggregation	Obfuscation	Encryption
Definition	Delinking data from identifiable persons, preventing relinking	Minimizing information shared and/ or number of parties who can access it	Preventing third-party observation and profiling across sites, apps, and devices	Notifying users about how personal data are collected, processed, and/or shared	Information flows adhere to norms of different contexts	Random noise is added to a dataset to hide individuals	Data or identifiers are processed or readable only on the user's device or browser	Multiple computers process parts of a dataset so that no one party sees all the data	Reporting data about a population rather than individuals	Obscuring details in the data	Hiding information from anyone without an authorized key
Related mechanisms	Aggregation Differential privacy Multi-party computation Obfuscation	Encryption Multi-party computation Local processing	Local processing Obfuscation	Consent (opt-out)	CI heuristic						
Related meanings						Anonymity	Limiting access Tracking	Anonymity Limiting access	Anonymity Limiting access	Anonymity Tracking	Limiting access

CI: contextual integrity.

The privacy meanings we coded were defined as follows:

Anonymity: Any effort to prevent information from being associated with an identifiable person. This includes the initial anonymization of personal data, as well as subsequent defenses against adversaries trying to deanonymize that data. Common mechanisms for achieving anonymity include aggregation, obfuscation, MPC, and differential privacy.

Limiting access: Any effort to limit the information collected, processed, shared, or revealed about an individual. This includes references to secrecy and confidentiality, and it corresponds to mechanisms such as encryption and on-device data processing. Access can be limited along two dimensions: the *amount of information* about a user that is accessible; and the *number of parties* able to access information about a user.

Preventing third-party tracking and profiling: These documents often define privacy inversely, by referencing privacy violations. We coded instances in which the companies claim that their solutions are privacy-preserving because they prevent third-party tracking and profiling. This anti-tracking category is, in fact, a subset of Limiting Access; but it focuses particularly on third-parties and appeals directly to popular anxieties about “creepy” surveillance by unknown companies. We determined that it is important to capture the tendency among adtech companies to claim the prevention of third-party tracking as a privacy trump card.

Control: The ability of users to control information about themselves. This is typically related to consent mechanisms that let users opt-out of or opt-into commercial data collection and usage.

CI: CI defines privacy as the appropriate flow of information according to the norms, priorities, and values of a given social sphere or “context.” None of the solutions align with a rigorous CI definition; nevertheless, we marked instances where context and/or user expectations were mentioned.

In addition to these privacy meanings, we coded references to technical means of achieving privacy, which we call “privacy-supporting mechanisms.” As with the privacy meanings, we categorized mechanisms through inductive and deductive coding. The mechanisms most evident in the corpus are: data aggregation; differential privacy; encryption; MPC; obfuscation; and local (on-device/browser) computing and storage. Depending on the context of implementation, any given privacy-supporting mechanism may correspond to more than one privacy meaning, and multiple mechanisms and meanings may be activated simultaneously. To give one example, some of these solutions obfuscate metadata about recorded events, such as the exact time when an anonymous user purchased an advertiser’s product; this both limits the amount of information made accessible and it helps prevent deanonymization.

Findings

Our clearest finding is that “privacy” is wielded throughout these documents with a positive valence but with few direct statements of its specific meaning. We find frequent expressions of normative commitments to privacy, without much elaboration of

privacy's normative content (i.e. why it is important), apart from implied benefits of information security. Privacy is also treated in a "descriptive" sense (see Nissenbaum, 2009: 68–69), as a property that these attribution solutions will "enhance," "increase," or "protect." Across these normative and descriptive claims, though, almost none of the proposals expressly defines what privacy means. Instead, privacy meanings are implicit and often vague, evoked through reference to practices that violate privacy (e.g. "tracking") or mechanisms that protect against privacy harms. Furthermore, although all three companies align themselves with privacy as a value, they suggest that the extent of privacy must be balanced against economic priorities (which are themselves justified through normative appeals—namely, that advertising is an essential guarantor of the open Internet).

We find that the solutions all converge primarily around definitions of privacy as *anonymity*, as *limiting access* to individuals' data, and as the *prevention of third-party tracking and profiling*. The following sections describe each attribution solution and the privacy meanings and mechanisms encoded therein.

Meta/Mozilla's Interoperable Private Attribution

Meta (then Facebook) signaled its intention to use PETs for more "private" measurement of ad effectiveness at least as early as 2021, and, together with Mozilla, it published an overview of the Interoperable Private Attribution (IPA) system in January of 2022. The proposed solution uses local identifiers called "write-only match keys" to link "source events," such as viewed impressions, with "target events," such as purchases or app installations, from the same user. Match keys are set on a user's device or browser by designated "providers," such as Facebook, Google, and Twitter, when users log in to those platforms or apps. Any participating website can use those match keys to associate what happens on their site with an individual user, but the identity of the match key is only readable by the local device or operating system. Upon leaving the device, event records are matched in a confidential way via a MPC arrangement involving double encryption-decryption by "trusted helper" servers. The purpose of MPC is to collectively process data about source and target events without letting any single party access or reconstruct the behavioral records associated with each user. Finally, the system produces aggregated attribution reports for advertisers and publishers. Access to the reports is limited by a "privacy budget," imposed on each interested party, that gets depleted as they ask for information. The privacy budget prevents anyone from repeatedly querying the servers that process individual information so as to disaggregate and deanonymize conversion reports. User identity is further masked using differential privacy, a technique which adds a calibrated amount of distortion to a dataset so that insights may be derived about a population while concealing each individual's data.

The dominant privacy meanings applied in the IPA documentation are anonymity and limiting access. A key privacy promise is that the identifiers used to measure each individual's activities across sites, apps, and devices—"write-only match keys"—are not readable by third-parties, and so they "cannot be used for tracking or profiling" (Savage et al., n.d.). The attribution reports are considered "private" because advertisers and adtech vendors see aggregate data and are unable to re-identify individuals.

Meta comes close to articulating an explicit definition of privacy as limiting access, but with some discrepancies that bear noting. Documents predating the IPA proposal discuss how PETs will minimize the amount of data that the company collects or processes. “Ensuring privacy throughout our apps while reducing the data we collect is a long-term effort,” one text explains. It later alludes to the sophistication of this class of privacy mechanisms and their ability to satisfy advertisers’ business demands: “PETs involve advanced techniques drawn from the fields of cryptography and statistics. These techniques help minimize the data that’s processed while preserving critical functionality like ad measurement and personalization” (Facebook, 2021).

This initial position stands in subtle but critical contrast to Meta’s eventual proposal (with Mozilla) for IPA, which offers perhaps the clearest definition of privacy in the whole corpus: “Our privacy goal is to limit the total amount of information IPA releases about an individual over a given period of time” (Taubeneck et al., 2022a; emphasis added). One of the key questions motivating the IPA design is: “How can we make sure fewer companies have access to our personal data?” (Savage et al., n.d.: 20).

Our findings thus document a shift from the promise of *data minimization*—reducing the amount of data *collected* and *processed*—to the promise of limiting the amount of data that is *released* or *shared* and the number of parties involved. This is a much more permissive approach to privacy than preventing personal data from being generated and stored in the first place. It sidesteps questions about the legitimacy of the information flow and instead purports to make that flow “more private” by limiting access.

Despite this hedge on data minimization, IPA is the most ambitious of the attribution solutions we examined. Compared with the others, Meta/Mozilla make the boldest privacy claims and propose the most demanding computational and cryptographic mechanisms. That said, this proposal is also the most prospective. Key details remain indefinite—such as whether or how data from attribution reports are fed back into the optimization of ad targeting, and who will operate the “trusted” servers. Since Meta/Mozilla begin one document by stating (as if a matter of fact), “Advertisers need accurate reporting about how their ad campaigns are performing” (Savage et al., n.d.: 3), we should expect tensions and compromises to arise as IPA enters the messy politics of implementation.

Google’s attribution reporting API

Google’s solution has been in use since 2021. The documentation and publicity surrounding it also portray privacy mainly as limiting access and anonymity. Like IPA, Google’s Attribution Reporting API links source and target events while “minimizing” information sharing and adding noise to the produced reports. Implicit here is the notion that cross-context measurement (i.e. the joining of ad exposure or clicking events with conversion behaviors by a unique user) does not constitute tracking if it is executed locally on the user’s device or browser. “No cross-site identifier is used and no detailed cross-site browsing activity leaves the device,” Google explains. “A small amount of information is joined across sites—enough to measure conversions, but not enough to track [a user’s] activity across sites in detail” (Nalpas et al., 2023).

Google's solution provides two types of reports: Event-level reports attribute an ad click or view (source event) with limited data related to ad conversion (target event); and summary (aggregated) reports provide high-level insights into the link between source and target events. Event-level reports employ a delay in delivery, the addition of noise, the generation of "fake reports," and limitations on the number of reports generated per click and view. The data that serve as the basis for summary reports are encrypted when sent from a user's browser to the aggregation service, which produces the ultimate summary report for the advertiser. The aggregation service provides access to the report by decrypting the report and adding noise in a differentially private way. To further protect user anonymity, this report is sent with random delays and report queries are limited. These reports allow advertisers to quantify campaign spending, number of conversions, ROI, geographic location of conversions, and publisher site where conversions occurred.

Google's solution appears to make the most concessions to commercial demands. Its documentation does not clarify cross-site visibility and the mechanisms to limit identification to the same degree as the other solutions analyzed. How the source and trigger event data are matched in source-level reporting documentation, beyond a general overview, is also unclear. Google does not mention how the company's own view of data is impacted and if data are shared across its suite of products. In addition, event-level reports allow advertisers to optimize toward more efficient ROI, using key performance indicators to determine the optimal type of ad to serve and to train machine learning capabilities (Nalpas et al., 2023). Thus, the system still facilitates information flows and inferences that affect individuals' experiences and opportunities in online environments.

Apple's SKAdNetwork and private click measurement

Apple introduced the first iteration of SKAdNetwork in 2018 to enable attribution measurement while reducing the amount of information collected by adtech intermediaries. It updated the system in 2020 and 2022, with additional features, following the initiation of its App Tracking Transparency (ATT) protocol, which governs the sharing of device identifiers. For Apple, the meaning of privacy is overwhelmingly related to the prevention of "tracking," which refers to individual-level, cross-app or cross-site surveillance by third-parties. "It is becoming clear that business models that rely on tracking aren't sustainable," Apple admits in a video presentation to developers. "We recognize the importance of providing a more private way to measure ads to help you thrive in this changing ecosystem" (Apple, 2021). The company suggests that users "don't expect to have an invasive experience on the web, where their interests, behavior, and personal information is stored and tracked" (Apple, 2021). Apple (2021) also emphasizes control to a greater degree than its rivals: "Tracking infringes on a user's privacy without giving them the ability to identify, understand, or consent to what's being shared about them."

Apple's own attribution services report anonymized data to advertisers, websites, or apps. The company offers two "privacy-preserving" solutions, whose technical details depend on the venues of ad delivery and ad conversion. For "web-to-web" and "app-to-web" attribution, Apple uses "Private Click Measurement" (PCM), an "on-by-default"

solution that promotes privacy as limiting access, linking source and target events locally, on users' browsers, up to 1 week from ad impression to conversion. Only the browser on the user's device can match source and target events to actual users, and that data, according to Apple, never leaves the local device. To ensure anonymity and confidentiality, the reports are encrypted and signed to prevent fraud, provided to both ad impression and conversion outlets, and are delayed by 24–48 hours to further obfuscate user identity (Apple, n.d.). Conversion destinations are only registered at the top-level domain to prevent tracking users through a chain of subdomains that can reveal information about the attribution source.

For “app-to-app” and “web-to-app” attribution, Apple deploys SKAdNetwork 4.0, a solution for measuring the impact of advertising on app downloads and engagement. Any click on an ad for an app generates a report that is stored locally; the report includes the unique IDs of the publisher, advertiser, and ad network involved and a “hierarchical id”—a 4-digit number that can include information on the campaign, approximate user location, and the type of ad served (Apple, 2021). Once the user engages with the app, conversion reports are sent in three different time stamps to capture multiple conversions over a long period of time—sometimes more than a month after downloading the app. To ensure anonymity, attribution data that are sent to the ad network have no identifying user information, and are usually sent in a delay of hours or days. Level of detail in the conversion data is linked to the number of app installations, known by Apple as the “crowd anonymity function” (Apple, n.d.). Apple considers the SKAdNetwork solution private-by-design and thus does not apply its ATT consent control mechanism before calling those attribution APIs (Apple, n.d.).

Apple has long deployed privacy rhetoric within its branding and public relations, and the company leverages this in promoting its attribution solutions. The privacy meanings emphasized in Apple's documentation disadvantage its rivals—namely, Meta—by defining third-party attribution systems as illegitimate “tracking” and thus subjecting them to prohibitions and user control. By contrast, these privacy meanings authorize Apple's first-party ad “measurement” services.

Discussion

Limited privacy perspectives

Our research shines light on the meanings embedded in “privacy-preserving” attribution and the mechanisms that Google, Meta/Mozilla, and Apple are using to encode those meanings into technical systems and corporate policies. These companies are vague and selective in how they define privacy, yet they are leveraging the term's positive connotations to justify self-regulatory solutions that will structure data governance relations with users, customers, and competitors. Rather than grappling with these deeper issues, the “privacy-preserving” attribution solutions we looked at focus largely on complying with the strictures of a post-cookie world and on disavowing the “creep factor” associated with third-party tracking. Consequently, they exclude from concern the ongoing data collection and usage conducted by “first-parties,” and they neglect the larger “surveillant assemblage” (Haggerty and Ericson, 2000) in which adtech is embedded. In relation to

the latter point, all solutions boast that they limit the information revealed about individuals through technical restrictions, but they do not acknowledge the flexibility advertisers still have to specify characteristics for targeting on the “line-item” level. A line item is a string of taxonomic descriptors that an advertiser or demand-side platform uses to define certain features about the delivery and targeting of an ad campaign, such as publisher site, geography, demographic details, and creative content. It is possible that clever manipulation of line items will let advertisers evaluate ad performance in granular detail, regardless of the technical restrictions imposed through an attribution system. For example, Google Summary Reports allow advertisers to see conversion counts and campaign spending broken down by targeting categories. This privacy feature promises to protect users by only sharing campaign level IDs, rather than user IDs; however, by applying targeting categories at the line-item level before interacting with the platform, advertisers could potentially compromise de-identification efforts through permutations of line-item targeting. They could, in effect, turn the campaign ID into something that works more like a pseudonymous user or cohort ID.

At a broader level, platforms’ privacy perspectives appear to be inspired by cybersecurity threat models. Consequently, many of the touted features are designed to be robust against malicious activity. Implicit here is the claim that privacy violations are, almost by definition, associated with unsanctioned actions. Attribution, in and of itself, raises no concerns in this account, and the record-keeping required for attribution is justified by business needs. This orientation lends itself toward discrete (if highly creative) solutions, wherein privacy becomes an objective property that can be “increased” with cryptographic techniques. Preventing abuse is beneficial, of course; but coming to terms with adtech’s privacy problems requires a more holistic approach. Privacy is not just about preventing data breaches or identity theft. To be useful in thinking critically about socio-technical systems, a theory of privacy needs to grapple with the political economy of data and the integrity of social life (Nissenbaum, 2009; Viljoen, 2021). The corporate, security-oriented, and self-regulatory approaches to privacy that dominate the discourse around data governance in adtech tend to obscure or elide these power relations (see also Marwick, 2022).

Differences among solutions

Despite many commonalities, the proposals do exhibit differences. For example, Meta/Mozilla emphasize general features of the solution, while Apple’s texts convey a significant amount of technical information. Google’s documentation instructs advertisers on how to incorporate its solution into their routines and describes the solution’s impact on campaign activation and reporting processes.

These differences reflect variations in implementation status: Apple’s solutions have not yet been widely adopted, and the Meta/Mozilla proposal is entirely prospective, while the Google solution is already used across parts of the digital advertising industry—although the protracted deprecation of cookies softens the urgency for change. In addition, each company’s market positions may help explain differences in these solutions. Apple and Google, which own browsers, devices, and operating systems, restrict the view of what is captured to their own ecosystems, effectively legitimizing the

enclosure of data within their walled gardens. This gestures at privacy compliance while also making their adtech products and services more valuable to advertisers, as compared with competitors whose “third-party” status excludes them from the data assets secured behind the garden walls. In the case of Meta, which is a social media platform, a partnership with Mozilla helps advance a cross-device and cross-browser solution that could require other browser and device operators to organize around a single standard (Hercher, 2022). Having recourse to less infrastructural power than Google and Apple, Meta may be trying to shift the ecosystem’s dependencies to the software or application level, where it has advantages of scale and reach. Finally, Apple’s emphasis on control (via consent) and tracking prevention reflects (1) Apple’s position in the adtech stack, as an operating system that can administer privacy permissions and (2) its comparatively minor stake in digital advertising. Apple exploits its position by imposing consent requirements on competitors like Meta, whose business model is deeply dependent upon what Apple defines as third-party “tracking,” while exempting its own first-party “measurement” from this definition of tracking and the corresponding mechanisms of user control. Apple is translating this advantage into a growing market share in digital advertising (McGee, 2021). Overall, Apple’s approach impairs other surveillance advertising companies while sanitizing its own expansion in that business.

Contexts of measurement and tracking

A brief consideration of CI—a theory of privacy as a social good—and CPE—a framework sensitive to power relations and democratic norms—can sharpen our analysis of the problems and contradictions within “privacy-preserving” attribution solutions. We observed some instances in which companies use CI vocabulary, promising to prevent “cross-context” tracking or information flows and/or to respect user expectations. For example, Meta/Mozilla claim that their solution prevents “cross-context tracking” and the revelation of “any cross-context information” (Taubeneck et al., 2022b). What constitutes a “context,” however, is left undefined. Apple’s (2021) documents also mention contexts and user expectations, with slightly more detail. Describing its PCM service, Apple says it “is intended to support privacy-preserving *measurement* of clicks across websites or from apps to websites. It is not intended to be used to *track* users, events, or devices across those contexts” (Wilander, 2021; emphasis added). This seems to imply that each site or app represents a context. Crucially, though, there is a distinction lurking between “measurement” and “tracking.” The documents encode tracking with a negative connotation, as a way for someone to follow and record individuals’ activities without justification or consent. By contrast, “measurement” is presented as a legitimate necessity, authorized by the premise that advertisers need, and deserve, to know how efficiently their campaigns are achieving sales or other objectives. Thus, Apple has license to “measure” user clicks across contexts, while ad networks’ methods of following the causal chain of attribution are barred as illegitimate “tracking.” This rhetorical move has concrete implications for anyone who makes a living from Apple’s app ecosystem, and it illustrates the power of policy language to materialize certain relationships, privileges, and obligations (Gillespie, 2010; Greene and Shilton, 2018; Scharlach et al., 2023).

It appears that the principle heuristic for defining appropriate information norms in adtech is the proximity relationship between parties (or, to put it another way, the ownership of the sites where data extraction and usage occurs). In this formulation, a first-party relationship assures the integrity of a context, while a third-party relationship is a *de facto* violation. Surely, third-party tracking activates problematic information flows. But, the legitimacy of the information flows here called “measurement” is not secured simply by being conducted within a first-party relationship; per CI theory, that determination must be rooted in considerations of social values and purposes. Attention to the political economy of media, platforms, and data further enhances those considerations.

We suggest that the industry’s treatment of “contexts” does important political work. The legitimization of attribution represents a silent extension of media marketization and the platform enclosure of social life (Wu and Taneja, 2021). The embedding of attribution processes in digital media effectively renegotiates the implied transaction between audiences, publishers, and advertisers: from an exchange based on attention, to one based on buying behavior. Attribution implies that marketers are not just entitled to measure “audience attention,” to confirm that their ads are distributed properly; rather, marketers are entitled to measure the effects of advertisements, by following audiences beyond the sites of ad exposure and into the marketplaces where those audiences become active consumers. This is a corporate-imposed shift in relationships that requires scrutiny. For attribution to be “privacy-preserving,” in the sense of comprising legitimate information flows, we would have to accept that media and marketplaces are coterminous—that a prevailing purpose of news, entertainment, and social media is to produce not just audiences but consumers. The industrial logic of commercial media in the United States has always centered around *bona fide* consumers (Meehan, 2005), but its implementation is a site of social struggle, as people resist commodification of their leisure time and attention (Smythe, 1981). Justification for this emergent attribution arrangement is not assured by techno-solutions that configure privacy as anonymity or limiting access, and its normalization should be considered part of the corporate cultivation of resignation to commercial surveillance (Draper and Turow, 2019; McGuigan et al., 2023). CI and CPE are useful analytics—and troublesome ones from adtech’s perspective—because they demand an account of the assumption at the core of all these solutions: Why is the measurement of advertising effects, and the relationships necessary for joining media and marketplace data, integral to the socio-technical systems that mediate our social and personal lives and our access to news and culture? The rhetoric in the documents we analyzed does not answer this question. Accepting that attribution can be private requires an admission that the production of consumers deserves pride of place among the values and priorities commonly attached to media systems in a democracy (see, for example, Napoli, 2019; Pickard, 2019).

Conclusion

Google, Meta, and (to a lesser extent) Apple are advertising giants. They have benefited from perverse data collection practices for years. Their executives declared the death of privacy and invested heavily in data-extractive technologies and commercial relations. While we welcome initiatives to reverse this trend, our examination shows that what

these companies refer to as privacy is a vague and narrow conception thereof. The addition of PETs into attribution systems may mitigate some existing problems (such as leaking de-anonymized data), but those technical solutions do not obviate the problems raised by the for-profit use of behavioral data to intervene in people's experiences within digital environments.

Our evidence helps answer the research questions posed above. We see that anonymity, limiting access, and prevention of third-party tracking are the most dominant privacy meanings evoked in the corpus. Furthermore, all the companies we examined are implementing these definitions via similar mechanisms, although some of the details vary. The variations we observed may be related to differences in companies' market positions, their recourse to infrastructural power, and the implementation status of their solutions.

Finally, we can consider whether to be satisfied with adtech's "privacy" reformism. The answer, in short, is no. We have suggested that attribution should be taken seriously because of its importance to key business use cases and its relative omission from analyses of marketing surveillance. Another critical reason for scrutinizing these solutions is that they *do* improve the current state of the art. In some sense, this might be adtech taking its best shot. And yet the information flows fundamental to attribution remain in tension with expectations about the social and personal purposes of media. Adtech proponents argue that privacy concessions must be tolerated to maintain the current shape of the online economy; if marketers lose confidence in the efficiency of their advertising investments, then consumers will lose "free" access to websites and apps. Even setting aside the structural problems and internal contradictions that render that arrangement highly inequitable and unstable (Hwang, 2020), the presumption that advertisers are entitled not just to distribute ads but also to confirm and optimize their efficacy represents an asymmetric extension of corporate priorities deeper into the mediation of everyday life. "Privacy-preserving" attribution further cinches people's engagement with online environments to the accumulation of private profit by claiming to "solve" a set of problems that are admittedly real but, nevertheless, not at the heart of the issue. The upshot of these maneuvers may be further empowerment of leading companies. The advertising industry is constantly selling reinvented versions of itself to clients, investors, and other publics; as symbols of technoscientific progress, adtech's optimization techniques have provided handy discursive resources for those efforts. We argue that PETs are being put to similar use, making privacy a wedge to open even more space for enclosure by the walled-garden companies that are positioned to exploit advantages in computing capacities and data access. In sum, not only do these solutions fall short of achieving privacy, they also exploit privacy rhetoric in ways that may compound corporate concentrations of power and wealth.

This study showcases a particular example of the general tendency for technology companies to define and operationalize concepts and values in self-serving ways (e.g. DeCook et al., 2022; Gillespie, 2010). Adtech's attribution proposals are, among other things, a form of public relations, part of a corporate movement toward platform-imposed "privacy." This movement raises important implications for policy making and platform governance. Framing privacy as a technical achievement empowers platform companies to insert PETs or other self-regulatory mechanisms as

“fixes” to the “privacy problem,” and it diverts attention toward the configuration of internal details rather than the values, social relations, and power dynamics congealed within adtech infrastructures. The frameworks of CI and CPE collectively force these issues into the open, making the purposes and priorities at the root of attribution systems into matters of concern and collective political action. By accepting these proposals on their own terms (however well-meaning their proponents may be), we risk further normalizing the platform enclosure of personal and population-level data and deepening ad-supported media’s highly contestable relations of commodification, discrimination, and exploitation. Challenging adtech’s privacy meanings is a critical step for denying platform companies’ claims of ownership over a privatized—but not privacy-preserving—digital sphere, where social mediation and cultural production are collapsed into an encompassing commercial context.

Limitations and future directions

There is much more to know, and our study has limitations. Some of the texts in our corpus are fairly technical. They are also quite vague, both in that they are written for developers who may be implementing these systems across different software configurations and use cases, and in that some elements of these systems remain prospective or experimental. We tried to compensate by assembling an interdisciplinary research team with diverse competencies. However, there are parts of the corpus that we cannot definitely interpret without inside access to these companies. Future research should use multi-modal methods to triangulate our findings and to further clarify these companies’ definitions and mechanisms of privacy.

This study is also limited by its focus on a small part of a much larger ecosystem. Nevertheless, our findings provide a basis for ongoing work. Most obviously, our procedure could be replicated to analyze the broader discourse about privacy in adtech. Building on analyses of technical and commercial standards development in digital advertising (Cluley, 2020; Gehl, 2014), this work could be extended to study the whole of Google’s Privacy Sandbox initiative, as well as inter-organizational forums like the W3C’s Private Advertising Technology Community Group and the IAB Tech Lab’s Project Rearc. Based on background investigation into those initiatives, we think our conclusions are applicable to this broader discourse; but looking beyond attribution solutions, and beyond communication venues controlled by the biggest adtech companies, would likely yield additional insights.

Acknowledgements

This article benefited from feedback at the 4th Annual Symposium on Applications of Contextual Integrity. The authors also thank Aaron Shapiro, Alison Miller, and the anonymous reviewers for highly constructive feedback.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

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

References

- Apple (2021) Meet privacy-preserving ad attribution. *Apple Developer*. Available at: <https://developer.apple.com/videos/play/wwdc2021/10033/>
- Apple (n.d.) Attributing ads with SKAdNetwork and private click measurement. *Apple Developer*. Available at: <https://developer.apple.com/app-store/ad-attribution/>
- Baik JS (2020) Data privacy against innovation or against discrimination? The case of the California Consumer Privacy Act (CCPA). *Telematics and Informatics* 52: 101431.
- Barocas S and Nissenbaum H (2014) Big data's end run around anonymity and consent. In: Lane J, Stodden V, Bender S, et al. (eds) *Privacy, Big Data, and the Public Good*. New York: Cambridge University Press, pp. 44–75.
- Bindra C (2021) Building a privacy-first future for web advertising. In: Google Ads, 25 January. Available at: <https://blog.google/products/ads-commerce/2021-01-privacy-sandbox/>
- Binns R (2022) Tracking on the web, mobile and the Internet of things. *Foundations and Trends in Web Science* 8(1/2): 1–113.
- Citron DK and Solove DJ (2022) Privacy harms. *Boston University Law Review* 102: 793–863.
- Clark K (2021) Google ditches last-click attribution in favor of machine learning-based model. *The Drum*, 27 September. Available at: <https://www.thedrum.com/news/2021/09/27/google-ditches-last-click-attribution-favor-machine-learning-based-model>
- Cluley R (2020) The politics of consumer data. *Marketing Theory* 20(1): 45–63.
- Cohen JE (2013) What privacy is for. *Harvard Law Review* 126: 1904–1933.
- Crain M (2021) *Profit Over Privacy: How Surveillance Advertising Conquered the Internet*. Minneapolis, MN: University of Minnesota Press.
- Cyphers B (2021) Google's FloC is a terrible idea. *Electronic Frontier Foundation*, 3 March. Available at: <https://www.eff.org/deeplinks/2021/03/googles-floc-terrible-idea>
- DeCook JR, Cotter K, Kanthawala S, et al. (2022) Safe from “harm”: the governance of violence by platforms. *Policy & Internet* 14(1): 63–78.
- Deighton J and Kornfeld L (2020) *The Socioeconomic Impact of Internet Tracking*. New York: Interactive Advertising Bureau. Available at: <https://www.iab.com/wp-content/uploads/2020/02/The-Socio-Economic-Impact-of-Internet-Tracking.pdf>
- Draپر NA and Turow J (2019) The corporate cultivation of resignation. *New Media & Society* 21(8): 1824–1839.
- Epstein D, Roth MC and Baumer EPS (2014) It's the definition, stupid! Framing of online privacy in the Internet governance forum debates. *Journal of Information Policy* 4: 114–172.
- European Commission (2022) The digital services act: ensuring a safe and accountable online environment. Available at: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act-ensuring-safe-and-accountable-online-environment_en
- Facebook (2021) What are privacy-enhancing technologies (PETs) and how will they apply to ads?, 11 August. Available at: <https://about.fb.com/news/2021/08/privacy-enhancing-technologies-and-ads/>
- Federal Trade Commission (FTC) (2022) Commercial surveillance and data security rulemaking. *Federal Trade Commission*, 11 August. Available at: <https://www.ftc.gov/legal-library/browse/federal-register-notice/commercial-surveillance-data-security-rulemaking>
- Gandy OH (2021) *The Panoptic Sort*. 2nd ed. New York: Oxford University Press.

- Gehl RW (2014) *Reverse Engineering Social Media: Software, Culture, and Political Economy in New Media Capitalism*. Philadelphia, PA: Temple University Press.
- Gillespie T (2010) The politics of “platforms.” *New Media & Society* 12(3): 347–364.
- Graham M (2022) More changes loom for online marketers. *Wall Street Journal*, 25 January. Available at: <https://www.wsj.com/articles/more-changes-loom-for-online-marketers-11643150679>
- Greene D and Shilton K (2018) Platform privacies: governance, collaboration, and the different meanings of “privacy” in iOS and Android development. *New Media & Society* 20(4): 1640–1657.
- Haggerty KD and Ericson RV (2000) The surveillant assemblage. *The British Journal of Sociology* 51(4): 605–622.
- Hercher J (2022) Mozilla and meta submit (yet another) privacy ad tech proposal in new W3C group. *AdExchanger*. Available at: <https://www.adexchanger.com/online-advertising/mozilla-and-facebook-submit-yet-another-privacy-ad-tech-proposal-to-new-w3c-group/>
- Hwang T (2020) *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet*. New York: FSG/Logic.
- Kak A and West SM (2023) AI now 2023 landscape: confronting tech power. *AI Now Institute*, 11 April. Available at: <https://ainowinstitute.org/2023-landscape>
- Kollnig K, Shuba A, Van Kleek M, et al. (2022) Goodbye tracking? Impact of iOS app tracking transparency and privacy labels. In: *ACM conference on fairness, accountability, and transparency*, Seoul, Republic of Korea, 21–24 June, pp. 508–520. New York: ACM.
- MacKenzie D (2021) Cookies, pixels and fingerprints. *London Review of Books*, 1 April. Available at: <https://www.lrb.co.uk/the-paper/v43/n07/donald-mackenzie/cookies-pixels-and-fingerprints>
- Maréchal N (2018) Targeted advertising is ruining the internet and breaking the world. *Motherboard*, 16 November. Available at: <https://www.vice.com/en/article/xwjden/targeted-advertising-is-ruining-the-internet-and-breaking-the-world>
- 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*, 17 October. Available at: <http://libproxy.lib.unc.edu/login?url=https://www.proquest.com/trade-journals/apple-s-privacy-changes-create-windfall-own/docview/2597904702/se-2>
- McGuigan L (2023) *Selling the American People: Advertising, Optimization, and the Origins of Adtech*. Cambridge, MA: MIT Press.
- McGuigan L, West SM, Sivan-Sevilla I, et al. (2023) The after party: cynical resignation in adtech’s pivot to privacy. *Big Data & Society* 10: 205395172312036.
- McNealy J (2022) Platforms as phish farms: deceptive social engineering at scale. *New Media & Society* 24(7): 1677–1694.
- Meehan ER (2005) *Why TV Is Not Our Fault*. Lanham, MD: Rowman & Littlefield.
- Mizarhi-Borohovich I, Newman A and Sivan-Sevilla I (2023) The civic transformation of data privacy implementation in Europe. *West European Politics*. Epub ahead of print 15 March. DOI: 10.1080/01402382.2023.2184108.
- Mosco V (2009) *The Political Economy of Communication*. 2nd ed. Thousand Oaks, CA: Sage.
- Mudd G (2021) Privacy-enhancing technologies and building for the future. In: Facebook for Business, 11 August. Available at: <https://www.facebook.com/business/news/building-for-the-future>
- Nalpas M, Dutton S and White A (2023) Attribution reporting. *Chrome Developers*, 20 January. Available at: <https://developer.chrome.com/docs/privacy-sandbox/attribution-reporting/> (accessed 2 February 2023).

- Napoli PM (2019) *Social Media and the Public Interest*. New York: Columbia University Press.
- 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 University Press.
- Pickard V (2019) *Democracy Without Journalism? Confronting the Misinformation Society*. New York: Oxford University Press.
- Savage B, Taubeneck E and Thomson M (n.d.) A non-technical introduction to Interoperable Private Attribution (IPA). Available at: https://docs.google.com/presentation/d/1NpQz0Wm73eEKw24V7B0yCjq4Tw2qPgeezhMfS0-P-TY/edit#slide=id.gf172a1733b_0_251
- Scharlach R, Hallinan B and Shifman L (2023) Governing principles: articulating values in social media platform policies. *New Media & Society*. Epub ahead of print 7 March. DOI: 10.1177/14614448231156580.
- Smith H (2019) People-based marketing and the cultural economies of attribution metrics. *Journal of Cultural Economy* 12(3): 201–214.
- Smythe DW (1981) *Dependency Road: Communications, Capitalism, Consciousness, and Canada*. Norwood, NJ: Ablex.
- Solove DJ (2013) Introduction: privacy self-management and the consent dilemma. *Harvard Law Review* 126: 1880–1903.
- Srinivasan D (2020) Why Google dominates advertising markets. *Stanford Technology Law Review* 24: 56–175.
- Taubeneck E, Savage B and Thomson M (2022a) Interoperable Private Attribution (IPA). Available at: <https://docs.google.com/document/d/1KpdSKD8-Rn0bWPTu4UtK54ks0yv2j-22pA5SrAD9av4s/edit>
- Taubeneck E, Thomson M, Savage B, et al. (2022b) IPA end to end protocol. *GitHub*, 28 July. Available at: <https://github.com/patcg-individual-drafts/ipa/blob/main/IPA-End-to-End.md#ipa-end-to-end-protocol> (accessed 2 February 2023).
- Thomson M (2022) Privacy preserving attribution of advertising. In: Mozilla, 8 February. Available at: <https://blog.mozilla.org/en/mozilla/privacy-preserving-attribution-for-advertising/> (accessed 2 February 2023).
- Turow J (2011) *The Daily You*. New Haven, CT: Yale University Press.
- 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): 20539517211025061.
- Veale M (2022) Adtech's new clothes might redefine privacy more than they reform profiling. *Netzpolitik.org*, 2 February. 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 FZ (2022) Adtech and real-time bidding under European data protection law. *German Law Journal* 23(2): 226–256.
- Viljoen S (2021) A relational theory of data governance. *Yale Law Journal* 131: 573–654.
- West SM (2019) Data capitalism: redefining the logics of surveillance and privacy. *Business & Society* 58: 20–41.
- Wilander J (2021) Introducing private click measurement, PCM. In: WebKit, 1 February. Available at: <https://webkit.org/blog/11529/introducing-private-click-measurement-pcm/>
- 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.
- Zuboff S (2019) *The Age of Surveillance Capitalism*. New York: Public Affairs.

Author biographies

Lee McGuigan is an Assistant Professor in the Hussman School of Journalism and Media at the University of North Carolina at Chapel Hill. He is the author of *Selling the American People: Advertising, Optimization, and the Origins of Adtech* (MIT Press, 2023).

Ido Sivan-Sevilla is an Assistant Professor of information studies at the University Maryland, where he bridges computer science and public policy by developing measurements and theories of policy implementation across a range of cybersecurity, privacy, and machine learning issues.

Patrick Parham is a Ph.D. student at the College of Information Studies, University of Maryland. He has been studying advertising and media technology, and proposals addressing the deprecation of third-party cookies. Patrick previously worked in the programmatic advertising industry.

Yan Shvartzshnaider is an Assistant Professor and the director of the Privacy Rhythm research lab in the Department of Electrical Engineering and Computer Science, Lassonde School of Engineering at York University. His research focuses on sociotechnical systems that incorporate a socially meaningful conception of privacy which meets peoples' expectations and is ethically defensible. His work addresses the fundamental mismatch between programmable privacy frameworks and the ever-shifting privacy expectations of computer system users.