Can Privacy Technologies Replace Cookies? Ad Revenue in a Field Experiment

Zhengrong Gu, Garrett A. Johnson, & Shunto J. Kobayashi *

June 6, 2025

Abstract

As regulators seek to balance user privacy with publisher sustainability, Google's Privacy Sandbox offers a potential replacement for third-party cookies. This paper presents the first independent estimates of the publisher-side economic effects of Privacy Sandbox, a suite of privacy-enhancing technologies for online advertising. Leveraging an open, industry-wide field experiment, we partner with a major ad management firm to evaluate over 200 million ad impressions across more than 5,000 publishers. We find that removing third-party cookies reduces publisher revenue by 29.1%, while Privacy Sandbox preserves just 4.2% of this lost revenue. We further document that Privacy Sandbox increases ad latency and reduces impression delivery by 2.9%. The results underscore the continued economic tension between privacy and the funding of online content.

Keywords: privacy, online advertising, privacy-enhancing technologies, online content

*Questrom School of Business, Boston University: Gu <gzr@bu.edu>; Johnson <garjoh@bu.edu>; and Kobayashi <sjkobaya@bu.edu>. Authors are listed in alphabetical order. We thank our industry partner Raptive including Paul Bannister, Patrick McCann, and Joshua Zhang for their assistance with the data and our research. We thank the Center for Industry Self-Regulation for financial support.
Disclosures: Prof. Johnson served as an outside expert for the Competition and Markets Authority (CMA) in the United Kingdom for its Privacy Sandbox case. This research was conducted in his personal capacity and does not reflect the views of the CMA.

1 Introduction

Digital advertising is the economic engine behind the internet that funds online content creation and journalism. Third-party cookies have long enabled digital advertising's essential functions—namely targeting and measurement—and increased revenue for both publishers and advertisers (CMA 2020a; Wernerfelt et al., 2025). Yet, critics raise concerns about user privacy and liken pervasive tracking to corporate surveillance (Zuboff, 2019). In response, regulations like Europe's General Data Protection Regulation (GDPR) sought to empower consumer choice while browsers like Safari and Firefox block third-party cookies altogether. These privacy gains come at a cost: publishers collect less ad revenue while consent dialogs clutter the web. The dilemma posed by competing privacy and economic interests calls for new solutions.

Privacy-Enhancing Technologies (PETs) balance these competing interests and are gaining traction among academics, firms, and regulators (Gadotti et al., 2024; ICO 2022). Examples include differential privacy, which adds statistical noise to data, and on-device computation, which processes user data locally to prevent data leakage. PETs have been applied in government statistics (Kenny et al., 2024), medical research (Warnat-Herresthal et al., 2021; Zhou et al., 2024), social science research (Evans & King, 2023) and big tech firms (Hard et al., 2018). Though such PET applications are often narrow in scope, the ad industry is now attempting to replace third-party cookies with PETs that underpin a variety of firms and industry use cases. Examples include Apple's AdAttributionKit, Firefox's Privacy-Preserving Attribution, and Microsoft's Ad Selection API.

Among these efforts, Google's Privacy Sandbox stands out for its scope and potential impact, given Chrome's dominant browser share. It comprises PETs for ad targeting (Topics, Protected Audience) and ad measurement (Attribution Reporting) that are available on Chrome browsers since 2023. Though Chrome abandoned plans to deprecate cookies, Google continues to invest in the Privacy Sandbox, which services the over 20% of Chrome traffic that already lacks cookies.

A key concern among policymakers is whether Google's approach will further entrench its dominance (Alcobendas et al., 2023). The UK's Competition and Markets Authority (CMA 2022) is investigating Privacy Sandbox on the grounds that it could distort competition and enable self-preferencing. These concerns are magnified by Chrome's two-thirds share of global browser traffic

(CMA 2022) and Google's substantial market presence across the ad tech stack (CMA 2020b). At the same time, regulators recognize the promise of privacy-enhancing technologies in improving legal compliance and reducing harm (ICO 2022).

To inform this policy debate, the CMA and Google designed an open, industry-wide field experiment. This experiment assigns millions of Chrome users globally into one of three browser environments: (1) the status quo with third-party cookies, (2) Privacy Sandbox with cookies deactivated, and (3) a fully cookieless environment. This design allows for a clean comparison of the economic impact of cookie deprecation and the efficacy of Privacy Sandbox as a potential replacement. Importantly, the experiment is open: browser-facing firms can observe each user's treatment group and can analyze the experimental outcomes. This allays concerns that big tech firms only release flattering research (Bak-Coleman et al., 2023).

We partner with Raptive—the largest ad management firm—to evaluate publisher outcomes within this experiment. Raptive represents over 5,000 publishers and fully manages their ad monetization. This enables us to observe comprehensive publisher-side outcomes across more than 200 million ad impressions served to Chrome users globally in August 2024.

Our findings reveal that third-party cookies materially increase publisher revenue. Removing cookies reduces revenue by 29.1% relative to the status quo, and by 35.3% for the subset of impressions that have third-party cookies. These estimates are somewhat lower than those in prior studies (Johnson et al., 2020; Alcobendas et al., 2023; Ravichandran & Korula, 2019), likely because Raptive implements cookie-replacement identifiers that partially offset the loss.

Privacy Sandbox preserves only a modest share of cookie-enabled revenue. Publisher revenue in the Sandbox condition falls 27.9% relative to the status quo—recovering just 4.2% of the loss from deprecating cookies. Moreover, we find longer ad latency under Privacy Sandbox, with the 75th percentile exceeding 5.6 seconds—more than double that of other conditions. This increased latency reduces impressions served by 2.9% in the Sandbox condition, undercutting the modest revenue gains per impression. While Sandbox performance is weak, its limited industry adoption (Johnson, 2024a) likely constrains its effectiveness.

We contribute to several research streams. First, we add to the emerging academic literature on Privacy Sandbox. Kobayashi et al. (2024) is closest to our work: they evaluate advertiser performance within the same open experiment and find more encouraging results. Several firms

participated in the experiment and shared their findings (Google, 2024c,a; Index Exchange product team, 2024; Pascoe, 2024; Selman, 2024). Alcobendas et al. (2023) and Sousa (2024) instead take a structural approach to modeling Privacy Sandbox's impact. Jerath & Miller (2023) explore consumer perceptions of privacy-enhanced advertising. Johnson (2024a) documents the adoption of Privacy Sandbox across websites and vendors, and Johnson et al. (2022) reviews privacy-centric online advertising more broadly.

Second, we contribute to the literature on the economics of privacy (e.g., Acquisti et al., 2016; Goldfarb & Que, 2023). Our work complements prior studies on the value of cross-site identifiers for publishers (Johnson et al., 2020; Alcobendas et al., 2023; Miller & Skiera, 2024) and advertisers (Aridor et al., 2024; Goldfarb & Tucker, 2011; Wernerfelt et al., 2025). We extend this literature by providing the first publisher-side evidence from a randomized, industry-wide field experiment. Related research also documents the impact of privacy regulation on content creation (Johnson et al., 2023; Kircher & Foerderer, 2024; Lefrere et al., 2022). The economics literature on privacy-enhancing technologies has focused on differential privacy (e.g., Hotz et al., 2022; Komarova & Nekipelov, 2020; Steed et al., 2022), while Privacy Sandbox involves a broader set of privacy-enhancing technologies.

Third, our study contributes to the literature on self-regulation. Privacy Sandbox represents an unusual case, in which a dominant platform preemptively deploys technical standards under regulatory oversight. Prior work shows that self-regulation can arise in response to anticipated regulation (Maxwell et al., 2000), and that oversight can enhance its effectiveness (DeMarzo et al., 2005). Self-regulation can also encourage private investment in public goods such as privacy-enhancing technologies (Baron, 2010). While privacy self-regulation generally employs a notice-and-choice regime (Acquisti et al., 2016), Privacy Sandbox shifts compliance into the browser architecture itself. Privacy Sandbox is also unusual in that it is led by a single firm rather than a self-regulatory organization (DeMarzo et al., 2005).

This article is organized as follows. The next section provides background on our data provider and Privacy Sandbox. Section 3 explains the design of the experiment and Section 4 describes our data. Section 5 explains our estimation methodology and Section 6 presents our results. Section 7 places our results in context and Section 8 concludes.

2 Background

Online display advertising relies on real-time auctions to allocate individual ad impressions—instances in which an ad is shown to a user. These auctions are technically complex and require coordination across multiple intermediaries on both the supply and demand sides. As a result, many publishers delegate ad-selling functions to ad management companies, which handle the entire process in exchange for a commission on revenue.

Our data partner, Raptive, is the largest ad management company in the online display ecosystem. As of 2024, Raptive reports serving over 5,200 creators and reaching 191 million monthly site visitors, with more than $2.5 billion paid out to its partners to date. Collectively, Raptive ranks among the top ten ComScore media properties overall, and leads in content categories such as Food, Home, Lifestyle, and Family. Market analysts estimate that ad management companies account for 25.7% of web bid requests, up from 11.8% in 2022, with Raptive alone representing 43.1% of those requests (Jounce Media, 2024). Our own analysis of Sincera data shows that Raptive has a 11.6% market share by site count—third among its peers—but is skewed toward higher-traffic publishers. The company's client roster spans a range of publishers from Feel Good Foodie and The Woks of Life to DJ Jazzy Jeff and MacRumors. News category content represents about 3% of Raptive's ad volume.

The transition away from third-party cookies marks a foundational change in digital advertising. While Safari and Firefox disabled cookies by default in 2018 and 2019, respectively. Chrome announced plans to follow suit in 2019 but reversed course in April 2025. Despite this reversal, an estimated 36% of global browser traffic already lacks third-party cookies, as a substantial share of Chrome traffic operates without them.[1] Third-party cookies support cross-site tracking, which enables behavioral targeting, retargeting, and ad measurement. Beyond advertising, cookies also support user authentication, content personalization, and fraud prevention. In response to growing privacy concerns, Google introduced Privacy Sandbox—a suite of privacy-preserving technologies intended to substitute for cookies while retaining core advertising functionality.

[1] Chrome traffic can lack third-party cookies due to user actions such as site-level opt-outs, clearing cookie caches, or enabling browser-level blocking. In our data, 21.8% of impressions lack cookies—though this is a conservative estimate, as it reflects only users who have not disabled cookies in their browser settings. Based on global browser market shares from StatCounter (August 2024), we estimate that at least 36% of traffic is cookieless—comprising all Firefox (2.74%), Safari (18.57%), and 21.8% of Chrome (65.20%) traffic. See: https://gs.statcounter.com/browser-market-share.

Privacy Sandbox offers technologies for targeting (Topics, Protected Audience), measurement (Attribution Reporting), and fraud prevention (Private State Tokens). These tools leverage privacy-enhancing techniques such as on-device processing, differential privacy, and zero-knowledge proofs. For example, the Topics API provides coarse-grained interest-based targeting based on a user's recent browsing activity, while the Protected Audience API enables on-device auctions for behavioral advertising and retargeting. The Attribution Reporting API offers privacy-preserving conversion measurement by adding noise.

Johnson (2024a) documents modest adoption of these APIs in our August 2024 sample period: 27.8% of sites called the Topics API, 16.6% ran Protected Audience auctions, and 13.3% added users to Protected Audience interest groups to show targeted ads off-site. Google plays a central role in deploying Sandbox technologies. In August 2024, Google's market share was 21.5% in Topics API calls, 38.3% in Protected Audience interest group creation, and 88.7% in Protected Audience auctions. Raptive also participated and ranked fourth in interest group creation with a 7.9% share.

Google announced the Privacy Sandbox initiative in August 2019 alongside its intent to deprecate third-party cookies in Chrome. In 2021, the UK's Competition and Markets Authority (CMA) launched an investigation and took on oversight of the initiative, citing Chrome's browser dominance and Google's centrality in adtech. Chrome began testing the APIs in 2021 and made them generally available in 2023. In January 2024, Google disabled third-party cookies for 1% of users (approximately 30 million users) as part of the field experiment we study. In July 2024, Google postponed full deprecation, instead proposing a browser-level consent model. In April 2025, Google abandoned its deprecation plan altogether, citing feedback from industry and regulators, but reaffirmed its commitment to continued investment in Privacy Sandbox.

Privacy Sandbox leads a broader industry shift towards privacy-enhancing ad technologies. Microsoft is piloting a targeting solution in its Edge browser (Ad Selection API) modeled on the Protected Audience API.[2] The World Wide Web Consortium (W3C), spurred by a joint proposal from Firefox and Meta,[3] is advancing a browser-based measurement standard. Apple continues to evolve its own measurement tool, AdAttributionKit. Among these efforts, Privacy Sandbox re-

[2]Source: https://learn.microsoft.com/en-us/microsoft-edge/web-platform/ad-selection-api.
[3]Source: https://blog.mozilla.org/en/mozilla/privacy-preserving-attribution-for-advertising.

mains the most comprehensive in scope and functionality, with the broadest adoption of privacy-preserving technologies to date.

Raptive engages with Privacy Sandbox both as a publisher and as a vendor, interacting with Sandbox APIs through Google's ad server (Google Ad Manager). Furthermore, Raptive contributed a public analysis of the Topics API (Marti, 2022). Raptive's early adoption and public engagement position it as a key participant in the ecosystem's transition to Privacy Sandbox.

3 Experimental design

The Privacy Sandbox field experiment is a collaboration between Google and CMA, which is the lead regulator overseeing Google's browser privacy changes. Launched in January 2024, the experiment encompasses approximately 60 million Chrome users—2% of eligible users—across desktop and Android mobile devices (CMA 2023b; 2023a). Users are randomly assigned to one of three groups: (1) 1% in Status Quo (third-party cookies enabled), (2) 0.75% in Privacy Sandbox (Sandbox APIs active and cookies disabled), and (3) 0.25% in Cookieless (both APIs and cookies disabled).[4] This design facilitates comparisons between Privacy Sandbox and two key benchmarks: the current ecosystem and a cookieless environment. Moreover, the comparison between Status Quo and Cookieless conditions identifies the effect of cookie removal. Appendix A provides balance checks.

Assignment to experimental groups depends on several eligibility criteria. Users must run a recent version of Chrome that supports Privacy Sandbox and the experimental labels, and cannot be in incognito mode or have manually disabled cookies. The design also excludes both Chrome Enterprise users and iOS users. While Chrome enforces assignment via technical implementation, users retain partial control. They may override cookie settings or disable Privacy Sandbox. Those in Sandbox and Cookieless groups receive a one-time notification about enhanced tracking protection but are otherwise unaware of their assignment.

The design is notable in its scale and openness. Browser-level randomization and open labeling facilitates coordinated industry participation in the experiment, involving both supply and

[4]Google and the CMA instead refer to these groups as control_1, treatment_1, and control_2 respectively. For additional technical explanations on the testing labels, see https://developers.google.com/privacy-sandbox/private-advertising/setup/web/chrome-facilitated-testing.

demand sides of the display ad market. This enables general equilibrium evaluation across many firms over several months and allows browser-facing firms to conduct independent analyses. The CMA requested interim results from participants in June 2024 and many firms submitted their findings, including Google. Nevertheless, experimental labels remain accessible through at least the first half of 2025.

We exploit this open experiment in partnership with Raptive to evaluate the revenue impact of Privacy Sandbox for thousands of publishers. Raptive continues its standard ad operations across all experimental arms. In the Status Quo group, Raptive transmits third-party cookie identifiers to buyers. In the Sandbox group, it uses the Protected Audience API for serving ads, while buyers may also access the Topics and Attribution Reporting APIs alongside conventional ad calls. Notably, Raptive deploys alternative identifiers across all groups, including deterministic identifiers based on hashed email addresses and probabilistic identifiers based on device and network characteristics in all experimental groups. These alternatives help attenuate, but do not fully offset, the loss of cookies due to limited industry adoption. In all conditions, buyers can also rely on contextual signals like page URL.

Raptive is well positioned to assess the impact of Privacy Sandbox. Raptive directly interfaces with users and can directly observe browser-assigned group labels without relying on upstream intermediaries, avoiding data gaps that would otherwise bias inference from the experiment. By capturing all impressions served and revenue earned across conditions, Raptive offers a comprehensive and internally valid estimate of publisher-side effects.

4 Data

We analyze impression-level data from Raptive covering all 207.2 million ads served to experimental users in August 2024. These data span over 5,000 publishers across 44 content categories—including entertainment, food, gaming, and news. The data include impressions served to users worldwide on desktop, smartphones, and tablets. This scope helps mitigate concerns about selection and strengthens the external validity of our findings.

We impose two additional data filters beyond those discussed in the experimental design. First, we focus on August 2024, when Google Ad Manager (GAM) began fully reporting Protected

Table 1: Summary Statistics (Sold Impressions)

Variable	Observations	Mean	S.D.	Min.	25%	Median	75%	Max.
All impressions								
Scaled ad revenue[a]	207,157,196	0.382	1.196	0	0.043	0.236	0.521	1328.571
Ad latency (sec.)[b]	207,157,196	7.370	24.196	0	0.639	1.131	3.288	188.682
Desktop	207,157,196	0.408	0.491	0	0	0	1	1
Smartphone	207,157,196	0.568	0.495	0	0	1	1	1
Tablet	207,157,196	0.024	0.153	0	0	0	0	1
EU	207,157,196	0.067	0.250	0	0	0	0	1
US	207,157,196	0.609	0.488	0	0	1	1	1
Remaining countries	207,157,196	0.324	0.468	0	0	0	1	1
Status Quo group								
Scaled ad revenue[a]	104,510,637	0.440	1.196	0	0.057	0.307	0.586	1328.571
Ad latency (sec.)[b]	104,510,637	6.720	23.653	0	0.6187	0.947	2.244	188.682
Cookie indicator	104,510,637	0.781	0.413	0	1	1	1	1
Privacy Sandbox group								
Scaled ad revenue[a]	76,283,492	0.327	1.188	0	0.036	0.186	0.450	1328.571
Ad latency (sec.)[b]	76,283,492	8.4579	25.000	0	0.784	1.821	5.641	188.682
PA API sold indicator	76,283,492	0.072	0.2658	0	0	0	0	1
Topics category returned[c]	76,283,492	0.271	0.444	0	0	0	1	1
Cookieless group								
Scaled ad revenue[a]	26,363,067	0.312	1.211	0	0.029	0.174	0.436	1328.571
Ad latency (sec.)[b]	26,363,067	6.703	23.830	0	0.547	0.821	2.082	188.682

Notes: The unit of observation is a sold impression, so all variables are conditional on sale. [a]Ad revenue is scaled by a factor between 0.1 and 1 to preserve confidentiality. [b]Ad latency is measured in seconds and winsorized at the 99th percentile to mitigate the influence of outliers. [c]Topics category returned is an indicator for the Topics API returning a behavioral category label when Raptive called the API.

Audience API (PA API) revenue to Raptive.[5] Second, we restrict attention to banner ads and exclude native and video formats, which were not yet supported by PA API.

Table 1 presents summary statistics overall and across experimental groups. These include scaled ad revenue, ad latency, and device and geography shares. Impressions are distributed across desktop (40.8%), mobile (56.8%), and tablet (2.4%) devices. Most impressions are from US users (60.9%), with 6.7% from the EU.

Ad revenue is scaled for confidentiality: expressed as cost per thousand impressions in U.S. dollars and scaled by a number between 0.1 and 1. The Status Quo group shows the highest average scaled revenue ($0.44), followed by Privacy Sandbox ($0.33) and Cookieless ($0.31). These provide preliminary evidence that cookies and Privacy Sandbox add value.

[5]Prior to August 2024, GAM excluded PA API impressions where a non-Google seller won the meta-auction, so the data was incomplete: see https://support.google.com/admanager/answer/14639079.

Ad latency, defined as the time between an ad request and its delivery, influences both user experience and monetization. Median latency is longest in the Privacy Sandbox group at 1.82 seconds—about double the latency in the Status Quo (0.95s) and Cookieless (0.82s) groups. Longer latency can reduce the likelihood of a successful auction before page load. However, all statistics in Table 1 are conditional on sold impressions. We address the potential impact of unsold impressions on monetization in Sections 5 and 6.

Some variables are treatment-specific by design. Third-party cookie use is exclusive to the Status Quo group, while PA API and Topics API data are unique to the Sandbox group. Table 1 shows that only 78.1% of Status Quo impressions contain third-party cookies. Eligible Chrome users in the experiment permit cookies at the browser level, but may lack cookies because of site-level consent, cache clearing, or browser extensions.

In the Sandbox group, we observe limited adoption of Privacy Sandbox technologies. Table 1 shows that 7.2% of impressions are sold via PA API, as the PA API auction must compete with the conventional auction. For Topics API, we obtain an indicator for whether Raptive observes a behavioral category label for a given impression. Vendors like Raptive can query the Topics API for a user's interest categories, but the browser only returns a label if that vendor has previously encountered the user on a site associated with that category. This design prevents vendors from gaining access to new user information that they would not already have under the status quo. Raptive observes Topics API category data for 27.1% of impressions in the Sandbox group. Google is present on Raptive-affiliate websites as well as a wider network of sites (Johnson, 2024a), so 27.1% represents a lower bound on how often Google observes Topics API category data. We do not observe Attribution Reporting API or Trust Token API usage, although these may be functioning in the background.

In sum, these data offer a rich basis for analyzing how Privacy Sandbox technologies affect publisher revenue and ad delivery performance.

5 Estimation methodology

We use the experimental design to estimate the intent-to-treat (ITT) effects of two key interventions: (1) removing third-party cookies and (2) replacing cookies with Privacy Sandbox technolo-

gies. Let i denote the impression, $D \in \{SQ, PS, CL\}$ denote experimental assignment, and y denote the outcome variable (either ad revenue or ad latency). The ITT effects respectively are defined as:

$$ITT_{CL} = E[y_i|D_i = CL] - E[y_i|D_i = SQ] \tag{1}$$

$$ITT_{PS} = E[y_i|D_i = PS] - E[y_i|D_i = SQ] \tag{2}$$

To assess the degree to which Privacy Sandbox mitigates revenue loss from cookie removal, we compute what we term the *recovery share*:

$$\text{Recovery share} = \frac{E[y_i|D_i = PS] - E[y_i|D_i = CL]}{E[y_i|D_i = SQ] - E[y_i|D_i = CL]} \tag{3}$$

As we saw in Section 4, the Status Quo group includes both cookied and cookieless impressions. Since only the cookied impressions generate differential value relative to the Cookieless group, the ITT estimate for cookie removal understates the effect on cookied traffic. We therefore estimate the effect of removing cookies for the subset of cookied impressions (i.e., the average treatment effect on the treated or ATET) using:

$$ATET_{CL} = \frac{ITT_{CL}}{\Pr\left[Cookie_i = 1|D_i = SQ\right]} \tag{4}$$

where $Cookie_i$ is an indicator for whether impression i has a third-party cookie (see e.g., Manski, 2007). Thus, the ATET scales the ITT by the share of cookied impressions in the Status Quo group. Note that we do not estimate an analogous ATET for Privacy Sandbox, since the technologies are designed to apply regardless of cookie availability.

While most quantities in the above expressions can be directly estimated using sample counterparts, missing data in the Privacy Sandbox condition presents a challenge. Longer ad latency—especially under the Protected Audience API—may cause users to navigate away before the ad loads. In these cases, no ad is served and no revenue is generated. Our dataset contains only sold impressions, so these instances represent missing outcomes. To address this, we use the law of

iterated expectations:

$$\begin{aligned} E\left[y_i|D_i = PS\right] =& E[y_i|D_i = PS, sold_i] \cdot \Pr\left[sold_i|D_i = PS\right] \\ &+ E[y_i|D_i = PS, unsold_i] \cdot (1 - \Pr\left[sold_i|D_i = PS\right]) \end{aligned}$$

Since unsold impressions generate zero revenue by definition, the revenue outcome simplifies to:

$$E\left[y_i|D_i = PS\right] = E[y_i|D_i = PS, sold_i] \cdot \Pr\left[sold_i|D_i = PS\right] \tag{5}$$

We estimate the Sandbox sale probability ($\Pr\left[sold_i|D_i = PS\right]$) by leveraging the known randomization probabilities and assuming all impressions are observed in the Status Quo and Cookieless groups. Given the experimental design—1% of users assigned to Status Quo, 0.25% to Cookieless, and 0.75% to Privacy Sandbox—we expect the Privacy Sandbox group to have $\frac{3}{5}$ as many total impression opportunities as the combined Status Quo and Cookieless groups $\left(\widehat{N^0_{PS}} = \frac{3}{5}\left(N_{SQ} + N_{CL}\right)\right)$. Thus, we estimate the Sandbox sale probability as the ratio of the total (observed) impressions (N_{PS}) and the total (implied) impression opportunities $\left(\widehat{N^0_{PS}}\right)$:

$$\widehat{\Pr}\left[sold_i|D_i = PS\right] = \frac{N_{PS}}{\widehat{N^0_{PS}}} = \frac{N_{PS}}{\frac{3}{5}\left(N_{SQ} + N_{CL}\right)} \tag{6}$$

This adjustment allows us to estimate the average scaled ad revenue in the Sandbox group, accounting for impressions lost to latency. The resulting estimates provide a more accurate comparison of monetization performance across treatment groups.

6 Results

We evaluate the impact of third-party cookie deprecation on publishers and assess whether Privacy Sandbox mitigates the resulting revenue loss. We begin by analyzing how Sandbox affects ad latency and the share of impressions sold, and then estimate average revenue after adjusting for unsold impressions. Additional heterogeneity analyses are provided in Appendix B.

Privacy Sandbox relies on on-device computation in the Protected Audience API to enhance privacy, but this can increase ad latency. Figure 1 shows the empirical cumulative distribution

functions (CDFs) for ad latency in the experiment.[6] Figure 1a shows that ad latency is consistently higher in the Sandbox group. The 75th percentile latency is 5.64 seconds in the Sandbox condition, more than twice that of the Status Quo (2.24 seconds) and Cookieless (2.08 seconds) groups. Appendix B.1 shows that this latency effect is more pronounced on mobile devices, likely due to more limited processing power.

To isolate the effect of PA API, Figure 1b separates Sandbox impressions based on whether they were sold via PA API. Latency remains elevated even for impressions not sold through PA API (75th percentile: 5.47 seconds), while impressions sold via PA API exhibit even higher delays (75th percentile: 7.89 seconds). These results suggest that PA API auctions contribute to latency regardless of sale because they compete alongside the traditional ad auctions.

Longer latency reduces the chance that an impression is served before the user leaves the page. While our dataset includes only sold impressions, we estimate the share of missed opportunities using the experiment's randomization structure and group sizes. As detailed in Section 5, we use Equation (6) to estimate that 2.85% of impression opportunities go unsold in the Privacy Sandbox group relative to expectations under random assignment.[7]

We next turn to ad revenue. Table 2 summarizes average scaled ad revenue per *sold* impression: 0.440 in the Status Quo, 0.327 in the Sandbox, and 0.312 in the Cookieless condition. However, because our data only include sold impressions, the observed revenue in Sandbox is upwardly biased due to the exclusion of zero-revenue unsold impressions.

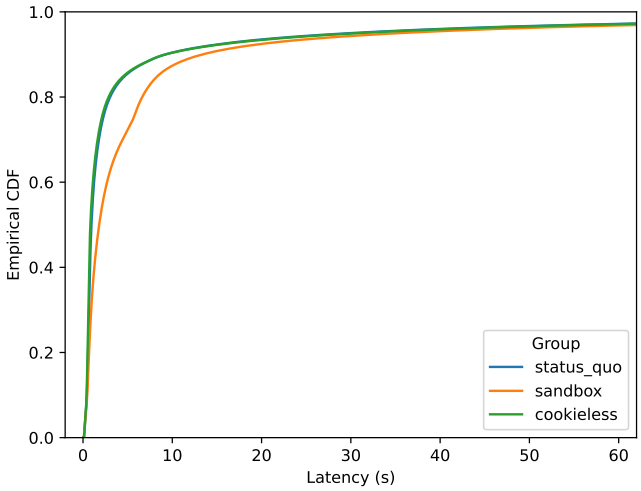
Using Equation (5) from Section 5, we adjust the Sandbox revenue to account for the estimated 2.85% of unsold impressions. Given that the mean revenue per sold impression is 0.327 and the estimated sale probability is 97.15%, the adjusted revenue across all impression opportunities in the Sandbox condition is 0.317. Note that the means in Table 2 are estimated precisely (i.e., standard errors on the order of 10^{-4}) due to our large sample size.

In Figure 2, we normalize the mean scaled revenue for each group against the Status Quo

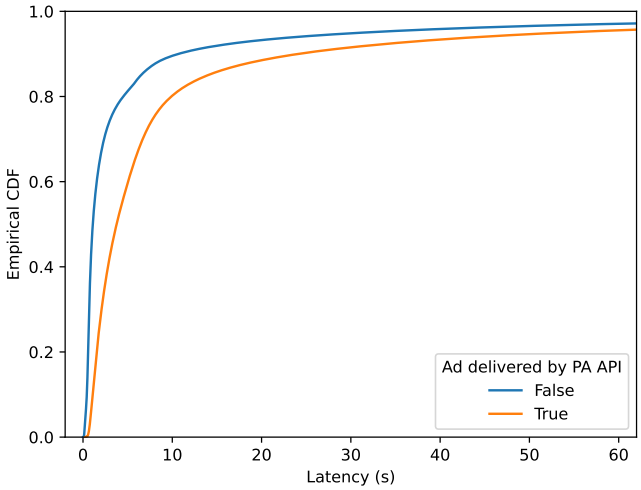
[6]Note that latency here is conditional on sold impressions. We do not adjust for the unobserved ad latency among unsold impressions in the Sandbox group. However, we expect that the latency distribution for unsold impressions must strictly dominate of the latency in the Status Quo and Cookieless group in order for the Sandbox impressions to time out and go unsold. As such, are our qualitative findings are unchanged.

[7]Substituting our group-level impression counts from Table 1 into Equation (6), we have:

$$\widehat{\Pr}\left[sold_i | D_i = PS\right] = \frac{N_{PS}}{\frac{3}{5}\left(N_{SQ} + N_{CL}\right)} = \frac{76,283,492}{\frac{3}{5}\left(104,510,637 + 26,363,067\right)} = 97.15\%$$



(a) Ad Latency by Experimental Group



(b) Sandbox Group: Ad Latency by PA API Delivery

Figure 1: Empirical CDFs for Ad Latency

Table 2: Ad Revenue by Experimental Group (Adjusted for Unsold Impressions)

	Mean estimate	Standard error
Status Quo		
Scaled ad revenue	0.440	1.17×10^{-4}
Cookieless		
Scaled ad revenue	0.312	2.36×10^{-4}
Privacy Sandbox		
Scaled ad revenue (conditional on sold)	0.327	1.36×10^{-4}
Unsold impression share	2.85%	-
Adjusted ad revenue (all impression opportunities)	0.317	1.67×10^{-4}

Note: Ad revenue is scaled by a factor between 0.1 and 1 to preserve confidentiality. For the Sandbox group, adjusted ad revenue accounts for unsold impressions using Equation (5) and the share of unsold impressions. Asymptotic standard errors given for means.

mean scaled ad revenue to enable relative comparisons. The results show that the Sandbox group generates 72.12% of the Status Quo revenue, while the Cookieless group achieves 70.91%. Note that both publishers and Raptive see equal relative revenue reductions given that ad management companies earn a percentage commission on revenue.

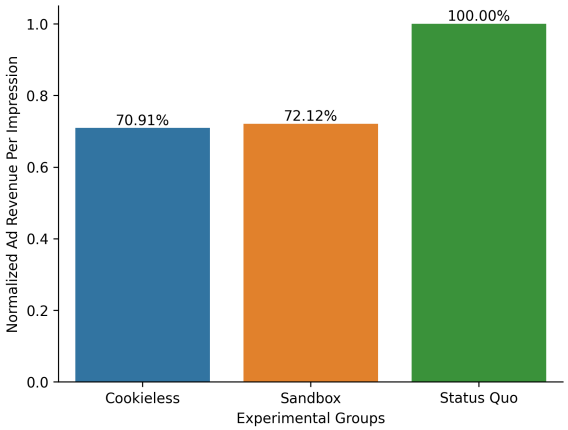
Table 3 summarizes our key estimated effects.. The intent-to-treat (ITT) estimate for cookie removal—the difference between Cookieless and Status Quo revenue—is a 29.09% decline. Because only 78.09% of impressions in the Status Quo group have cookies, we calculate the effect of removing cookies from the subset of cookied impressions in the Status Quo group. From Equation (4), the corresponding average treatment effect on the treated (ATET) is -35.29%.[8]

For the Sandbox group, the ITT revenue effect is a 27.88% decline relative to Status Quo, which is only a 1.21% improvement over the Cookieless group. This implies that Privacy Sandbox recovers just 4.18% ($\frac{1.21\%}{27.88\%}$, see Equation 3) of the revenue lost due to cookie deprecation.

Appendix B.2 explores heterogeneity by publisher size and country. We find competitive consequences from deprecating cookies in that smaller publishers show greater revenue loss. We also demonstrate that publisher revenue from EU users experiences a greater impact from cookie loss (ITT of -49.79%), despite having a lower cookied impression share (56.61%). This occurs because

[8]In absolute terms, the corresponding ITT estimate for scaled revenue is -0.128. Given Equation (4), the absolute ATET is -0.128/78.09%, as 78.09% of Status Quo impressions have cookies. To compute the ATET in relative terms, we further normalize by the average scaled ad revenue among Status Quo cookied impressions (0.465), yielding -35.29%.

Figure 2: Normalized Mean Ad Revenue by Experimental Group



Note: The normalized ad revenues are calculated by dividing each experimental group's mean scaled ad revenue by that of the Status Quo in Table 2. For the Sandbox group, we use the mean ad revenue adjusted to account for unsold impressions.

Table 3: Treatment Effect Estimates

	Estimate	Standard Error
Effect of Removing Cookies		
Intent to Treat	-29.09%	0.057%
Share of impressions with cookies (Status Quo)	78.09%	0.0040%
Average Treatment Effect on the Treated	-35.29%	0.062%
Effect of Privacy Sandbox		
Intent to Treat	-27.88%	0.036%
Recovery Share	4.18%	0.20%

Note: The Intent-to-Treat effects are calculated using Equations (1) and (2), and are normalized by the estimated mean ad revenue of the Status Quo group. The Sandbox ITT accounts for unsold impressions in the Sandbox group. The Average Treatment Effect on the Treated is obtained using Equation (4) and normalized by the estimated mean ad revenue of *cookied* impressions in the Status Quo group. The recovery share is computed using Equation (3). Standard errors are heteroskedasticity-robust. For the ATET, robust standard errors are computed using two-stage least squares. The delta method is used to account for normalization in all estimated effects.

cookies deliver even more value per cookied impression in the EU market (ATET of -66.43%).

In summary, publisher revenue falls meaningfully without third-party cookies. Privacy Sandbox offers some revenue mitigation, but increases ad latency and reduces impression delivery. After adjusting for these effects, Sandbox recovers only a small fraction of the loss that occurs without cookies.

7 Discussion

7.1 Value of third-party cookies to publishers

As shown in Section 6, our intent-to-treat (ITT) estimate indicates that deactivating third-party cookies reduces publisher revenue by 29.1% relative to the Status Quo group. However, only 78.1% of impressions in the Status Quo group include third-party cookies. To account for this, we compute the average treatment effect on the treated (ATET), which shows a 35.3% revenue reduction for impressions with cookies. This effect is notably larger for European Union traffic, where the estimated ATET reaches 66.4% (see Appendix B.2).

The overall estimates are somewhat smaller than prior work, which often reports declines of around 50% (e.g,. Johnson et al., 2020; Ravichandran & Korula, 2019). Most prior studies rely on observational comparisons of cookied and cookieless impressions and assume unconfoundedness. Laub et al. (2024) find such estimates are sensitive to model specification. Structural modeling approaches yield similar conclusions using ad auction data (Alcobendas et al., 2023; Sousa, 2024). Ravichandran & Korula (2019) present experimental evidence from deactivating cookies within Google's own ad buying, though the CMA (2020a) notes this was a partial equilibrium test in that competing buyers retained access to cookies. The CMA's reanalysis estimated a potential 70% reduction in ad prices—rather than 52%—under full cookie deprecation. Our study extends this literature by estimating general equilibrium revenue effects in a field experiment where all participants lose cookies using data from over 5,000 publishers.

Our findings must be interpreted in light of widespread adoption of alternative user identifiers. These include both probabilistic identifiers based on IP address and device recognition, and deterministic identifiers such as hashed email addresses.[9] While these alternatives aim to replicate

[9]See IAB Tech Lab (2023) for further details.

the functionality of cookies, their effectiveness is more limited due to lower industry adoption. For example, effective cross-site tracking using deterministic identifiers requires that websites employ a common identity provider, and that the user logs in with a consistent email address.

Raptive makes extensive use of such alternative identifiers. On average, our data show that Raptive employs approximately 11 such identifiers per impression, with only about 3.5% of impressions lacking any identifier. Raptive has publicly stated that these identifiers boost monetization.[10] As such, our estimates reflect the incremental impact of deactivating third-party cookies in an environment that continues to leverage alternative identifiers.

If the goal is to measure the effect of eliminating all cross-site identifiers—including cookies, probabilistic IDs, and hashed emails—our estimates likely understate the full monetization loss. Both contexts are relevant for policy: partial identifier removal reflects industry responses under current enforcement norms, while complete de-identification informs debates over stricter privacy regimes. The latter is also of academic interest as it isolates the value of tracking technologies in a more principled way.

Alternative identifiers remain controversial. In principle, many privacy regulators impose the same legal restrictions on cookies and alternative identifiers—whether probabilistic or deterministic (e.g., ICO 2019). Google initially committed to blocking probabilistic identification methods with the launch of Privacy Sandbox but reversed course in 2024 by allowing its use—sparking regulatory concern. Our findings thus offer policy-relevant evidence for evaluating not only cookies but the broader tracking ecosystem.

7.2 Value of Privacy Sandbox to publishers

Our results illustrate that Privacy Sandbox has a modest impact on publisher revenue: it recovers just 3.6% of the revenue lost from removing third-party cookies. While both regulators and industry stakeholders are interested in long-run general equilibrium outcomes, our findings reflect the industry's partial adoption and adaptation as of August 2024. We view our estimates as directionally informative for policy evaluation, but necessarily incomplete.

Raptive fully adopted Privacy Sandbox across its publisher network. However, revenue out-

[10] Raptive found that one alternative identifier increased its revenue for Safari impressions by 7%-8% on Google's ad exchange whereas another identifier increased ad revenue as much as 500% on Safari and 30% on Chrome. Source: https://www.adexchanger.com/data-exchanges/cafemedia-dishes-on-alternative-ids/.

comes also depend on adoption elsewhere in the ecosystem. In our data, the Protected Audience API accounts for only 7.2% of sold impressions, and Topics API data is returned for 27.1%—a lower bound due to the API's design. This limited uptake reflects broader trends: as of mid-2024, Privacy Sandbox was adopted by a minority of websites and vendors (Johnson, 2024a).

Other firms participating in the experiment report more favorable results. Kobayashi et al. (2024) find that Privacy Sandbox retained 82–93% of advertiser outcomes on a per-dollar or per-impression basis. Criteo estimates it would need to spend 60% less under Privacy Sandbox to match Status Quo outcomes (Selman, 2024). Google (2024c) reports that ad spend fell 14% and conversions per dollar fell 32% under Sandbox. On the publisher side, Google estimates that revenue fell 21% for Ad Manager users and 18% for AdSense publishers in the Sandbox group. These compare to 34% and 22% declines, respectively, in the Cookieless group—implying modest but higher recovery shares than our own.

These findings suggest that Privacy Sandbox has *potential* viability if advertiser and platform participation increases. Greater adoption could shift more demand through Privacy Sandbox APIs and improve publisher monetization. However, the gulf between the demonstrated value of cookies and the soft performance of Privacy Sandbox may help explain Google's retreat from cookie deprecation.

Latency remains a barrier to Sandbox performance. As shown in our results and corroborated by industry experts (Martin, 2024; Selman, 2024), the Protected Audience API introduces delays that reduce impression delivery. While new approaches to combining traditional auctions with PA API show promise,[11] wider adoption could exacerbate latency issues.

The future of Privacy Sandbox remains uncertain. With Chrome's cookie deprecation no longer on the horizon, a major incentive for industry adoption has disappeared. Although Google continues to develop the initiative, its timeline is unclear. The CMA maintains regulatory oversight and has not yet approved final implementation. Google's ongoing legal challenges—including the potential divestiture of Chrome or segments of its ad business—add further complexity. Lastly, Sandbox must compete with cookies and alternative identifiers, unless privacy regulators opt to

[11] Google Ad Manager has faced criticism for running the traditional and PA API auctions in sequence rather than in parallel, as the latter may reduce latency. Google (2025) reports that PA API auction speed improved 35% in 2025 and that greater usage of parallelized should yield additional improvement. Moreover, Chrome has proposed an approach (i.e., bidding and auction services) for moving PA API auctions off user devices and onto "trusted servers", that could increase speed without writing user data to disk.

curtail user identifiers in favor of privacy-enhancing ad technologies.

8 Conclusion

This paper makes several contributions to the empirical literature on privacy and the digital economy. First, we provide the first experimental estimates of the revenue impact from deactivating third-party cookies at an industry-wide scale. Our findings demonstrate that third-party cookies continue to play a significant role in supporting publisher revenue and that their removal imposes a substantial economic cost. These results inform the tradeoff between safeguarding user privacy and sustaining the advertising revenue that funds online content, including independent creators and journalists.

Second, we offer the first independent academic estimates of publisher revenue under Google's Privacy Sandbox. By evaluating publisher-side performance, our findings expand the literature on the economic effects of privacy regulation to encompass technological responses that aim to mitigate these tradeoffs. This contributes to ongoing policy discourse concerning the viability of privacy-enhancing technologies for firms in general and the digital advertising industry in particular.

Our study has several limitations. First, the limited adoption of Privacy Sandbox during the study period constrains its effectiveness. Our findings should therefore be interpreted as directionally informative rather than as long-run general equilibrium estimates.

Second, the treatment combines multiple Privacy Sandbox technologies. While this bundled approach reflects its real-world implementation, this precludes us from isolating the contribution of specific components. In the future, we hope to examine individual technologies—such as the Protected Audience API and Topics API—to assess their distinct economic effects.

Third, the analysis is based on data from a single ad management company. Although this firm represents a large and diverse portfolio of publishers, the data capture only one point in time and one segment of the market. Given the scope and complexity of the digital advertising ecosystem, a comprehensive evaluation of Privacy Sandbox requires additional studies spanning multiple firms and roles within ad tech.

Finally, our analysis reflects a setting in which third-party cookies were either fully active or

fully deactivated. With Google's decision to halt cookie deprecation, Privacy Sandbox will now coexist alongside third-party cookies. In this hybrid environment, performance outcomes may diverge from those estimated in our study. Future research should explore this mixed adoption scenario to better understand its implications for publishers, advertisers, and platforms.

References

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–92.

Alcobendas, M., Kobayashi, S., Shi, K., & Shum, M. (2023). The impact of privacy measures on online advertising markets. Available at SSRN 3782889.

Aridor, G., Che, Y.-K., Hollenbeck, B., Kaiser, M., & McCarthy, D. (2024). Evaluating the impact of privacy regulation on e-commerce firms: Evidence from Apple's App Tracking Transparency. Available at SSRN.

Bak-Coleman, J., Bergstrom, C. T., Jacquet, J., Mickens, J., Tufekci, Z., & Roberts, T. (2023). Create an IPCC-like body to harness benefits and combat harms of digital tech. *Nature*, 617(7961), 462–464.

Baron, D. P. (2010). Morally motivated self-regulation. *American Economic Review*, 100(4), 1299–1329.

Competition and Markets Authority (2020a). Appendix F: the role of data in digital advertising. In *Online platforms and digital advertising market study*.

Competition and Markets Authority (2020b). *Online platforms and digital advertising market study*. Technical report, Competition and Markets Authority.

Competition and Markets Authority (2022). *Decision to accept commitments offered by Google in relation to its Privacy Sandbox Proposals*. Technical report, Competition and Markets Authority.

Competition and Markets Authority (2023a). *Quantitative testing of Google's Privacy Sandbox technologies – Additional CMA guidance to third parties on testing*. Technical report, Competition and Markets Authority.

Competition and Markets Authority (2023b). *Quantitative testing of Google's Privacy Sandbox technologies – CMA guidance to third parties on testing*. Technical report, Competition and Markets Authority.

DeMarzo, P. M., Fishman, M. J., & Hagerty, K. M. (2005). Self-regulation and government oversight. *The Review of Economic Studies*, 72(3), 687–706.

Dubé, J.-P., Bergemann, D., Demirer, M., Goldfarb, A., Johnson, G., Lambrecht, A., Lin, T., Tuchman, A., Tucker, C. E., & Lynch, J. G. (2025). The intended and unintended consequences of privacy regulation for consumer marketing. Available on SSRN.

Evans, G. & King, G. (2023). Statistically valid inferences from differentially private data releases, with application to the Facebook URLs dataset. *Political Analysis*, 31(1), 1–21.

Gadotti, A., Rocher, L., Houssiau, F., Creţu, A.-M., & de Montjoye, Y.-A. (2024). Anonymization: The imperfect science of using data while preserving privacy. *Science Advances*, 10(29), eadn7053.

Goldfarb, A. & Que, V. F. (2023). *The Economics of Digital Privacy*. Working Paper 30943, National Bureau of Economic Research.

Goldfarb, A. & Tucker, C. (2011). Privacy regulation and online advertising. *Management Science*, 57(1), 57–71.

Google (2024a). Google Marketing Platform Non-O&O/Display ARA E2E CMA Test Study. Google ads-privacy Github.

Google (2024b). *Privacy Sandbox Progress Report: Q1 Reporting Period - January to March 2024. Prepared for the CMA, 25 April 2024*. Technical report.

Google (2024c). *Testing privacy preserving signals in the absence of third-party cookies on Google's display ads platforms*. Technical report, Google.

Google (2025). *Privacy Sandbox Progress Report: Q1 Reporting Period - January to March 2025. Prepared for the CMA, 29 April 2025*. Technical report.

Hard, A., Kiddon, C. M., Ramage, D., Beaufays, F., Eichner, H., Rao, K., Mathews, R., & Augenstein, S. (2018). Federated learning for mobile keyboard prediction.

Hotz, V. J., Bollinger, C. R., Komarova, T., Manski, C. F., Moffitt, R. A., Nekipelov, D., Sojourner, A., & Spencer, B. D. (2022). Balancing data privacy and usability in the federal statistical system. *Proceedings of the National Academy of Sciences*, 119(31), e2104906119.

IAB Tech Lab (2023). *Identity Solutions Guidance*. Technical report, IAB Tech Lab.

Index Exchange product team (2024). Insights from Privacy Sandbox testing. Index Exchange Blog.

Information Commissioner's Office (2019). *Guidance on the use of cookies and similar technologies*. Technical report.

Information Commissioner's Office (2022). Privacy-enhancing technologies (PETs). In *Draft anonymisation, pseudonymisation and privacy enhancing technologies guidance* chapter 5. Information Commissioner's Office.

Jerath, K. & Miller, K. M. (2023). Using the dual-privacy framework to understand consumers' perceived privacy violations under different firm practices in online advertising. Available at SSRN.

Johnson, G. (2024a). Unearthing privacy-enhancing ad technologies (PEAT): The adoption of Google's Privacy Sandbox. *Available at SSRN*.

Johnson, G. A. (2024b). Economic research on privacy regulation: Lessons from the GDPR and beyond. In A. Goldfarb & C. E. Tucker (Eds.), *Economics of Privacy* chapter 4. University of Chicago Press.

Johnson, G. A., Lin, T., Cooper, J., & Zhong, L. (2023). COPPAcalypse? The YouTube settlement's impact on kids content. Available at SSRN 4430334.

Johnson, G. A., Runge, J., & Seufert, E. (2022). Privacy-centric digital advertising: Implications for research. *Customer Needs and Solutions*, 9(1), 49–54.

Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1), 33–51.

Jounce Media (2024). *March 2024 RTB Supply Path Benchmarking Report*. Technical report, Jounce Media.

Kenny, C. T., McCartan, C., Kuriwaki, S., Simko, T., & Imai, K. (2024). Evaluating bias and noise induced by the U.S. Census Bureau's privacy protection methods. *Science Advances*, 10(18), eadl2524.

Kircher, T. & Foerderer, J. (2024). Ban targeted advertising? An empirical investigation of the consequences for app development. *Management Science*, 70(2), 1070–1092.

Kobayashi, S. J., Johnson, G. A., & Gu, Z. (2024). Privacy-enhanced versus traditional retargeting: Ad effectiveness in an industry-wide field experiment. Available at SSRN.

Komarova, T. & Nekipelov, D. (2020). Identification and formal privacy guarantees. *arXiv preprint arXiv:2006.14732*.

Laub, R., Miller, K., & Skiera, B. (2024). The economic value of user tracking for publishers. SSRN working paper.

Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (2022). Does privacy regulation harm content providers? A longitudinal analysis of the impact of the GDPR. Available at SSRN.

Manski, C. (2007). *Identification for prediction and decision*. Harvard Univ Pr.

Marti, D. (2022). Topics API data: what a million impressions tell us.

Martin, A. (2024). Privacy Sandbox's latency issues will cost publishers. AdExchanger.

Maxwell, J. W., Lyon, T. P., & Hackett, S. C. (2000). Self-regulation and social welfare: The political economy of corporate environmentalism. *The Journal of Law and Economics*, 43(2), 583–618.

Miller, K. M. & Skiera, B. (2024). Economic consequences of online tracking restrictions: Evidence from cookies. *International Journal of Research in Marketing*, 41(2), 241–264.

Pascoe, A. (2024). Results and recommendations from our Privacy Sandbox testing. The NextRoll Blog.

Ravichandran, D. & Korula, N. (2019). *Effect of disabling third-party cookies on publisher revenue*. Technical report, Google Inc.

Selman, E. (2024). Criteo's Privacy Sandbox market testing. Criteo R&D Blog.

Sousa, T. (2024). Competition and targeting in auctions for online ads. Working paper.

Steed, R., Liu, T., Wu, Z. S., & Acquisti, A. (2022). Policy impacts of statistical uncertainty and privacy. *Science*, 377(6609), 928–931.

Warnat-Herresthal, S., Schultze, H., et al. (2021). Swarm learning for decentralized and confidential clinical machine learning. *Nature*, 594(7862), 265–270.

Wernerfelt, N., Tuchman, A., Shapiro, B. T., & Moakler, R. (2025). Estimating the value of offsite tracking data to advertisers: Evidence from Meta. *Marketing Science*, 44(2), 268–286.

Zhou, J., Chen, S., Wu, Y., Li, H., Zhang, B., Zhou, L., Hu, Y., Xiang, Z., Li, Z., Chen, N., Han, W., Xu, C., Wang, D., & Gao, X. (2024). PPML-Omics: A privacy-preserving federated machine learning method protects patients' privacy in omic data. *Science Advances*, 10(5), eadh8601.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power, edn*. PublicAffairs, New York.

Appendix

A Experimental validity

We assess the internal and external validity of the field experiment used in this study.

A.1 Internal validity

We evaluate internal validity by examining the balance of observations and covariates across experimental groups. Our analysis excludes subgroups exhibiting clear imbalances due to bot activity or other reason. After filtering, remaining observations display only economically small differences across treatment groups.

We identify several sources of potential imbalance. We drop Unix-based users—a common vector for bot traffic[12]—and impressions from outdated Chrome versions (v119 and earlier), which exhibit minor imbalances. Despite efforts to filter bot traffic, some residual bot activity is expected in an open experiment on the web.

Ad latency, which varies by treatment group, introduces small causal differences in the number of successfully loaded impressions (see Figure 1)—our unit of analysis. This is likely due to variation in the mix of selling channels and buyer-side data requests across groups.

Table 4 shows that the Status Quo, Privacy Sandbox, and Cookieless groups include 104.5M, 76.3M, and 26.4M impressions respectively. These closely match the expected randomization shares of 1%, 0.75%, and 0.25%,[13] although the Sandbox group falls short (−2.85%) of its counterparts. As discussed in Sections 5 and 6, we interpret this deviation as a causal consequence of Privacy Sandbox on the share of sold impressions.

Table 5 displays covariate balance by user device and region. Differences are small among sold impressions. These are likely driven by longer ad latency in the Sandbox group, which reduces impression delivery—especially on mobile devices. Table 6 confirms that latency effects are more

[12]This issue is discussed in Google (2024b) and here: https://github.com/WICG/turtledove/issues/982.

[13]The Status Quo group slightly under-delivers (−0.22%) and the Cookieless group slightly over-delivers (+0.66%) relative to the average of these two "control groups." These discrepancies may reflect longer ad rendering times for impressions with third-party cookies. However, the differences are small enough that we retain our assumption from Section 5 that the Status Quo and Cookieless group impression counts represent the number of impression opportunities.

Table 4: Observation Counts by Treatment Group

Group	Randomization Probability	Impression Counts	Deviation from Controls† (%)
Status Quo	1%	104,510,637	-
Privacy Sandbox	0.75%	76,283,492	-2.85%
Cookieless	0.25%	26,363,067	-

Notes: †Deviation from controls refers to the group's percentage deviation from the average number of impressions in the Status Quo & Cookieless groups, given the randomization probabilities (see Footnote 7 for calculation). Note that we exclude the Privacy Sandbox from this comparison as we expect that the Privacy Sandbox has a causal effect of reducing the number of impressions sold.

Table 5: Covariate Balance by Devices and Regions

	Status Quo	Privacy Sandbox	Cookieless
Device Types			
Desktop	40.31%	41.56%	40.25%
Smartphone	57.32%	56.10%	57.37%
Tablet	2.37%	2.34%	2.38%
Countries & Regions			
EU	6.65%	6.79%	6.76%
US	61.13%	60.66%	61.01%
Others	32.24%	32.55%	32.22%

Notes: Sold impression shares by experimental group.

severe on mobile, which may explain slight over-representation of desktop impressions in the Sandbox group.

While these differences are statistically significant due to the large sample size, they are small in magnitude and consistent with the causal effect of serving ads through different channels on latency and impression delivery. Taken together with the covariate and group-size balance checks, we view the data as consistent with proper randomization.

A.2 External validity

The CMA-Google field experiment enhances external validity by taking place in a live market setting, where buyers and sellers actively interact across randomized user groups. Vendors had prior access to Privacy Sandbox APIs, allowing time for integration and testing, and the broader advertising ecosystem had already partially adapted to a cookieless environment due to changes in Safari and Firefox.

Still, our results reflect market conditions and technology maturity as of 2024. The effectiveness of Privacy Sandbox depends on adoption and investment decisions that hinge on perceived

performance—creating a classic coordination problem. As such, the experiment cannot capture the long-run equilibrium effects of cookie deprecation or the full potential of Sandbox technologies at scale. To the extent that the APIs' use remained concentrated within <1% of Chrome users during the experiment, the findings may not generalize to a fully transitioned ecosystem.

Nonetheless, our analysis directly addresses the key policy question: can privacy-enhancing advertising tools sustain publisher revenue? In this respect, our findings provide directional evidence for ongoing regulatory evaluations of Privacy Sandbox and related technologies.

B Heterogeneity estimates

B.1 Ad latency

We assess the impact of Privacy Sandbox on ad latency across experimental groups and device types. Table 6 reports that Privacy Sandbox impressions exhibit consistently higher latency at both the median and 75th percentile. Median latency in the Sandbox group is 1.81 seconds, compared to 0.95 seconds in the Status Quo and 0.82 seconds in the Cookieless group.

Latency differences are more pronounced on mobile devices. The gap between the median latency in the Sandbox and Status Quo groups is 0.41 seconds on desktop, 1.41 seconds on smartphones, and 1.82 seconds on tablets. These patterns suggest that the latency burden of Privacy Sandbox—especially the Protected Audience API—is exacerbated on mobile devices, likely due to more limited computational resources required for on-device processing.

B.2 Ad revenue

We explore heterogeneity in the effects of cookie removal and enabling Privacy Sandbox across publisher size and geographic regions. Note that the Privacy Sandbox results here include the impact of unsold impressions under Privacy Sandbox, which depress average revenue per impression opportunity.

We proxy for publisher size using impression volume and divide publishers into four quartiles: small, medium, large, and very large. For cookie removal, our Intent-to-Treat (ITT) estimates in Table 7 show that larger publishers experience smaller proportional revenue declines. This

Table 6: Ad Latency Heterogeneity Estimates

(a) Ad Latency Median Estimates

	All	Smartphone	Desktop	Tablet
Status Quo	0.95	1.11	0.75	1.70
Privacy Sandbox	1.81	2.52	1.16	3.52
Cookieless	0.82	0.98	0.65	1.53

(b) Ad Latency 75th Percentile Estimates

	All	Smartphone	Desktop	Tablet
Status Quo	2.24	2.80	1.54	4.48
Privacy Sandbox	5.664	6.60	3.15	8.77
Cookieless	2.08	2.67	1.39	4.22

Notes: Ad latency is with respect to sold impressions (see Footnote 6).

Table 7: Treatment Effect Estimates: Publisher Size Heterogeneity

	Publisher size			
Effect of Removing Cookies	Small	Medium	Large	Very large
Intent to Treat (%)	-36.27	-36.19	-34.55	-30.92
Share of impressions with cookies (%)	76.71	77.73	77.64	78.12
Average Treatment Effect on the Treated (%)	-44.53	-45.80	-44.31	-39.89
Effect of Privacy Sandbox				
Intent to Treat (%)	-34.88	-33.35	-33.05	-29.63
Recovery Share (%)	3.83	7.83	4.34	5.11

Notes: Publisher size is determined by impression volume and split into four size quartiles.

partly reflects greater prevalence of cookied impressions among larger publishers, though these differences are modest. Average Treatment Effect on the Treated (ATET) estimates reveal greater losses for small- and medium-sized publishers.

Privacy Sandbox ITT in Table 7 estimates also vary by publisher size. Very large publishers experience smaller revenue losses than smaller counterparts. Nevertheless, the recovery share estimates fluctuate (ranging from 3.83% to 7.83%) across publishers of different sizes.

These patterns echo findings from prior privacy economics research, which often shows smaller firms are more vulnerable to privacy restrictions than larger firms (Dubé et al., 2025; Johnson, 2024b). This suggests that privacy regulation that restricts cookies may have unintended competitive effects. Nevertheless, Privacy Sandbox improves revenue for all publisher groups.

We then evaluate heterogeneity by user country and region. Table 8 reports treatment effects for users in the US, EU and remaining countries. The ITT estimate for cookie removal in the EU

Table 8: Treatment Effect Estimates: Country Heterogeneity

Effect of Removing Cookies	EU	US	Remaining countries
Intent to Treat (%)	-49.79	-26.93	-45.25
Share of cookie imps. (%)	56.61	80.58	77.80
ATET (%)	-66.43	-32.61	-55.35
Effect of Privacy Sandbox			
Intent to Treat (%)	-38.60	-26.29	-40.67
Recovery Share (%)	22.47	2.37	10.12

(−49.79%) is nearly double that in the US (−26.93%), driven by both a lower share of impressions with cookies (56.61% vs. 80.58%) and larger ATET estimates (−66.43% vs. −32.61%). The latter pattern is consistent with prior observational work showing a greater impact of removing cookies from EU impressions (Johnson et al., 2020). These findings likely reflect the impact of GDPR-mandated consent banners, which may reduce the effective use of cookies in EU countries.

Table 8's Privacy Sandbox ITT estimates are also larger in the EU (-38.60%) than in the US (-26.29%). Yet the EU exhibits a much higher recovery share (22.47%) than the US (2.37%)) in Table 8. One possible explanation is that EU stakeholders—subject to stronger regulatory pressure—have adopted Privacy Sandbox technologies more extensively, contributing to greater relative performance.

Taken together, these results highlight how the effects of cookie deprecation and the promise of Privacy Sandbox vary meaningfully across publisher size and geography.