

Privacy Interventions and Publishers' Revenues: Evidence on Long-Run Effects and Ad-Traffic Heterogeneity

Tobias Kircher^a, Riley Grossman^b, Mike Smith^b, Alessandro Acquisti^c, Cristian Borcea^d, Yi Chen^b,
Cristobal Cheyre^e

^a TUM School of Management, Technical University of Munich, Campus Heilbronn, 74076 Heilbronn, Germany, tobias.kircher@tum.de.

^b Martin Tuchman School of Management, New Jersey Institute of Technology, 07102 Newark, NJ, United States.

^c MIT Sloan School of Management, Massachusetts Institute of Technology, 02142 Cambridge, MA, United States

^d Ying Wu College of Computing, New Jersey Institute of Technology, 07102 Newark, NJ, United States.

^e Ann S. Bowers College of Computing and Information Science, Cornell University, 14853 Ithaca, NY, United States

Abstract: A central concern in the online privacy debate is the possibility that privacy interventions such as Apple's App Tracking Transparency framework may jeopardize the business model of ad-supported news and media content providers. Prior empirical work on the impact of privacy interventions on publishers' revenues is limited to the direct and short-term effects, neglecting the possibility of increases in ad revenues for unprotected ad-traffic and of adaptation by the online advertising industry. If ad revenue increases from ad-traffic from unprotected users occurred, empirical research that examined the impact of privacy interventions on revenues from protected users would paint an incomplete picture of their overall impact on aggregate publishers' revenues. We study whether ad revenues recover over time and whether ad revenues from ad-traffic remaining addressable increase. For investigating our research questions, we exploit Apple's block of third-party cookies on all iPhone browsers, starting in September 2020. Our findings suggest that Apple's policy did not necessarily translate into lower revenues in terms of average ad compensation per impression: The ad compensation publishers received for protected ad-traffic decreased just temporarily while ad compensation increased for ad-traffic from users remaining addressable. These findings inform the current policy debate about the depreciation of third-party cookies, while also suggesting that previous privacy interventions may have influenced competition more than they have strengthened user privacy.

1. INTRODUCTION

A central concern in the online privacy debate is the risk that privacy interventions (including regulatory frameworks such as Apple's App Tracking Transparency framework, ATT), aimed at protecting personal data and constraining advertising networks' ability to track online users and build user profiles, may harm the business model of online publishers, especially ad-based news and media content providers. Online advertisers leverage users' information to infer their interests, track them across platforms and sites, and show them behaviorally targeted ads. When ad networks can track users and predict their interests (for instance, via third-party cookies), ads' relevance and effectiveness grow (Goldfarb and Tucker 2011). This increases advertisers' bids for ad impressions (Kraft et al. 2023), which in turn means higher per-impression revenues for publishers. Restricting ad networks' ability to track users may, therefore, harm publishers and their ability to produce quality content. 48% of publishers responding to the survey by Lotame said that they expect to have to cut jobs due to dropping ad revenues when third-party cookies were depreciated (Lotame 2021).

Much prior work, including both academic and industry research, as well as anecdotal evidence, does suggest that restrictions on tracking reduce publishers' ad revenues (CMA 2020; Wang et al. 2024). For instance, ads shown to iOS app users after tracking was restricted due to ATT were compensated much less than ads shown to Android users (Udonis 2024). This pattern aligns with research studying the advertising-centric negative impacts of privacy interventions on ad effectiveness (Goldfarb and Tucker 2011) and click-through rates (Aridor et al. 2024). A critical gap in the literature on the relationship between privacy protections and online advertising revenues is that previous studies did not account for the possibility of industry adaptation.

Empirical studies examining the effects of privacy protections on the online advertising ecosystem typically focus on interventions that apply to particular groups of users (for instance:

iOS users following the enactment of Apple's ATT), rather than to the entire population of users visiting a certain publisher website. In practice, existing privacy interventions restrict tracking and targeting for users with specific characteristics while leaving other - often substantial - segments of the market fully trackable and targetable (for instance, users of mobile operating systems other than iOS; and so forth). Because privacy interventions differentially affect different groups of ad-traffic, publishers may earn more revenues from ad-traffic where ads could still be targeted and users be tracked. Demand-side platforms may have to submit higher bids for ad-traffic remaining addressable due to smaller supply of impressions of desired audiences. Furthermore, the online advertising industry could develop alternative ad targeting technologies over time, which effectively replace the ad targeting technology restricted through privacy interventions, with revenues recovering to pre-intervention levels. If ad revenues from ad-traffic that remained addressable increased and ad revenues from protected ad-traffic recovered, then empirical research that examined the impact of privacy protections on revenues from protected users alone would paint an incomplete picture of the relationship between privacy protections restricting an ad targeting technology and publishers' revenues. Adaptation by the online advertising industry could offset the direct negative effects. Recent work that suggests that interventions such as the GDPR and ATT did not negatively impact publishers' ability to produce content (Lefrere et al. 2025; Cheyre et al. 2023) corroborates the possibility of industry adaptation.

In this manuscript, we aim to find out whether the restriction of an ad targeting technology necessarily results in lower revenues from online advertising for publishers and whether it could be possible to enhance users' privacy without compromising publishers' ability to provide high-quality content. We first investigate the question of whether the negative revenue impact is merely short-term. Next, we study potential increases in ad revenues from ad-traffic remaining addressable. In online advertising, addressability refers to the extent to

which ad traffic can be identified (e.g., via cookies, mobile ad IDs, login-based identifiers), targeted to user interests and behavior, and measured in view of frequency and attribution at the user level.

RQ 1: Do publishers' ad revenues from protected ad-traffic recover to pre-intervention levels after privacy intervention?

RQ 2: Do publishers' ad revenues from ad-traffic that remains addressable increase after privacy intervention?

We contribute to better understanding these relationships by studying the consequences of Apple's block of third-party cookie tracking on all iPhone browsers, starting in September 2020. We use a sample of 100,000 randomly selected ad impressions in each month from January 2020 to October 2021, provided by an anonymous publisher monetization platform. That is, we do not study total revenue effects – this would require population data, which we do not have – but the impacts on the revenues per impression. We create 7,960 website-month observations between January 2020 and September 2021. We compare ad-traffic from iPhone users to ad-traffic from Android-phone users. To study research question two, we compare ad-traffic from desktop users, which largely remained addressable, to ad-traffic from Android-phone users. Ad-traffic from desktop users contains audiences who were protected on their mobile devices by Apple's intervention and which arguably became more valuable as the number of times they were available for advertising fell.

Our empirical investigation yields one central finding: Apple's block of third-party cookies did not reduce publishers' revenue per impression in the long run. We present two main findings. First, publishers' revenues per impression from iPhone ad-traffic fell only in the short-term. Second, publishers' revenues per impression from desktop ad-traffic increased. Also the increase in revenues per impression from desktop ad-traffic was short-term. Therefore, on the bottom line, Apple's stop did not lead to a long-run decline in ad revenues per impression

on average. More broadly, our findings suggest that Apple's stop did not necessarily translate into lower revenues by publisher.

Our findings advance two strands of literature. First, we contribute to research on the revenue effects of privacy interventions in digital advertising. Prior studies document negative revenue impacts but rely almost exclusively on short-term observation windows and largely ignore industry adaptation (Kraft et al. 2023; Cecere and Lemaire 2023; Wang et al. 2024). We provide two key insights. Apple's stop of third-party cookies reduced publishers' revenues per impression only temporarily; in the longer run, revenues recovered, consistent with the development of alternative targeting technologies. Revenues increased for ad-traffic that remained addressable, implying that partial privacy interventions can generate supply-and-demand effects that raise prices for audiences that become harder to reach. Second, we contribute to the economics of privacy literature by documenting the effects of privacy interventions on publishers' revenues (Acquisti et al. 2016). While prior work examined impacts on app developers, start-ups, technology vendors, and content creators, evidence on publishers has been limited (Kircher and Foerderer 2024a; Cheyre et al. 2023; Kircher et al. 2024; Peukert et al. 2023; Johnson et al. 2023; Kircher and Foerderer 2024b). We show that the revenue effects of Apple's Intelligent Tracking Prevention (ITP) are short-lived, aligning with recent findings that major privacy regulations did not reduce content quantity or quality. More broadly, our results help reconcile conflicting findings in this literature. When alternative targeting technologies exist or some users remain unprotected, the economic effects of privacy interventions are limited and temporary (Cheyre et al. 2023; Lefrere et al. 2025). In contrast, when targeting is comprehensively restricted and no unprotected user groups remain, privacy interventions lead to persistent supply-side contractions (Kircher and Foerderer 2024; Johnson et al. 2024). The economic costs of privacy interventions may not extend to website publishers.

Our findings have implications for policymakers and platform managers. Contrary to common claims, platform-level privacy interventions do not necessarily reduce publishers' revenues per impression in the long run, which puts industry criticism into perspective. For operating system managers, the recovery of revenues per impression - apparently driven by the emergence of alternative tracking technologies - raises questions about the effectiveness of interventions such as third-party cookie blocking and ATT in meaningfully enhancing user privacy. Apparently, the online advertising industry found ways to bypass the restrictions of these policies. More broadly, our results inform debates on competition, supporting concerns that platform-led privacy interventions may affect competitive dynamics more than they improve privacy.

2. LITERATURE

Existing research in marketing has examined the revenue loss due to privacy initiatives that make users less trackable and by converse the value of cookies. The stream of research in marketing that evaluates the impact of privacy interventions on publishers' ad revenues so far has investigated the impacts of Apple's ATT (Kraft et al. 2023; Cecere and Lemaire 2023)) and Europe's GDPR (Wang et al. 2024). Kraft et al (2024) estimate a 20% decline in ad revenues per impression by app developers due to ATT. Using estimated ad price data from Facebook Ad Manager, find a decline in ad prices by 10% and a drop of the conversion rate by 7.5%. Wang et al. (2024) evaluated GDPR's impact on the ad compensations received by one US-based publisher in the five weeks after GDPR's enforcement. This stream of research is connected with a collection of studies by industry, which experimentally compared the value of ads targeted with cookies to ads that are cookieless, finding price differences of 21% to 70% (CMA 2020; Gu et al. 2025; Ravichandran & Korula 2019; Vargas 2024; Google 2024). However, these studies overlook industry adaptation. Especially ATT led to the development of alternative tracking technology (Cecere

and Lemaire 2023). The observation periods in previous research do not extend more than three months after the interventions took place, thereby preventing the study of industry adaptations that may take several months to take place (Cecere and Lemaire 2023; Wang et al. 2024). Another gap in this literature is the neglect of the impact of privacy interventions on ad revenues by publishers from ad-traffic that remains addressable. None of the studies investigated whether privacy interventions led ad buyers to spend more on ad-traffic remaining addressable.

The closest are three studies. A set of perspective articles assessed the consequences of privacy intervention to be local (Moradi et al. 2025; Acquisti 2023). According to them, the impacts of privacy interventions on aggregate, across-impression publisher revenues are nuanced (Acquisti 2023, Moradi et al. 2025). The model by Ding et al. (2025) predicts that if costs of developing an alternative ad targeting technology are low and a privacy intervention is stringent, then ad networks invest in the development of an alternative ad targeting technology, which can lead to an overall increase in ad revenues per impression by publishers. Their model also suggests that marketers shift their ad budgets from real-time bidding to direct contracts with publishers when a privacy intervention restricts tracking, a shift that raises ad revenues by publishers from direct contracts. We shed light on the longitudinal consequences of privacy interventions and whether privacy intervention leads to the increase of ad revenues from ad-traffic remaining addressable. In other words, we offer insight into whether the direct negative effects of privacy interventions are temporary or enduring.

The stream of research in marketing that examines the value of cookies offers several data points, but these are contentious and even collectively are limited in generalizability. There is a need for more data points. While some studies found large differences in per-impression revenues when cookies were or were not available (Johnson et al. 2020; Kraft et al. 2023;), others found only marginal differences (Marotta et al. 2019; Wang et al. 2024). In addition to

the contradicting findings from academic research, also evidence from media bloggers is mixed. iPad users who used a Safari version that still allowed third-party tracking were found to yield three times higher CPMs for publishers than iPad users who used the updated tracking-free version (Publir 2020). By contrast, other media bloggers raise the possibility that estimations of the value of cookies may be overstating the true value as estimations often rely on data on CPMs from the point of view of an advertiser. When tracking is disabled, advertisers pay fewer fees to advertising technology intermediaries (i.e., middlemen) who provide services, such as audience segmentation, that rely on the availability of a third-party cookie (Medium 2021; FouAnalytics 2021). Thus, some of the observed decline in advertisers' CPMs may not extend to publishers.

The previous research on the value of cookies is also hardly generalizable, even if considered together. The cross-sectional comparisons stem from time points around 2016 when most ad-traffic was trackable (Marotta et al. 2019; Johnson et al. 2020). Research in this domain primarily used data from desktop impressions (Johnson et al. 2020; Wang et al. 2024), whereas mobile ad impressions represent the majority of ad impressions today and appear to be much more affected than desktop impressions in industry research (Medium 2021; FouAnalytics 2021). Hence, existing research on the value of cookies is inconclusive. We examine to what degree the revenue per impression from a mobile user decreases when privacy intervention removes the availability of a third-party cookie. We add an important data point to the scholarly study of the value of cookies.

Another related stream of research is the economics of privacy (Acquisti et al. 2016). One body of research within this research domain is the research on the economic impact of privacy interventions. Previous work examined the consequences of privacy interventions for app developers (Kircher and Foerderer 2024a; Cheyre et al. 2023; Kircher et al. 2024), app start-ups (Kircher and Foerderer 2021), and web technology vendors (Peukert et al. 2023;

Johnson et al. 2023). Yet, little is known so far about how privacy interventions impact publishers. The closest are three insights. First, the ePrivacy Directive curbed the effectiveness of ads shown on publisher websites (Goldfarb and Tucker 2011). Second, the stop of targeted advertising to children on YouTube precipitated a considerable decline in the provision of child-directed video content on YouTube (Kircher and Foerderer 2024b). Third, the General Data Protection Regulation (GDPR) had little influence on the quantity and quality of website content (Lefrere et al. 2025). However, in order to better understand the impact of privacy interventions on publishers, we need evidence from privacy interventions other than the GDPR and evidence from publishers outside Europe as well evidence from a large number of websites (Wang et al. 2023). The data used by Wang et al. (2024) stems from one US-based publisher, which experienced little variation in the consent rate. It is known that there was little compliance with the GDPR in terms of consent banners in the early years after the GDPR took effect (Nouwens et al. 2020), and this noncompliance is still an issue (Bouhoula et al. 2024; Smith et al. 2024). We study the impacts of a stop of third-party tracking on the revenue from advertising for a large number of US websites.

Another gap in this domain within the literature on the economics of privacy evaluating privacy interventions is the set of privacy interventions that were studied by privacy researchers. So far, emphasis has been given to studying the GDPR (Peukert et al. 2023) and ATT (Cheyre et al. 2023) – opt-in privacy regimes. One series of research is devoted to studying the impact of children's privacy legislation in the US (Kircher and Foerderer 2024a; Kircher and Foerderer 2024b). Little research has been done on Apple's Intelligent Tracking Prevention (ITP) program, which restricts tracking on WebKit, the engine that runs the Safari browser. Like the ITP, ATT represents a technical restriction of third-party tracking, whereas the tracking protection from the GDPR hinges on website compliance.

Within the literature on the economic costs of privacy interventions, there is a seeming contradiction between different studies. While some studies report no effect on the provision of ad-supported digital goods (Cheyre et al. 2023; Lefrere et al. 2025), while others documented strong adverse consequences for the supply of ad-supported digital goods (Kircher and Foerderer 2024; Johnson et al. 2024). Lefrere et al. (2025) did not find evidence for a fall in the provision and quality of website content following the GDPR. Cheyre et al. (2024) did not reject the null hypothesis that Apple's ATT had no impact on the development of mobile apps. These findings can be reconciled with the string of papers reporting strong negative effects by two explanations that we look into: First, ATT only restricted one sort of ad targeting technology, the sharing of the IDFA, while other ad targeting technologies were unaffected. Second, both in the case of the GDPR and the case of ATT groups of ad-traffic remained addressable. This ad-traffic could have become more valuable, yielding higher ad revenues for firms and compensating the loss of ad revenues from protected ad-traffic. The documentation of negative consequences of privacy interventions stems from contexts, where the emergence of alternative targeting technologies was not possible with little alternative monetization options to advertising (Kircher and Foerderer 2024a, Kircher and Foerderer 2024b). Moreover, the demand-side reallocation effects to unprotected ad-traffic was not possible. The privacy interventions studied – the implementation of children's privacy regulation – was enforced across platforms (Kircher and Foerderer 2024, Kircher and Foerderer 2023; Johnson et al. 2024). An alternative explanation is little compliance, but this only applies to the GDPR, whereas ATT is a technical restriction of tracking. Website owners resorted to dark patterns in response to the GDPR, manipulating users to consent into tracking, often in a non-compliant way (Kirkman et al. 2022, Grossman et al. 2026). However, scholarly research has not yet provided evidence for that different ad-traffics are affected in different directions by privacy

interventions and that revenue drops of interventions restricting one ad targeting technology are merely short-term.

Our research also relates to the emerging literature on marketers reactions to consumer protections. In view of advertisers, existing research documented that advertisers shifted their campaigns towards Meta when TikTok was banned in the US (Donati and Fong 2025). The US Do Not Call registry led to increasing calls by telemarketers to unregistered consumers (Goh et al. 2015). However, little is known about how ad buyers react to the impaired ability to serve targeted ads and track conversions. We address this gap by studying whether prices paid for ad-traffic remaining addressable following privacy interventions increase.

3. DATA

We leverage data from an anonymous ad management service provider that partners with thousands of websites in the US. The data stems from programmatic advertising (Google 2019). More than 95% of the ads are shown to users located in the US.

This setting is of advantage due to three main reasons. First, the ad management service is leveraging advanced AdTech, including tracking technologies, which ensures effective treatment in a quasi-experiment studying the consequences of privacy interventions. Furthermore, the ad service provider possesses data for a large number of similar websites, which benefits both internal and external validity of an empirical quantitative investigation.

The data, obtained from the anonymous ad management service provider, contains 100,000 randomly selected ad impressions shown to mobile and desktop users for the second Monday in each month from January 2020 to September 2021 (2.1M impressions). The dataset includes the variables date, time, website (hashed), operating system (iOS, Android, etc.), browser (Safari, Chrome, etc.), advertiser (hashed), ad unit name (hashed), user ID (DoubleClick Cookie ID, hashed), country, and revenue in the unit CPM (cost per mille,

scaled). Our dataset comprises ads of creative size “300x250”, a very popular ad size that occurs both on mobile and desktop devices.

4. DESIGN

4.1. Setting

4.1.1. Apple’s policy change: Stop of third-party tracking on iOS

We leverage Apple’s block of third-party cookies on all browsers installed on iPhones, and on the Safari browser on MacOS, with the release of iOS14 on September 16, 2020 (Wilander 2020). Apple’s block came after a series of cookie restrictions on Safari under the Intelligent Tracking Prevention (ITP) program (Headerbidding 2024; Digiday 2017; Digiday 2019). For the first time, with the release of iOS14 browsers other than Safari were affected by Apple’s privacy interventions on browsers on iOS. The block applied to the browser engine (i.e., WebKit) on which browsers run on iOS; therefore, third-party cookies were removed from Chrome, and other browsers, on iPhones as well (Wilander 2020; Simo Ahava 2020).

Apple’s policy creates a desirable setting for exploring the impacts of privacy interventions for three reasons. First, unlike privacy legislation, Apple’s third-party cookie block did not require continuous enforcement to ensure website compliance and had immediate effects once implemented once implemented, as it technically removed websites’ ability to set third-party cookies. Apple’s policy change was automatically enforced, while legislation like the GDPR is associated with low compliance (Kirkman et al. 2022; Bouhoula et al. 2024; Smith et al. 2024; Grossman et al. 2026) which inhibits the effectiveness and impact on the digital advertising industry. Second, the time between the announcement and the enforcement was relatively short, making anticipation effects unlikely. The block was announced on March 24, 2020 and the block took effect in conjunction with the release of iOS14 on September 16, 2020. Third, this block only affected iPhone users, while not users of other devices. This allows us to

study potential increases in revenues from ad-traffic remaining trackable and construct control groups.

Industry analysts assessed Apple's restriction of third-party tracking on WebKit to impact the recognition of visitors of websites, and therefore, the profiling of users, retargeting, frequency capping, and ad attribution (Digital Power 2024; Simo Ahava 2020; Kevel 2022). Experts forecast that AdTech may likely respond by relying more on fingerprinting techniques (Simo Ahava 2020), server-side tagging for Google Analytics cookies, which allows publishers to identify the same user and measure ad statistics (Medium 2021), and extended identifiers. However, the deprecation of third-party cookies makes cross-site tracking more difficult, and substantially impacts the ability to build comprehensive user profiles, target specific users on publishers' websites, and attribute purchases to ads.

4.1.2. Method

We analyze the impacts of Apple's block of third-party cookies through a difference-in-differences (DiD) regression. We split sold ad impressions into four groups in Figure 1. Group (1) is the directly affected group: ad-traffic from iPhone users on Safari. While identification of affected ad-traffic is straightforward, the key challenge is the construction of an adequate control group, consisting of ad-traffic that is unaffected by the privacy intervention and that unlikely experienced indirect influence.

Group (2) is the control group for research question one: ad-traffic from Android-phone users. Because Apple's block only applied to iPhone users and not to Android-phone users, we exploit that Android was unaffected by the block of third-party cookies by Apple and use Android ad-traffic as control group. Like iPhone ad-traffic, Android ad-traffic pertains to mobile devices. In other words, Android ad-traffic is similar to iPhone ad-traffic.

Group (3) is the ad-traffic where we expect revenue increases: ad-traffic from desktop device users, including MacBook users and Windows computer users. While iPhone users

could not be tracked any more on their iPhone, many of them continued to be addressable by ad buyers on their desktop device. Many iPhone users own a Windows computer. In 2020, Apple's market share was 23% among desktop devices in the US, while it was 60% mobile devices (Statista 2024; Statista 2025a). Furthermore, many iPhone users not only remain trackable on their Windows desktop device, but also on their MacBooks, because a considerable share of MacBook users browses on Chrome. An indirect back-in-the-envelope calculation suggests that, given a market share of 23% of MacBooks in December 2020 and a Safari share of 14% in January 2020 in the US on desktop, 39% of MacBook users browsed on Chrome back then (Statista 2024; Statista 2025b). While advertising audiences that iPhone users are related to remained addressable, the number of times they could have been served ads shrank. After the block by Apple, these cohorts of iPhone users, which for example are interested in wealth management, high-end furniture, and luxurious cars, only were addressable when they used their desktop device. This likely increased the value of and demand for the instances when these ad audiences are available. Therefore, we expect revenue increases for desktop ad-traffic. Ad-traffic from desktop devices could also have experienced price increases due to proactive behavior by DSPs and advertisers, in addition to supply and demand dynamics.

At the same point in time, the advertising audiences that Android-phone users are associated with, including electronics, VPNs, and personal loans, were as often addressable as before. For this reason, we do not expect revenue increased for Android-phone ad-traffic. Hence, Android-phone ad-traffic is a suitable control group for research question one.

Group (4) is the control group for studying the impact on desktop ad-traffic: ad-traffic from MacBook users browsing on Safari. We exploit that less than 20% of MacBook users on Safari were addressable to ad buyers through a third-party cookie when the block occurred. The reason for this likely is that the lifetime of a HTTP cookie was restricted to seven days under the Intelligent Tracking Prevention program and people use their desktop device less

often, such that their respective cookie expires before they visit a website partnering with third-party ad networks again. Another reason may be the more pronounced use of tracking protection software by MacBook users browsing on Safari. This means that there was little negative influence on ad compensations for MacBook users. To ensure the unaffectedness of this control, we removed the remainder of traffic that was addressable before the block.

We aggregate the ad impression data to the website level as we are interested in understanding the average impact on ad revenues by publishers and because this ensures comparability between groups. The resulting matched-pair panel resembles perfect matching and follows prior research (Gao et al. 2025, Miller et al. 2024). To avoid potentially favoring a desired outcome, we do not match on the outcome variable (Stuart 2010). Figure 1 illustrates the research design.

One design choice is the definition of the pre- and post-periods of the DiD estimation framework. Apple made the announcement on March 24th, 2020 (i.e., the announcement date) and implemented the block on September 16th, 2020 (i.e., the enforcement date) in conjunction with the release of iOS14 and Safari 14. We use the enforcement date to define the pre- and post-periods because the direct negative effects and potential supply-and-demand-driven revenue increases on desktop would start to occur. Our descriptive evidence clearly demonstrates that the availability of third-party cookies started to plummet around the release of iOS14 (see Figure 6a). Our observation period begins on Monday, January 13, 2020, about 9 months before the enforcement date, September 16, 2020, and ends on September 13, 2021. This allows to understand the long-term consequences.

[Figure 1]

4.2. Variables

The central dependent variable of our empirical investigation is $\text{Log}(\text{AVG. PRICE})$. We log the variable AVG. PRICE to account for its skewness. This variable corresponds to the

log-transformed monthly average price per website i , against the background of our analysis on the website-month level. We use `AVG. PRICE` and `log(MEDIAN PRICE)` as alternative dependent variables. `MEDIAN PRICE` is the median of the ad compensations that website i received for showing ads in month t . A binary indicator that takes the values 0 and 1, `VIEWED` is useful for assessing the effective manipulation. Google Ad Manager considers an ad to be viewed when at least 50% of the creative was displayed for at least one second (Google 2021).

Two variables of the Google Ad Manager dashboard allow to measure the availability of third-party tracking cookies. Our primary measure is `DoubleClick Cookie`, a binary indicator that takes the value 1 if for an ad impression a tracking cookie was available. This variable is called the “Encrypted DoubleClick cookie ID” in the Google Ad Manager. On mobile devices, the `DoubleClick` cookie corresponds to Apples’s Identifier for Advertising (IDFA) and the Android Advertising ID (AdId) (Google 2025). The variable `SERVING RESTRICTION` also captures how the IDs can be used for advertising purposes (when available). This variable takes the values “Limited ads”, “Restricted data processing”, “non-personalized ads”, “Basic ads”, “Personalization disabled”, and “No restriction”. This variable can distinguish between users consenting to tracking and users who give their consent to some but not all purposes of cookie tracking (Google 2025).

Further variables include `SHARE IOS`, `SHARE WINDOWS`, and `PUBLISHER TRAFFIC`. `SHARE IOS` ranges between 0 and 1, measuring the share of users with an Apple device running on iOS (iPhone, iPod, iPad) per website in month t . The proportion of ad impressions on the website i in month t associated with Windows ad-traffic, `SHARE WINDOWS` also has values between 0 and 1. `PUBLISHER TRAFFIC` is the log of the number of ad impressions on website i in month t in our dataset. Table 1 presents the summary statistics of the variables used in our article.

[Table 1]

4.3. Sample

We start with 2.1M observations of ad impressions that we obtained. We dismiss impressions from outside the US as they could distort the results – more than 95% of website traffic stems from US-based visitors; visitors outside the US may be uninteresting to US-based advertisers and are using the web under different privacy regimes, especially when it comes to European visitors. 199,924 observations are deleted. In addition, 4,516 impressions are dropped as they are of type “interstitial”, arguably not comparably to the majority of ads as they interrupt usage and require a click to be closed. We drop 53,844 ad impressions on publisher websites for which we have less than five observations per month and treatment, control group, in order to get meaningful average ad compensations per month. 1,841,716 ad impressions remain.

Getting these impressions to the website-month level and calculating the average ad compensations per month and traffic entails the shift from 1,841,716 observations, which are based on ad impressions, to 11,727 observations, which are based on 1480 websites and 21 months. To keep the panel balanced and ensure comparability over time and across groups, we remove those websites that are not observed in the treatment group or control group as well as those sites that are not observed before the policy change or in the period after it until the end of the panel period, dismissing 1039 websites. The final panel consists of 7,960 website-months observations.

4.4. Estimation Model

We estimate the effects of the policy change on the revenue per impression by applying a difference-in-differences analysis in conjunction with fixed effects (Angrist and Pischke 2009):

$$\text{Log}(\text{PRICE})_{i,t} = \beta_0 + \beta_1 \text{TREAT}_i \times \text{AFTER}_t + \psi_i + \gamma_t + \alpha i_{i,t} + \varepsilon_{i,t} \quad (1)$$

where $\text{Log}(\text{PRICE})$ is the main dependent variable of interest in month t for website i , γ_t are month-day (i.e., time) fixed effects. AFTER is a perfect linear combination with the time fixed effects, thus absorbed. α_i are fixed effects on the level of the website, which absorb the factor TREAT . We include website-level and month fixed effects to control for any time-invariant heterogeneity and to account for time-varying heterogeneity constant across channels, while there is minimum heterogeneity due to our matched-pair design. The coefficient of interest is β_1 , the DiD coefficient. We also study the impact on the impression rate, which allows to understand impacts on the quality of ads being shown.

The vector $\psi_{i,t}$ represents control variables: We adjust for SHARE IOS , because the supply of relatively more iPhone user traffic could impact the treatment and control group ad-traffic differently. The control of SHARE WINDOWS is involved in regressions for the same reason: to adjust for the possibility of differences in the usage of the Internet across devices. By adding PUBLISHER TRAFFIC to the regression model, we account for a website's traffic. PUBLISHER TRAFFIC is the number of times we observe a website i in a given month t . We log this variable to account for its skewness. We do not control for AD AGENCY as variables related to ad agencies arguably are "bad controls" - they "might just as well be dependent variables too" (Angrist and Pischke 2009, p. 64). We cluster standard errors on the website level to address for the possibility of serial correlation in the residuals and heteroscedasticity (Bertrand et al. 2004).

5. RESULTS

5.1. The Impact of the Stop of Third-party Tracking on Revenues from iPhone Ad-traffic

Table 2 provides the results of comparing the ad-traffic from iPhone users to the ad-traffic from Android-phone users. Columns (1) and (2) report the regressions with $\text{Log}(\text{AVG.}$

PRICE) as outcome. Column (1) presents the coefficient of the DiD term when the regression does not involve controls. The DiD coefficient has the value -0.013 and is statistically significantly different from zero. Column (2) shows the results for when the regression involves controls. The effect decreases in magnitude to -0.012. With the dependent variable being log transformed, the effect in percent is -1.2%. Accordingly, Apple's block reduced the average ad compensation publishers received for showing ads to iPhone users by -1.2%.

[Table 2]

Figure 2 provides model-free evidence on the effect of Apple's stop of third-party cookies on the ad compensation for ad-traffic from iPhone users. Before Apple's policy change that became effective in September 2020, the ad compensation publishers received for iPhone traffic was mainly greater than the compensation received for Android traffic. After the policy change, the ad compensation for ad-traffic from iPhone users dropped below the level for ad-traffic from Android-phone users. The ad compensation for iPhone users stayed lower compared to the compensation for Android-phone users until June 2021. Then onwards, the compensation for the both mobile ad-traffics is roughly on the same level. This suggests that the negative consequences for ad compensations caused by Apple's block were short-term.

[Figure 2]

When comparing the weak effect magnitudes in the regressions with the graphical evidence, the question arises why the declines shown in the graphs suggest a much stronger decline in ad compensations than the coefficients in the regressions. We reconcile this seeming contradiction. The gap between the regression coefficient and the graphical evidence can be explained by that the majority of websites was little affected, while larger websites, which represent larger shares of ad impressions, were more affected. In other words, the regressions report the average impact on website publishers rather than the impact on the entire ad traffic on publisher websites which is driven by large website publishers. The larger impact for

popular websites aligns with previous research that found that privacy regulation in particular compounds the effectiveness of ads on general-interest websites (Goldfarb and Tucker 2011).

5.2. The Impact of the Stop of Third-party Tracking on Ad Revenues from Ad-traffic Remaining Addressable

In the regressions of Table 3, we explore how ad compensations for ad-traffic that remained trackable were affected. In Table 3 Columns (1) and (2), we compare publishers' ad-traffic from desktop device users to the same publishers' ad-traffic from MacBook users browsing on Safari. The rationale for this control group is expounded in the method section. Column (1) is the baseline with the results from the regression that does not include controls. The DiD-term is positive and statistically significant. Column (2) reports the DiD-term when the regression involves controls. The effect increases to 0.037. This suggests that Apple's block increased the average ad compensation publishers received for ads from desktop ad-traffic by 3.7% on average and *ceteris paribus* ($\exp(0.037)-1$). Therefore, Apple's block led to increases in ad revenues from ad-traffic that was not protected by Apple's block of third-party tracking.

[Table 3]

Figure 3 offers descriptive evidence of how Apple's policy affected ad compensation for the revenue increase from desktop ad-traffic. Figure 3 Panel (A) plots the revenues from traffic originating from Windows device users and compares Windows-traffic (black) to Mac-Safari-traffic (gray). Prior to the implementation of Apple's policy in September 2020, publishers generally received similar ad compensation for Windows and MacBook-Safari ad-traffic. Following the block, however, compensation for Windows ad-traffic increased sharply, while the ad compensation for MacBook-Safari ad-traffic did not increase much. Compensation from Windows ad-traffic remained greater than that from MacBook-Safari users until around May 2021. From this point, compensation levels for both Windows and MacBook traffic converged and evolved similarly. These patterns corroborate an initially relatively strong

reallocation effect, while also suggesting that the reallocation of ad budget to desktop users was temporary rather than persistent.

[Figure 3]

Figure 3 Panel (B) plots the revenues from traffic originating from ad-traffic from MacBook users on Chrome, comparing ad compensation for MacBook users on Chrome (black line) with ad compensation for MacBook users on Safari over time (gray line). While before the block both groups yielded similar ad compensation to website publishers, after the policy change took effect the ad compensation for ad-traffic from MacBook users on Chrome increased sharply. MacBook users on Chrome continued to be compensated at greater levels until July 2021. Both the increase in revenues from ad-traffic from Windows and the increase in revenues from ad-traffic from Mac users on Chrome do not prevail until the end of the observation period.

6. ROBUSTNESS

6.1. Treatment Day

Concerns could arise as to whether the policy change became effective before September 2020, when iOS14 and Safari14 were released. We have reason to believe that the policy change became effective with the release of iOS14 in September 2020 (Medium 2021; Simo Ahava 2020), whereas the policy change was announced on March 24 2020 in conjunction with iOS 13.4 (Wilander 2020). We examine when the availability of cookies declined and whether the decline was gradual or sharp. Figure 4 plots two variables that are indicative of when the policy change was enforced. Panel (A) visualizes the availability of the DoubleClick tracking cookie over time. It is evident from the graph that the availability of the DoubleClick declined sharply from September 2020 and quickly reached levels close to zero for ad traffic from iOS devices. By contrast, ad traffic from Android-powered devices did not

experience a decline during this period in terms of the availability of the cookie. This supports our design decision to use the iOS14 release date as the treatment day in our difference-in-differences framework. Furthermore, we examine the course of the variable PERSONALIZED ADS, which we derive from the GAM variable servingrestriction. We code this variable 1 if the variable servingrestriction is empty, in particular does not take the value “Non-personalized Ads”. For iOS ad-traffic, this variable experienced a sharp decline as well from late summer 2020, quickly reaching levels close to zero. This further corroborates our research design.

[Figure 4]

6.2. Effective Manipulation

Related concerns could emerge in view of the possibility that despite the enforcement of ITP, users with iOS-powered devices and of MacBooks continued to be trackable. We examine the variable AD VIEWED both for iOS ad traffic and for Android-phone ad traffic in Figure 6 Panel (A). After the policy change took effect in September 2020, fewer ads were counted as viewed compared to Android-phone ad traffic, while ad traffic from devices running on iOS and ad traffic from Android-powered devices had similar view rates before the policy change. The gap remained until April 2021. Panel (B) in Figure 6 illustrates the presence of a low-value ad network that wins the bid for the ad impression. These ad sources could comprise house ads as well. We identify low-value ad sources by looking at the ad compensations per ad network and we found patterns in terms of four ad networks that continuously paid very low and often the very same rates over time. The black line plots the proportion of the low-value networks for iOS ad traffic, the gray line the share of these networks among the users of Android-powered phones. Whereas before the policy change the low-value ad networks had similar shares for both ad-traffics, right after the policy change was enforced in September 2020 low-value ad networks were providing the ad in markedly more impressions for iOS ad-

traffic compared to Android-phone ad traffic. This gap remained until the end of the observation period.

[Figure 5]

6.3. Validity of the Control Group

It could be argued that our control group are affected, which would violate the Stabled Unit Treatment Variable Assumption (SUTVA) and upward bias our findings. Regarding the study of research question two – the impact on ad compensation for ad-traffic remaining addressable it could be argued that our control group of MacBook Safari ad-traffic is negatively affected even though we removed the remainder of ad impressions – less than 20% of impressions – where the cookie was available before the block, which would lead to an overestimation of the increase in revenues from ad-traffic from desktop users. To address this concern, we repeat the analysis for an alternative control group, precisely the control group from the study of research question one. Hence, we repeat the analysis with Android-phone ad-traffic as control group. Column (1) in Table 4 reports the coefficient. The coefficient on Treat x After is positive and significant and similar in magnitude like the coefficient documented in Table 3.

6.4. Sensitivity

We provide further checks that corroborate the increase in revenues from ad-traffic from desktop users. Analyzing the sensitivity of the effect, we repeat the regression with alternative dependent variables. Column (2) in Table 4 reports the results of the regression with Log(MEDIAN PRICE) as dependent variable. Next, Column (3) in Table 4 provides the estimates of the regression with AVG. PRICE as dependent variable. The results are consistent.

6.5. Parallel Trends Assumption

In relation to our research design, the application of a difference-in-differences estimation framework, the concern could be raised that the groups of ad-traffic had been on different trends. In other words, it could be argued that the groups were affected by the time-varying factors differently. We test the parallel trends assumption by employing a relative time model. The relative time model allows to extrapolate from the trends between the groups before the policy change to how the trends would have been after the policy change in the absence of the policy change. We plot the coefficients and the 95% error bars in Figure 7. Before the policy change took effect, there were only marginal differences in the trends. From January 2021 onwards, the $Treat \times After$ coefficient takes a negative sign and is clearly different from zero. This corroborates that the parallel trends assumption holds.

[Figure 6]

7. DISCUSSION

7.1. Main Findings and Interpretation

Our empirical investigation yielded two main findings. First, Apple's stop of third-party tracking on iOS only reduced ad revenues per impression for a limited time period. Second, Apple's stop of third-party tracking led to increases in ad revenues from desktop users. Also the increase in ad revenues per impression from unprotected ad-traffic was short-term. While the increase in ad revenues for unprotected ad-traffic was greater in absolute magnitude than the decline, the share of ad-traffic where prices declined was greater than the share of ad-traffic where prices increased. In other words, the relative increase in compensations for Windows ad-traffic was greater than the decrease for iPhone ad-traffic.

In summary, we conclude that Apple's decision to stop third-party tracking did not necessarily reduce publishers' ad revenues. This conclusion stems from two observations. First,

both the decline in ad revenues from protected ad-traffic and the increase in ad revenues from unprotected ad-traffic did not apply for a longer time period. Second, an analysis of the net effect of both impacts suggests that the decrease in ad compensation for ad-traffic from iPhone users (who were no longer trackable) is almost entirely offset by the increase in the ad compensation for ad-traffic for desktop ad-traffic. Following a back-of-the-envelope calculation based on the share of iPhone and desktop users in our sample, the net impact of Apple's block of third-party cookies on the average ad compensation per impression would be -0.2% ($55\% \times -1.2\% + 13\% \times 3.7\%$). This figure is much lower than estimates from many previous studies.

Apple's block of third-party tracking may still have reduced ad revenues by publishers even though it did not reduce the average ad compensation per impression. The reason is that we do not have population data. If we received data on the entire ad-traffic by publishers, we could study the overall revenue impact. Especially, the fill-rate (i.e. the proportion of ad slots that are actually sold), which determines the volume of ad impressions, would be necessary to understand the total revenue impact. The revenues might still have decreased when the fill-rate was negatively impacted. Therefore, we cannot study the total ad revenue impact. However, our data does offer some insight into how the fill-rate may have been affected. Panel B in Figure 6 plots the share of ad impressions bought by low-value ad networks. Low-value ad networks represent ad networks that make low bids for impressions and that fill ad slots that would otherwise not be sold. After the policy change, the share of iPhone ad traffic bought from low-value ad networks increased relative to the share of Android ad traffic purchased by low-value ad networks. This gap persisted through the end of the observation period.

7.2. Discussion of Effect Magnitude and Development of Cookie Alternatives

Compared to previous studies, we find a much smaller value of third-party cookies (Johnson et al. 2020, Wang et al. 2024). This can be explained by the development of cookie

alternatives. There are reports that Meta found other ways of targeting and measuring after Apple's ATT, effectively replacing the curbed availability of the IDFA. Meta representatives said that they built new AdTech tools in response to ATT, while this took time (Business Insider 2022). Ionut Ciobotaru, Chief Product Officer at Verve Group, said: [...] Attribution will shift to probabilistic approaches that rely far more on econometrics and media mix modeling than they do today [...] (Hunch 2023). AdTech, in the past years, developed identifiers alternative to the cookie in response to the decreasing availability of third-party cookies (YouTube 2024).

That third-party cookies would be restricted to a large extent has been known by the online advertising industry since January 2020. The industry expected that third-party cookies could even be phased-out. Chrome announced in January 2020 the plan to phase out third-party cookies (TechCrunch 2020). Likely, alternatives to third-party cookies, such as the UID2 by The Trade Desk and the Publisher Provided ID, were already in use when Apple's policy change took effect. By October 2020, one month after the implementation of the policy change, 25 supply-side platforms had adopted the LiveRamps's RampID, allowing advertisers to buy addressable audiences without cookies and mobile identifiers (LiveRamp 2020). The RampID was founded as early as 2016, developed in response to Google's plan to remove third-party cookies (Publift 2024). The Trade Desk's Unified ID2.0 was integrated into many supply-side platforms and publisher ecosystems from October 2020 and made available to the IAB Tech Lab as open source in May 2021 (AdExchanger 2021).

One trackability variable suggests that most iPhone users became addressable again for ad buyers. Figure 4 Panel (B) plots the variable PERSONALIZED ADS according to the Google Ad Manager (GAM) variable servingrestriction. The graph shows that for almost the entire iPhone ad-traffic personalized ads could be served again from May 2021 onwards. We attribute the rebound in ad personalization to the implementation of the Publisher Provided ID,

a first-party ad targeting technology, within Google's programmatic ad ecosystem from March 2021 onwards (Google 2021). This timeline aligns with that we across graphs observe that from May 2021 onwards ad compensation for iPhone traffic recovered to levels similar to Android-phone traffic. Therefore, there is reason to believe that the Publisher Provided ID plays an important role in replacing third-party cookies and recovering ad revenues for publishers..

7.3. Generalizability

We are discussing the generalizability of our findings. Increases in ad revenues from unprotected ad-traffic would not apply to global prohibitions of third-party cookies. Should policymakers prohibit third-party cookies globally, then we would still expect that revenues by publishers would not fall to a large extent. The industry has already developed alternative targeting technologies, which can replace third-party cookies at least partially. Should policymakers stop ad targeting technology other than third-party cookies we would still expect industry adaptation and the development of alternative targeting solution, especially when the privacy intervention is stringent and the development of tracking alternatives is not costly (Ding et al. 2015).

However, we think that our findings don't generalize to prohibitions of targeted advertising that span a diverse set of targeting technologies. Then, industry adaptation is not possible. This assessment is consistent with previous research that documented significant costs of banning targeted advertising (Kircher and Foerder 2024a; Kircher and Foerderer 2024b). While our findings point to smaller impacts on ad revenues than it is often claimed, publishers could be more adversely affected by a prohibition of tracking by public policymakers that remove tracking for all users in a certain jurisdiction. We expect stronger negative effects for opt-in privacy regimes that render all tracking technologies ineffective, such as the GDPR, which requires consent for serving targeted advertisements.

A pressing question surrounding generalizability is the transferability to news publishers. Our dataset consists of blogs, which provide content on food, parenting, lifestyle, etc, topics that are different from news content. Moreover, the blogs are arguably less popular than newspaper websites. Previous research documented that general-interest websites such as news websites rely more on user identifiers for targeting ads than niche websites (Goldfarb and Tucker 2011). Against the backdrop of our data including niche websites, we expect stronger negative effects for news websites, where contextual advertising is less effective. Furthermore, another reason for why our findings are little transferable to news websites is that large news publishers like tend to have their own AdOps teams, while the websites in our dataset cooperate with a professional ad management service provider, who arguably is quicker in adopting programmatic AdTech innovation than inhouse ad teams. Some tracking alternatives involve an hashed e-mail, which blogs may be more capable of obtaining than news websites.

Our findings also are limited to the impact on advertisements of the size “300x250”. There may be ad creative sizes that are more or less affected. Especially those ad creatives among performance advertisers, such as skyscraper, may be stronger affected. It also remains unknown how ad compensations that are CPC-compensated were impacted, an ad compensation scheme associated with performance campaigns and where tracking is of higher importance than in branding campaigns. Our data is limited to CPM-compensated campaigns. However, CPM campaigns and medium rectangle ads represent very large market shares in online advertising.

7.4. Theoretical Contributions

Our discoveries advance multiple literatures. The marketing literature on the revenue effects of privacy interventions offers several empirical findings (Kraft et al. 2023; Cecere and Lemaire 2023; Wang et al. 2024). However, the reported negative impacts are based almost exclusively on short-term observation periods. Furthermore, this literature so far has

overlooked the possibility of revenue increases from showing ads to unprotected users. In general, the prior literature in this domain neglected industry adaption. We contribute to this literature two main insights. First, Apple's block of third-party cookies only short-term reduced revenues per impression. This suggests that privacy interventions do not reduce publishers' revenues per impression in the long-run, possibly because of the development of targeting alternatives. Second, Apple's block led to an increase in the revenues from ad-traffic that remained addressable. Accordingly, privacy interventions that do not apply to all devices can entail supply-and-demand effects, leading to price increases for audiences that became harder to reach. Our findings provide empirical support for the theoretical conjectures of Ding et al. (2025). Their model predicts that when privacy intervention restricts one targeting technology ad networks develop alternatives, which can recover lost revenues. Our extensive analysis suggests that cookie-alternatives such as the PPID made revenues per impression return to pre-intervention levels.

Our findings also contribute to the scholarly research on the economics of privacy (Acquisti et al. 2016). Prior research in this area documented the consequences of privacy interventions for app developers (Kircher and Foerderer 2024a; Cheyre et al. 2023; Kircher et al. 2024), app start-ups (Kircher and Foerderer 2021), and web technology vendors (Peukert et al. 2023; Johnson et al. 2023), and YouTube creators (Kircher and Foerderer 2024b). We contribute to this research the impact on publishers' revenues, which is short-term only. This aligns with the recent research within this domain that found that the GDPR did not reduce the quantity and quality of website content (Lefrere et al. 2025). Additionally, we advance this literature evaluating the costs of privacy interventions by studying a stop of third-party cookies. Previous research mainly examined the GDPR (Peukert et al. 2023) and ATT (Cheyre et al. 2023) – opt-in privacy regimes – and children's privacy legislation in the US (Kircher and Foerderer 2024a; Kircher and Foerderer 2024b). We contribute the evaluation of the impact of

Apple's Intelligent Tracking Prevention (ITP) program. Our findings, furthermore, reconcile seemingly contradicting findings within the research on the economic costs of privacy interventions. While a string of research could not reject the null hypothesis that there were no effects on the provision of ad-supported digital goods (Cheyre et al. 2023; Lefrere et al. 2025), others found considerable declines in the provision of ad-supported digital goods (Kircher and Foerderer 2024; Johnson et al. 2024). Our findings can explain this contradiction. When the development of alternative ad targeting technology is possible, then the effects are limited in magnitude and short-term (Cheyre et al. 2023). When there is unprotected ad-traffic, publishers also experience increases in ad revenues from ad-traffic remaining addressable (Cheyre et al. 2023; Lefrere et al. 2025). This results in limited supply-side impacts (Cheyre et al. 2023; Lefrere et al. 2025). By contrast, when all ad targeting is restricted and there are no unprotected user groups, privacy interventions lead to market supply-side contractions (Kircher and Foerderer 2024; Johnson et al. 2024).

7.5. Implications

Our findings inform public policymakers and managers of operating systems such as Apple iOS. While it is often claimed, privacy interventions on certain platforms do not necessarily lead to lower revenues per impression for publishers in the long run. This puts the criticism of the publisher industry into perspective. For managers of operating systems, our findings of recovered ad revenues per impression – likely as a result of the development of tracking alternatives – raises the question of whether privacy interventions such as the focal block of third-party cookies and ATT are effective in strengthening user privacy. These findings also inform the debate surrounding privacy policies in terms of how they influence competition. Our findings support the concern that privacy interventions by platforms rather impact competition than benefitting privacy (Heise 2024; Soko and Zhu 2024). Future research can shed light on how privacy interventions impact the fill-rate and volume of ads, which

determine total ad revenues as well, and on how they impact advertisers' ability to acquire consumers and users.

REFERENCES

- Acquisti, A. (2023). *The economics of privacy at a crossroads*. Economics of Privacy. University of Chicago Press.
- AdExchanger (2021). *AdExchanger's Regularly Updated Guide To UID 2.0*
<https://www.adexchanger.com/online-advertising/adexchangers-regularly-updated-guide-to-uid-2-0/>
- AppleWorld (2025). Study: iPhones users use their iPhones more than iPad users use their iPads.
<https://appleworld.today/2025/01/study-iphones-users-use-their-iphones-more-than-ipad-users-use-their-ipads/>
- Aridor, G., Che, Y. K., Hollenbeck, B., McCarthy, D., & Kaiser, M. (2024). Evaluating the impact of privacy regulation on e-commerce firms: Evidence from apple's app tracking transparency.
- Bouhoula, A., Kubicek, K., Zac, A., Cotrini, C., & Basin, D. (2024). Automated Large-Scale Analysis of Cookie Notice Compliance. *Proceedings of the 33rd USENIX Security Symposium*.
- Business Insider 2022. The biggest losers and winners from Apple's major privacy clambdown, one year on, Access: April 03, 2025. <https://www.businessinsider.com/apples-privacy-changes-biggest-winners-and-losers-one-year-later-2022-4>
- Campbell, J., Goldfarb, A., & Tucker, C. (2015). Privacy regulation and market structure. *Journal of Economics & Management Strategy*, 24(1), 47-73.
- Casadesus-Masanell, R., & Hervas-Drane, A. (2015). Competing with privacy. *Management Science*, 61(1), 229-246.
- Cecere, G., & Lemaire, S. (2023). Have i seen you before? measuring the value of tracking for digital advertising. *Measuring the Value of Tracking for Digital Advertising (December 1, 2023)*.
- CMA (2020). Online platforms and digital advertising.
https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final_report_1_July_2020_.pdf.
- Digital Power (2024). The impact of ITP on analytics and the user experience. <https://digital-power.com/en/inspiration/impact-of-itp-on-analytics-and-the-user-experience/>
- Digiday (2017). Publishers are already feeling pain from Apple's move against ad tracking.
<https://digiday.com/media/publishers-already-feeling-pain-apples-move-ad-tracking/>
- Digiday (2019). Impacts of Apple's latest privacy update thwarting 'link decoration'.
<https://digiday.com/media/publishers-can-be-collateral-damage-impacts-of-apples-latest-privacy-update/>
- Donati, D., & Fong, H. (2025). The Cost of Banning TikTok: Implications for Digital Advertising. Columbia Business School Research Paper Forthcoming.
- FouAnalytics 2021. Media CPMs (CPMs paid to publisher) Does Not Differ Much Between Tracking Protection Browsers vs Tracking Browsers, Access: April 01, 2025 <https://fouanalytics.substack.com/p/media-cpms-cpms-paid-to-publisher>
- Headerbidding 2023. Safari ITP (Intelligent Tracking Prevention) – From 1.0 to 2.3, Access: May 01, 2025 <https://headerbidding.co/intelligent-tracking-prevention-itp/>
- Flurry (2020). The Who, What, and When of iPhone and iPad Usage <https://www.flurry.com/blog/the-who-what-and-when-of-iphone-and-ipad-usage/>
- Goh, K. Y., Hui, K. L., & Png, I. P. (2015). Privacy and marketing externalities: Evidence from do not call. *Management Science*, 61(12), 2982-3000.
- Gao, B., Wang, J., Ding, X., & Guo, Y. (2025). The pitfalls of review solicitation: evidence from a natural experiment on TripAdvisor. *Management Science*, 71(2), 1671-1691.
- Goldfarb, A., & Tucker, C. (2011). Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3), 389-404.
- Google (2019). What is difference between programmatic channels and non-programmatic? How to set programmatic deals.
<https://support.google.com/admanager/thread/15213177?hl=en&sjid=16321264093809701533-NA>.
- Google (2021). Active View reporting. <https://support.google.com/admanager/answer/3154105>.
- Google (2021b). Helping publishers thrive in today's privacy environment.
<https://blog.google/products/admanager/helping-publishers-thrive-todays-privacy-environment/>
- Google (2024). Results from Privacy Sandbox APIs testing.
<https://support.google.com/admanager/answer/15189422?hl=en>.
- Google (2025). Explore Data Transfer fields. <https://support.google.com/admanager/table/7401123?hl=en>.

- Grossman, R., Smith, M., Borcea, C., & Chen, Y. (2026). Using Salient Object Detection to Identify Manipulative Cookie Banners that Circumvent GDPR. Forthcoming in Proceedings of the International AAAI Conference on Web and Social Media (ICWSM).
- Gu, Z., Johnson, G., Shunto, K. (2025). Can privacy technologies replace cookies? Ad revenue in a field experiment. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5284526.
- Heise (2024). British regulators: Apple torpedoes innovation in mobile browsers. <https://www.heise.de/en/news/British-regulators-Apple-torpedoes-innovation-in-mobile-browsers-10108113.html/>
- Hunch 2023. Adapting to privacy measures like ATT: What's next for personalization?, Access: April 03, 2025 <https://www.hunchads.com/blog/att-update>
- 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.
- Johnson, G. A., Shriver, S. K., & Goldberg, S. G. (2023). Privacy and market concentration: Intended and unintended consequences of the GDPR. *Management Science*, 69(10), 5695-5721.
- Kevel (2022). How to Keep Safari ITP from Limiting Your Ad Revenue. <https://web.archive.org/web/20250327071403/https://www.kevel.com/blog/safari-itp>
- Kircher, T., & Forderer, J. (2023). Does privacy undermine content provision and consumption? Evidence from educational YouTube channels. *Evidence from Educational YouTube Channels* (March 17, 2023).
- Kircher, T., & Foerderer, J. (2024). Ban targeted advertising? an empirical investigation of the consequences for app development. *Management Science*, 70(2), 1070-1092.
- Kircher, T., Hui, K. L., & Forderer, J. (2022). Social disadvantages of privacy? its convergent consequences for the access to app-based education. *Its Convergent Consequences for the Access to App-based Education* (July 25, 2022). HKUST Business School Research Paper, (2023-116).
- Kirkman, D., Vanica, K., & Woods, D. W. (2023). DarkDialogs: Automated detection of 10 dark patterns on cookie dialogs. 2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P), 847–867.
- Kraft, L., Skiera, B., & Koschella, T. (2023). Economic impact of opt-in versus opt-out requirements for personal data usage: The case of apple's app tracking transparency. Available at SSRN 4598472.
- Lambrecht, A., Goldfarb, A., Bonatti, A., Ghose, A., Goldstein, D. G., Lewis, R., ... & Yao, S. (2014). How do firms make money selling digital goods online?. *Marketing Letters*, 25(3), 331-341.
- LiveRamp 2020. LiveRamp Enhances Its Authenticated Identity Infrastructure to Support Unified ID 2.0, Access: April 19, 2025 <https://liveramp.com/blog/liveramp-enhances-authenticated-identity-infrastructure-support-unified-id-2-0/>
- Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (2025). Does privacy regulation harm content providers? a longitudinal analysis of the impact of the GDPR. *Management Science*. Ahead of Print.
- Marotta, V., Abhishek, V., & Acquisti, A. (2019). Online tracking and publishers' revenues: An empirical analysis. *Workshop on the Economics of Information Security 2017*.
- Medium (2021). The impact of Apple's Intelligent Tracking Prevention (Part 2). <https://medium.com/sparkline/the-impact-of-apples-intelligent-tracking-prevention-part-2-c10840188843>
- Medium 2021. Impact of Loss of 3P Cookies on Publishers' Ad Revenue, Access: March 12, 2025 <https://acfou.medium.com/abstract-2fef374edb2>
- Moradi, P., Cheyre, C., & Acquisti, A. (2024). Are there economic grounds for regulating behavioral ads?. Working Paper.
- Nouwens, M., Liccardi, I., Veale, M., Karger, D., & Kagal, L. (2020, April). Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1-13).
- Peukert, C., Bechtold, S., Batikas, M., & Kretschmer, T. (2022). Regulatory spillovers and data governance: Evidence from the GDPR. *Marketing Science*, 41(4), 746-768.
- Publir (2020). Impact of 3rd Party Cookie Blocking on Ad Revenues: What our Data Shows. <https://publir.com/blog/2020/06/pre-post-third-party-cookie-blocking-browser-comparison-what-our-data-shows/>
- Publift (2024). What Is Unified ID 2.0?. <https://www.publift.com/blog/unified-id-2-0>
- Lotame 2021. STUDY: Nearly Half of U.S. Publishers Say Revenue Loss from Cookie Deprecation Could Lead to Reduced Workforce, Access: April 19, 2025 <https://www.lotame.com/resources/study-nearly-half-of-publishers-say-revenue-loss-from-cookie-deprecation-could-lead-to-reduced-workforce>
- Ravichandran, D., & Korula, N. (2019). Effect of disabling third-party cookies on publisher revenue. Google. https://services.google.com/fh/files/misc/disabling_third_party_cookies_publisher_revenue.pdf.
- Schmitt, J., Miller, K. M., & Skiera, B. (2024). The Impact of the General Data Protection Regulation (GDPR) on Online Usage Behavior. HEC Paris Research Paper No MKG-2021-1437.
- Simo Ahava (2020). Intelligent Tracking Prevention In iOS 14, iPadOS 14, And Safari 14.

<https://www.simoahava.com/privacy/intelligent-tracking-prevention-ios-14-ipados-14-safari-14/>
 Smith, M., Torres-Agüero, A., Grossman, R., Sen, P., Chen, Y., & Borcea, C. (2024). A Study of GDPR Compliance under the Transparency and Consent Framework. In Proceedings of the ACM Web Conference 2024 (WWW '24) (pp. 1227–1236).
 Pinto, B., Sokol, D. D., & Zhu, F. (2024). The Antitrust and Privacy Interface: Lessons for Regulators from the Data. *Geo. Mason L. Rev.*, 31, 1019.
 Statista (2024). Market share held by leading desktop internet browsers in the United States from January 2015 to December 2024 <https://www.statista.com/statistics/272697/market-share-desktop-internet-browser-usa/>
 Statista (2025a). Subscriber share held by smartphone operating systems in the United States from 2018 to 2024. <https://www.statista.com/statistics/266572/market-share-held-by-smartphone-platforms-in-the-united-states/>
 Statista (2025b). Market share held by the leading computer operating systems in the United States from January 2012 to January 2025. <https://www.statista.com/statistics/272667/market-share-held-by-operating-systems-in-the-us-since-2009/>
 Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 25(1), 1-21.
 TechCrunch (2024). Google wants to phase out support for third-party cookies in Chrome within two years <https://techcrunch.com/2020/01/14/google-wants-to-phase-out-support-for-third-party-cookies-in-chrome-within-two-years/>
 Udonis (2024). eCPMs for Rewarded Video, Interstitial & Banner Ads. <https://www.blog.udonis.co/mobile-marketing/mobile-apps/ecpms>.
 Vargas, A. (2024). Early tests of cookieless chrome ads are encouraging, but don't get too excited. AdExchanger. <https://www.adexchanger.com/online-advertising/early-tests-of-cookieless-chrome-ads-are-encouraging-but-dont-get-too-excited/>.
 Wang, P., Jiang, L., & Yang, J. (2024). The early impact of GDPR compliance on display advertising: The case of an ad publisher. *Journal of Marketing Research*, 61(1), 70-91.
 Wilander, J. (2020, March 24). Full third-party cookie blocking and more. WebKit. <https://webkit.org/blog/10218/full-third-party-cookie-blocking-and-more/>.
 YouTube 2024.Ep.1 | Why are 90% of publishers impressions invisible to DSPs? | DPO Mastery, Access: April 09, 2025 <https://www.youtube.com/watch?v=9-20rzIJfh4>

TABLES AND FIGURES

Table 1: Summary Statistics

			Mean	S.D.	Min	Median	Max
1	Price	Ad compensation website i received for showing an ad to a user	0.30	0.54	0.002	0.152	10.050
2	DoubleClick Cookie	1 if third-party tracking cookie was available for ad impression	0.43	0.50	0.000	0.000	1.000
3	Mobile	1 if ad is shown to a user of a mobile device	0.82	0.39	0.000	1.000	1.000
4	iOS	1 if ad is shown to a user with an iPhone, iPod, or iPad device	0.53	0.50	0.000	1.000	1.000
5	Android	1 if ad is delivered to a user with an Android-phone	0.29	0.45	0.000	0.000	1.000
6	MacBook	1 if ad is shown to MacBook-users	0.064	0.24	0.000	0.000	1.000
7	Windows	1 if ad is shown to Windows-users	0.11	0.32	0.000	0.000	1.000
8	Safari	1 if ad traffic is from Safari browser	0.55	0.50	0.000	1.000	1.000
9	Chrome	1 if ad traffic is from Chrome browser	0.41	0.49	0.000	0.000	1.000
10	Edge	1 if ad traffic is from Edge browser	0.02	0.14	0.000	0.000	1.000
11	Firefox	1 if ad traffic is from Firefox browser	0.02	0.14	0.000	0.000	1.000
12	Viewed	1 if ad shown to a user was actually viewed	0.91	0.28	0.000	1.000	1.000
13	Website Observations	Number of times website i is visited by users	408.54	5992.75	1.000	19.000	207,843.000

Note: Based on the full sample (N=1,332,265). The definitions of the variables DoubleClick Cookie and Viewed are adopted from the Google Ad Manager documentation The variable price was scaled with an unknown factor by the data provider, in order to obfuscate the exact values.

Table 2: Effect of the Block on Ad Revenue per Impression from iPhone Users

Note: The table reports the results of our study's first test, the effect of the stop of third-party tracking in all browsers on iOS on the compensations for ad-traffic from iPhone users. The treatment group comprises iOS ad-traffic for each website, averaged as per month. The control group consists of the Android ad-traffic for each website and month, taken to the average. FE means fixed effects. Observations are website-months. Heteroskedasticity-robust standard errors clustered at the channel-level are reported in parentheses. *, **, *** indicate significance at the 5%, 1%, and 0.1% levels, respectively. Adj. R-sq excludes the explanatory power of website and month FE.

DV = Log(Avg. Price)		
	(1)	(2)
	Baseline	Baseline + Controls
Specification	OLS	OLS
Treat x After	-0.013*** (0.003)	-0.012*** (0.003)
<i>Controls</i>		
Log(Publisher Traffic)		-0.004 (0.004)
Share iOS		-0.018 (0.009)
Observations	7,960	7,960
Channel FE	x	x
Month FE	x	x
Adj. R-sq	0.228	0.232
F	213.27***	189.31***

Table 3: Effect of the Block on Ad Revenue per Impression from Desktop Users

Note: The table reports the results of our study's second test, the effect of the stop of third-party tracking on iOS on the compensation for ads shown to ad-traffic remaining trackable. The treatment group comprises the ad impressions shown to desktop device users for each website, averaged as per month. The control group consists of the ad-traffic related to MacBook users on Safari for each website and month, taken to the average. FE means fixed effects. Observations are website-months. Heteroskedasticity-robust standard errors clustered at the website-level are reported in parentheses. *, **, *** indicate significance at the 5%, 1%, and 0.1% levels, respectively. Adj. R-sq excludes the explanatory power of website and month FE.

DV = Log(Avg. Price)		
	(1)	(2)
	Baseline	Baseline + Controls
Specification	OLS	OLS
Treat x After	0.033* (0.012)	0.037** (0.011)
<i>Controls</i>		
Log(Publisher Traffic)		-0.030** (0.008)
Share Windows		0.493** (0.177)
Observations	675	675
Channel FE	x	x
Month FE	x	x
Adj. R-sq	0.270	0.373
F	125.17	27.00

Table 4: Sensitivity of Revenue Increase from Desktop-traffic

Note: The table reports the results of different robustness checks. The concern could arise that the impact on ad compensation from desktop ad-traffic is partly the result of a negative impact on the control group, ad-traffic from MacBook users on Safari. Column (1) reports the impact on ad compensation from desktop ad-traffic when we use an alternative control group. In particular, we reuse the control group from the first analysis: ad-traffic from Android-phone users. Columns (2) and (3) present the results when we use alternative dependent variables, in particular the logged median price per website, group, and month and the average price per website, group in the estimation framework, and month. FE means fixed effects. Observations are website-months. Heteroskedasticity-robust standard errors clustered at the website-level are reported in parentheses. *, **, *** indicate significance at the 5%, 1%, and 0.1% levels, respectively. Adj. R-sq excludes the explanatory power of website and month FE.

	DV = Log(Avg. Price)	DV = Log(Median Price)	DV = Avg. Price
	(1)	(2)	(3)
	Alternative Control Group	Alternative DV	Alternative DV
Specification	OLS	OLS	OLS
Treat x After	0.039** (0.012)	0.044*** (0.012)	0.072*** (0.020)
Observations	1,975	675	675
Channel FE	x	x	x
Month FE	x	x	x
Adj. R-sq	0.384	0.351	0.315
F	44.64***	17.633	19.541

Figure 1: Research Design

Note: The figure illustrates the research design. Two groups of ad-traffic are influenced by Apple's policy, while two serve as control groups.

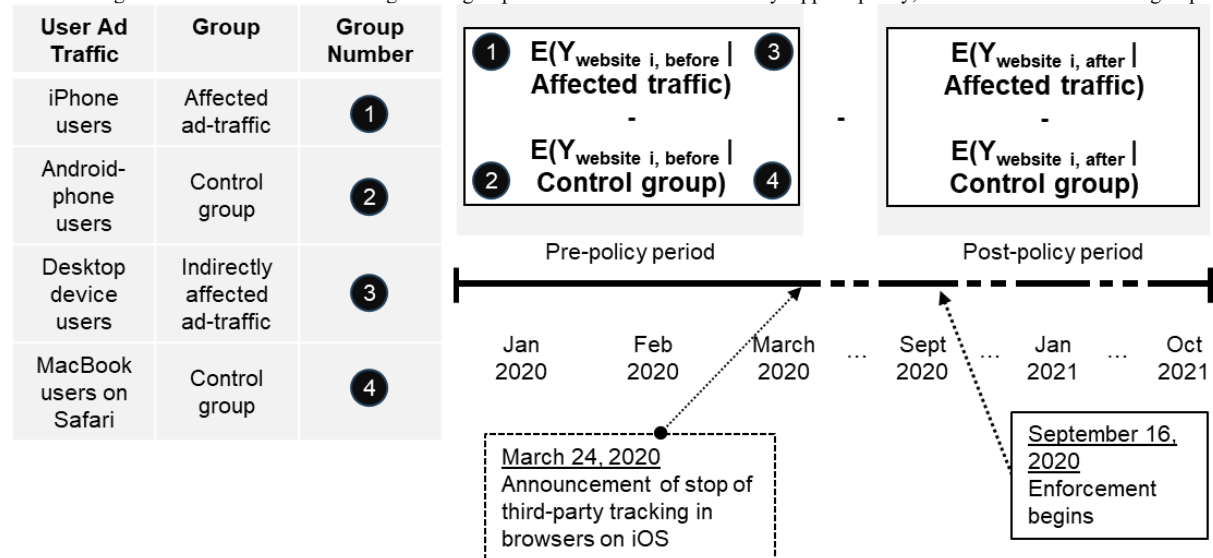


Figure 2: Impact of the Block on Ad Revenue per Impression from iPhone Users

Note: The figure illustrates the impact of Apple's policy on ad-traffic originating from iPhone users. The black line represents ad-traffic associated with iPhone users, the gray one ad-traffic associated with Android-phone users.




Figure 3: Impact of the Block on Ad Revenue per Impression for Desktop Ad-traffic

Note: The figure provides descriptive evidence on the indirect impact of Apple's policy on ad-traffic originating from desktop users for. Panel (A) compares ad-traffic from Windows users to ad-traffic from MacBook users on Safari. Panel (B) compares ad-traffic from MacBook users on Chrome with MacBook users on Safari. The black line represents ad-traffic related to users of Windows devices or users of MacBooks who browse on Chrome. The gray line plots the ad compensation for ad-traffic of MacBook users on Safari.

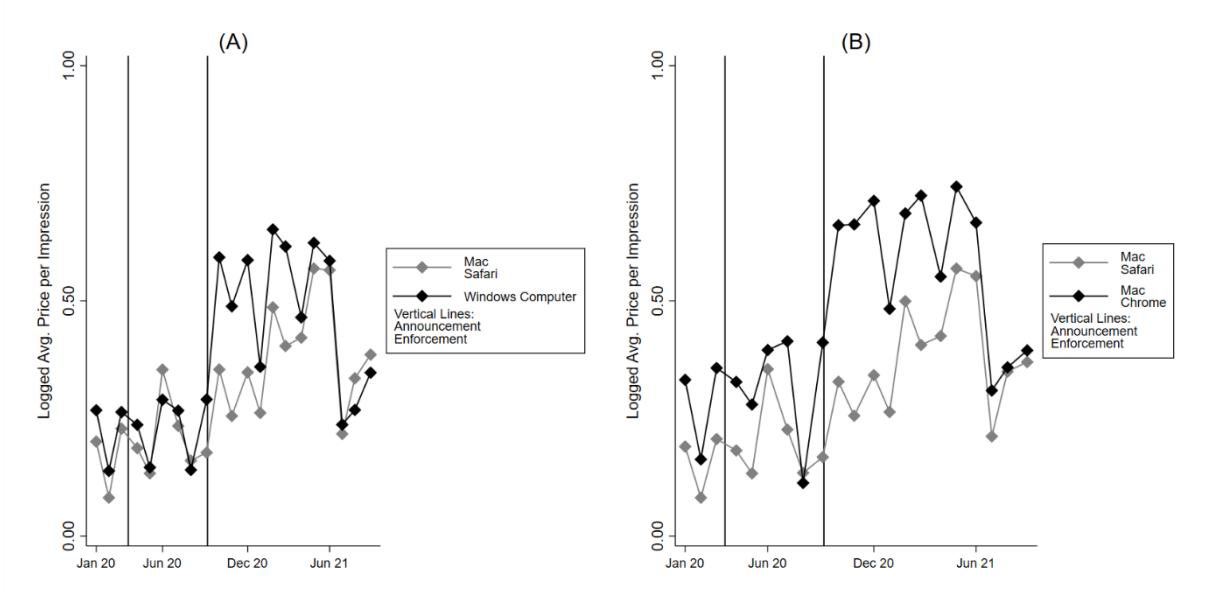


Figure 4: Treatment Day and Effective Manipulation

Note: The figure provides descriptive evidence on that the policy took effect in September 2020 when iOS14 was released. Panel A illustrates the availability of the DoubleClick Cookie ID. Panel B shows the development of the personalization of ads according to the variable SERVING RESTRICTION. Precisely, we code the variable PERSONALIZED ADS 1 if the variable SERVING RESTRICTION is empty instead of taking the value "Non-personalized Ads". The black lines denote iOS ad-traffic, the gray lines Android-phone ad-traffic.

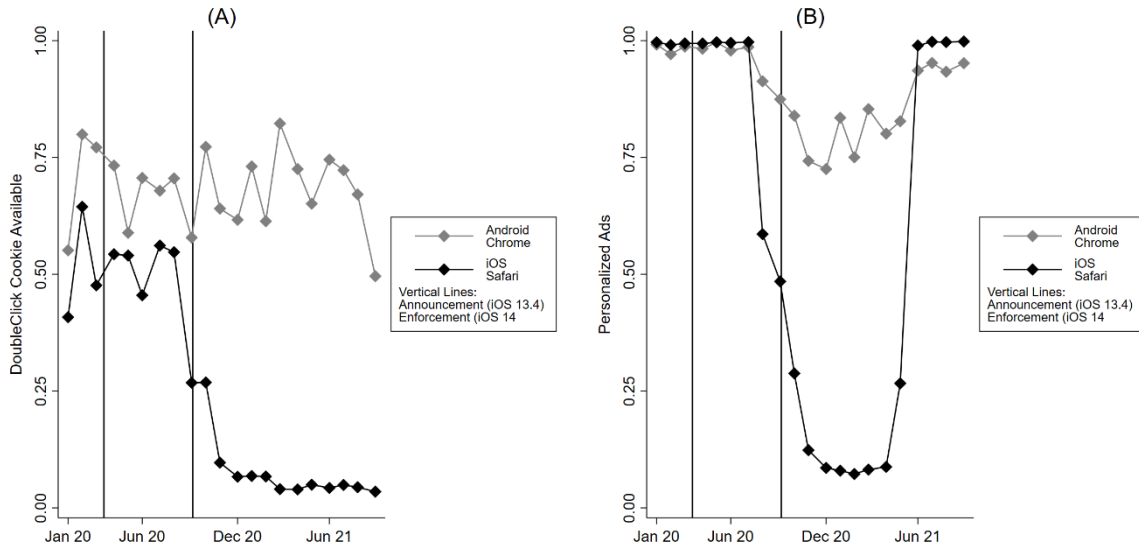


Figure 5: Effective Manipulation

Note: The figure provides descriptive evidence for the decline in trackability. The black lines represent ad-traffic from iOS devices, including the iPhone, the iPod, and the iPad. The gray lines stand for the ad-traffic from Android devices. Panel (A) plots the variable AD VIEWED over time. This variable takes the value 1 if the ad was viewed by the user (at least 50% of ad creative was displayed more than one second). Panel (B) visualized the share of low-value ad networks over time among all ad impressions.

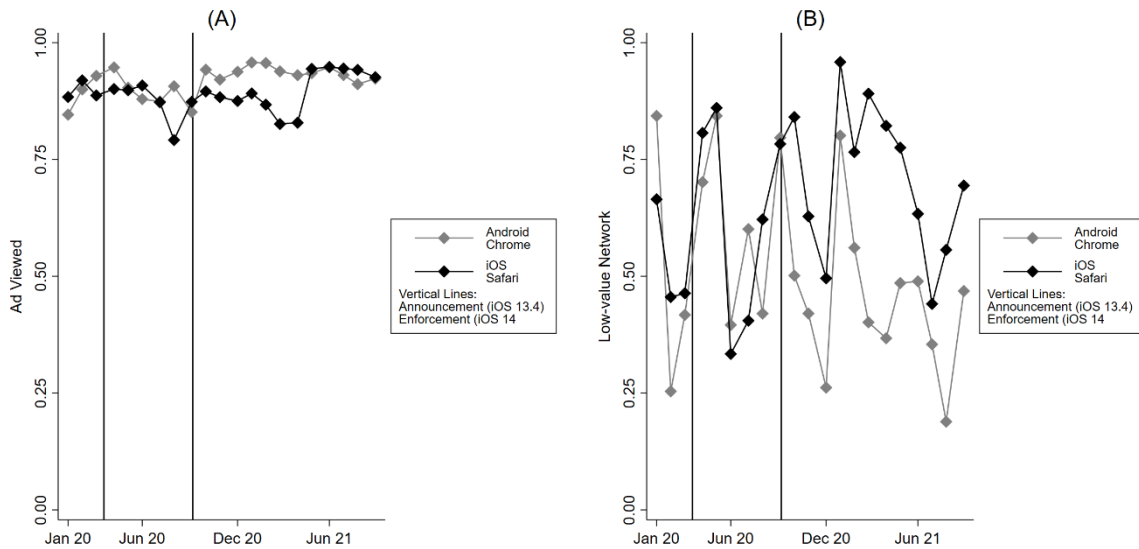


Figure 6: Relative Time Model

Note: To address concerns over anticipation, we examined potential anticipation effects and the validity of the parallel trends assumption in a relative time model.

