

Have I Seen you Before? Measuring the Value of Tracking for Digital Advertising*

Grazia Cecere[†] Sarah Lemaire[‡]

March 30, 2024

Abstract

Privacy regulation aiming to reduce the ability of ad platforms to aggregate user data can decrease the quality of ad display and thus challenge data-driven business models. We investigate the effect of privacy protection rules on the market for ads. We leverage a change in Apple’s privacy policy, the *App Tracking Transparency*, to compare predictions based on ad campaigns targeting iOS users versus Android users. To assess the effect of the policy, we use an original database of estimated ad outcomes on a social network in the US market. The results suggest a relative reduction in targeting efficiency and ad prices.

1 Introduction

Data collection has raised significant privacy concerns among consumers. These concerns have resulted in regulatory initiatives, not only by regulators aiming to protect consumers, but also by digital companies increasingly trying to attract consumers

*We thank Marc Bourreau, Alexandre de Cornière, Pierre Dubois, Isis Durrmeyer, Daniel Ershov, Renato Gomes, Ulrich Hege, Christian Hellwig, Doh-Shin Jeon, Ilja Kantorovitch, Laura Lasio, Nour Meddahi, Charles Pebereau, Guillem Roig, Wilfried Sand-Zantman, and Catherine Tucker for their helpful comments. We thank the participants of the EARIE Conference 2021, the IIOC Conference 2022, the Doctoral Workshop on The Economics of Digitization 2022, the EEA-ESEM Congress 2022, the Digital Economy Workshop 2023.

[†]Institut Mines Telecom, Business School, France, email: grazia.cecere@imt-bs.eu

[‡]Toulouse School of Economics, France, email: sarah.lemaire@tse-fr.eu

with privacy-friendly products or services. Although reassuring at first glance, these initiatives may have an adverse effect on economic outcomes (Goldfarb and Que, 2023; Johnson, 2022; Choi et al., 2022). In particular, limiting the collection and/or aggregation of data may hinder the efficiency of targeted advertising and deteriorate advertising revenues. This may in turn prevent firms from developing and maintaining qualitative internet services, hence reducing consumer surplus. It is important to examine this side of the trade-off, to be able to design efficient policies. This paper intends to shed light on how a change to privacy regulation impacts targeting efficiency, and to interpret the associated shift in equilibrium ad prices. More precisely, we look at how the *App Tracking Transparency (ATT)*—an initiative by Apple to offer more control over app tracking on iOS—affects predicted ad outcomes on META’s social network.

To conduct our study, we use an original database of estimated ad performances collected via Facebook Marketing Application Programming Interface (Facebook Marketing API). This tool is used by practitioners to run real ad campaigns, we provide several evidences that ad estimated predictions reflects the real market used by META. We collect these ad campaign predictions for a wide range of ad audiences from mid-March 2021 to mid-July 2021 and construct measures of the conversion rate (CR) and of the cost-per-mille (CPM). We use the CR as a proxy for targeting efficiency and the CPM as an indicator for price. Our empirical design exploits the fact that the *ATT* policy involves only iOS users and not Android users. The ability of firms to profile Android users hence does not change in the short run. We use a difference-in-differences design, using ads targeting iOS users as the treatment and ads targeting Android users as the control. We interpret the change in the conversion rate as the causal effect of the policy on the ability of Facebook to efficiently target consumers. We cannot, however, claim causality regarding prices,

because of their highly strategic nature. Nonetheless, as we restrict the period of study around the introduction of the policy, the change in predicted prices can be interpreted as a combination of the effects on quality, on demand, and on strategic market interactions. It provides rich intuitions on the market forces at play after the change in privacy policy.

We find that the introduction of Apple’s privacy scheme has been followed by a decrease in the effectiveness and prices of ads predictions targeted at iOS users relative to Android users. Ads targeted at iOS users triggered 7.5% fewer actions per impression after the introduction of the *ATT* policy. At the same time, the CPM to target iOS users decreased by 10% compared to the CPM to target Android users. Note that this decrease in prices is only relative, as our results suggest a general upward trend in prices during the period of study.¹ We find that the effect on CR strengthens over time. This trend aligns with the growing number of users adopting the new operating system iOS 14.5, which allows them to opt-in to *ATT*. The effect also intensifies when ads aim for an action harder to trigger. Indeed, we find that ads optimizing distribution with respect to app installs are more affected than ads optimizing on link clicks. Moreover, our robustness checks validate that the observed effect varies with respect to audience sizes. This suggests that the decline in ad targeting effectiveness cannot be attributed to technical difficulties in measuring conversions, as there is no reason for the technical ability to observe conversions to depend on audience sizes. As a consequence, we do capture a genuine decrease in targeting efficiency.

Our article contributes to two streams of literature. First, it contributes to the general empirical literature documenting the effects of privacy regulations on ad

¹This is consistent with official figures of META, suggesting an overall increase in the ad revenues in the US. See the Q3-2021 Earnings Presentation slides, p.10: https://s21.q4cdn.com/399680738/files/doc_financials/2021/q3/FB-Earnings-Presentation-Q3-2021.pdf, November 2023.

outcomes. Goldfarb and Tucker (2011) studied the impact of privacy regulation on ad effectiveness. Our paper distinguishes itself by analyzing publicly available data, used by advertisers for marketing purpose. Additionally, we study a privacy policy enforced in the market of mobile apps, characterized by pervasive data collection. A large strand of empirical studies look at the effect of privacy policies on publisher revenues in the context of open display advertising (Marotta et al., 2019; Ravichandran and Korula, 2019; Johnson et al., 2020; Alcobendas et al., 2021; Laub et al., 2022). Our article complements these, as we look at ad outcomes in another segment of display advertising: the social media owned and operated channel. Indeed, as opposed to publishers in the open channel, social media platforms such as Facebook have a fully integrated digital advertising system and hence control over the full value chain. This represents a large share of display advertising: according to practitioners, Facebook and Instagram accounted for about 24% of total digital ad revenues in the US in 2021.² Some of the existing studies focus on smaller players, but studying the largest player in the ad display market brings a different perspective. Our article explains how ad networks with important market power may be affected by privacy policies. Using insider data from META, Wernerfelt et al. (2022) show that limiting access to users' offsite data on a social media platform reduces ad campaign effectiveness. Similarly, Aridor and Che (2024) find a negative effect on offsite conversions relative to onsite ones, after the introduction of the ATT. While they both use sample data on ad campaigns, we look at largely aggregated ad outcomes at the audience level with an exogenous policy change. Moreover, the ATT policy permits us to investigate the effects of an *opt-in* policy: it offers privacy by default to consumers and does not request that they make an extra effort to do so. Most of the above-mentioned articles, on the contrary, look at

²<https://www.insiderintelligence.com/content/google-facebook-amazon-account-over-70-of-us-digital>. November 2023.

opt-out privacy policies with a rather low rate of pro-privacy choices. Aridor et al. (2020), for example, document that the introduction of the GDPR resulted in a 12.5% drop in total cookies. In the same way, Johnson et al. (2020) find that consumers' privacy preferences result in a 0.23% increase in ad impressions among those who opted out of behavioral advertising through the self-regulatory AdChoice program. Compared to these previous empirical works, the proportion of consumers making a pro-privacy choice in the case of *ATT* is estimated to be much higher, and therefore much less subject to selection biases.

Second, our article contributes to the nascent empirical work investigating the impact of Apple's recent privacy-friendly moves. Most papers of this strand of literature have studied the market consequences of a drop in targeting efficiency and/or ad revenues following the introduction of privacy labels (Bian et al., 2021) and the *ATT* (Li and Tsai, 2022; Kesler, 2022; Cheyre et al., 2023). They document the effect of *ATT* on the structure of the app market or on app revenues and business models. Like Kraft et al. (2023) and Ahmadi et al. (2023), we complement these analyses by providing evidences that the *ATT* had a negative effect on ad performances. We show that it affected ad outcomes for one of the most prominent players in the advertising market.

2 Context and Data

2.1 The App Tracking Transparency (*ATT*)

In April 2021, Apple introduced a feature on its new Operating System (OS), iOS 14.5, offering more control to consumers: the *ATT*. This feature is included in version 14.5 and all subsequent versions.³ It requires all apps to collect users' explicit consent

³We will denote 14.5 and later versions by 14.5+ throughout the paper.

before tracking them, with a simple yes-or-no form.⁴ As such, the *ATT* sets privacy by default, as the simplicity of the form makes it equally costly for consumers to choose tracking or no tracking.

More precisely, the policy requires app developers to ask users for their consent to access Apple’s *Identifier For Advertisers* (IDFA). The IDFA is a unique code assigned to mobile devices that allows data brokers and advertisers to track a given user across apps and to aggregate his data. Although ID tracking is not the only way to profile consumers (Jeon, 2021), it was expected that without the users’ consent, companies (including META) would be less efficiently able to match a given consumer to a precise ad content.⁵

iOS users account for almost 50% of mobile users in the US.⁶ While we do not have official statistics provided by Apple, the figures provided by practitioners measuring the rate of adoption of iOS 14.5+ and the rate of opt-in for tracking show high rates of adoption and pro-privacy choice. During the time span we studied (mid-March to mid-July 2021), adoptions of versions 14.5+ increased to around 75% (See evolution in Appendix A.2). Moreover, practitioners indicated a low rate of consent to tracking—between 12% and 25%—in September 2021.⁷ The change could hence concern roughly 1/3 of consumers’ privacy settings in the US mobile market.

2.2 Ad Outcome Predictions

We use Facebook *Ad Set Delivery Estimate* data, collected using Facebook Marketing API. This API is a tool developed by META, used by advertisers to plan massive ad

⁴See Figure 3 in the Appendix.

⁵<https://www.facebook.com/business/help/331612538028890?id=428636648170202> March 2023.

⁶According to Kantar, iOS had a market share of 46.8% in the US in September 2021 <https://www.kantar.com/campaigns/smartphone-os-market-share>, September 2021.

⁷<https://9to5mac.com/2021/07/23/app-tracking-transparency-opt-in-snap/> and <https://www.flurry.com/blog/att-opt-in-rate-monthly-updates/>, September 2021.

campaigns. It works through requests and allows advertisers to obtain predictions of ad outcomes—depending on a specified target audience—before running ad campaigns. A request returns estimates on the number of *impressions*, the number of *people reached* by an ad, the number of *actions* (e.g., clicks, app installs) triggered by the ad, for a large set of daily budgets (see a sample request output on Figure 5 in Appendix A.3).⁸ More precisely, each request returns three times 15 points forming three full supply functions — *Impressions*, *Reach*, and *Actions* — called the Ad Set Delivery Estimate. Compared to Facebook Ad Manager, the META’s API gives the size of the target audience on the day or month of the request (*DAU: Daily Average Users* and *MAU: Monthly Average Users*⁹). These predictions depend on the specified audience’s demographics, on the platform used to reach the audience (e.g. Facebook or Audience Network) and on what action the ad placement should maximize, i.e. which “optimization goal” (e.g. clicks, app installs).

2.3 Relevance of Facebook Delivery Estimate Data

While relying on performance data estimated by META itself may trigger worries regarding the black-box AI model with which these predictions are computed, we have several reasons to believe that this data is of interest for researchers and can complement the insights provided by other studies.

First, although the data generating process is not fully transparent, it is most likely based on AI models using the full amount of META past and current advertising performance data. This is much more exhaustive in fact than common datasets accessible for research on advertising or privacy. As a consequence, even if we cannot fully alleviate the risk of strategic bias, the selection bias is likely to be much

⁸The supply functions are always concave in spending, putting a cap on the performance of an ad with too high a budget.

⁹We express these measures in million users.

lower than in a sample of real advertising outcomes.

Second, the data provided by the API has one of the advantages provided by controlled field experiments: it allows us to compare performances of campaigns aimed at the exact same audience on iOS versus Android, thus alleviating any concerns relative to composition effects.

Third, the criticism formulated by practitioners regarding the imprecision of the prediction at the advertiser level is actually good news in the case of research.¹⁰ Indeed, this means that the API predictions are not too “account-dependent” and closer to general results than to the performances of an individual advertiser. They are much less “account-dependent”, it seems, than predictions on the simplified interface accessible through one’s Facebook account. Moreover, to address concerns regarding potentially inflated audience sizes that were raised in 2018¹¹, we compare in Table 1 audience size estimates given by the API to de-duplicated traffic numbers from SimilarWeb. The comparison does not indicate that these concerns are still valid in 2021.

Table 1: Estimated MAU produced by Facebook Marketing API compared to SimilarWeb

Date	MAU (Facebook)	Total Deduplicated MAU (SimilarWeb)
March 2021	145,301,412	191,750,266.33
April 2021	141,810,976	187,042,379.73
May 2021	137,612,432	196,331,274.95
June 2021	136,820,400	202,878,815.88
July 2021	138,302,816	208,025,326.07

Notes: Source for SimilarWeb: <https://pro.similarweb.com/>

¹⁰“Facebook’s estimations are based on complex algorithms and historical data, but they cannot account for the specific circumstances of each campaign.” See https://www.toolify.ai/ai-news/maximize-your-facebook-ads-uncover-estimated-daily-results-208566#google_vignette/

¹¹<https://www.cnbc.com/2021/02/18/facebook-knew-ad-metrics-were-inflated-but-ignored-the-problem-1.html>

Fourth, we believe that predictions provided by the Facebook Marketing API need to reflect real trends, as it is used by practitioners to decide whether to engage into an ad campaign *ex-ante*. This should motivate truthful prediction on META's side in the long run, as total inconsistency would have made the tool very quickly obsolete. As, the wave of criticism by practitioners in 2018 were followed by regular updates.¹² It seems hence that Facebook has been trying since then to make estimates closer to reality. This is supported, *inter alia*, by Grow et al. (2022) which validate statistically the quality of the data. Additionally, our results are in line with other studies by Aridor and Che (2024); Wernerfelt et al. (2022), with the perception of practitioners, as well as with META's official earning statements. They do not seem to be matching fully the best interest of META that could have hidden the effect altogether.

Finally, the aim of this study is not only to measure effects in absolute values but also to discuss the relative effects and strategic mechanisms at play. For the latter, the concern that predictions might be under or over-estimated should be less of a concern as long as the tendency is uniform.

2.4 Main Outcomes of Interest

With the data provided by Facebook Delivery Estimates based on ad performance, we build two indicators that measure *targeting efficiency* and *ad prices*. To be able to compare delivery estimates across audiences, we arbitrarily fix our daily budget to

¹²In March 2019, META declared that potential reach was based on how many people matching an advertiser's criteria had been shown an ad in the past 30 days. In March 2020, they also based the potential reach on the number of accounts used by each person. <https://www.ft.com/content/c144b3e0-a502-440b-8565-53a4ce5470a5> Additionally, since 2017 they have declared the number of fake accounts they delete from the platforms. <https://www.ft.com/content/98454222-fef1-11e9-b7bc-f3fa4e77dd47>

100 €¹³ and compute *impressions* and *actions* for each of the audiences. We compute the CR and CPM according to

$$CR_{100\text{€}} = \frac{Actions_{100\text{€}} \times 100}{Impressions_{100\text{€}}} \quad \text{and} \quad CPM_{100\text{€}} = \frac{100\text{€} \times 1,000}{Impressions_{100\text{€}}}$$

The CR denotes the proportion of *impressions* that trigger an *action* by consumers. We multiply it by 100 so that it could be expressed in percentage points. This variable measures targeting efficiency in our analysis. The CPM denotes the ratio between our arbitrary budget of 100 € and the number of *impressions*, expressed in thousands. It is a measure of the price of impressions.¹⁴

We use the CR and CPM as our main outcome variables to measure the effect of Apple’s new privacy policy on targeting efficiency and on the price of ads, respectively. Given the claims that privacy protection prevents firms from offering relevant advertising, we expect the *ATT* scheme to have a negative impact on conversion rates. The effect on prices, however, may be subject to different market effects. Indeed, a change in the conversion rate could have a direct quality effect on prices. It can change the proportion of demand between ads targeting iOS and Android users on META’s platforms. In fact, in the short run, Android registered a slight increase in ad spending compared to iOS.¹⁵

¹³The Delivery Estimates are not all provided for the same range of budgets. Hence, increasing the daily budget reduces the number of audiences for which the ad outcome estimates are available. On average per month, 63% of advertisers are likely to spend between 1\$ and 500 \$. <https://medium.com/@adamwilsonwebmaxy/how-much-does-a-facebook-ad-cost-f65039362bba>

¹⁴Note that with a fixed budget, what we will describe as a change in price is more precisely a shift in the supply curve of impressions.

¹⁵<https://www.forbes.com/sites/johnkoetsier/2021/05/19/ios-145-causes-massive-shift-in-ad-spend-to-android/?sh=bc2f63c20d74>

2.5 Data Collection

We collect daily data on estimated ad prices for target audiences in the US, from March 11, 2021 to July 11, 2021. Facebook Ad Manager provides a wide range of audiences. The advertiser can choose to include more or fewer criteria to define the target audience. These criteria include several socio-demographic characteristics such as age, gender, education, income, and user location. Facebook Ad Manager also provides individual interests from a list of several hundred predetermined interests, such as games or religion. An advertiser can also target users depending on their operating system. In the remainder of the paper, the word “*audience*” designates the aggregation of criteria at the exclusion of the OS, such that we observe each audience for iOS and for Android. Our requests use a combination of the audience criteria in Table 2, using one criterion per column, or not specifying (*Not spec.*) a criterion in the given category.

To use the API as in the real ad campaign, an advertiser must also specify the *optimization goal*, that is, the type of *action* that an impression should trigger. We chose to include *link clicks*¹⁶ and *app installs*. Each request should specify a *publisher platform*. We collect data for ads published on Facebook only (Fcbk) and both on Facebook and on the Audience Network (Fcbk&AN).¹⁷ Note that the API does not provide outcomes for ads exclusively displayed on the Audience Network.

Table 2: Targeting Criteria in Collected Audiences

Goal	Pub. Platform	Age	Gender	Education	Income	Interest/Behavior
App Install	Fcbk	18-34	Male	Bachelor degree	Top 10%	Games
Link click	Fcbk&AN	35-49 50-65	<i>Not spec.</i>	<i>Not spec.</i>	Top 10 to 25% Top 25 to 50% <i>Not spec.</i>	Fcbk page admin Religion <i>Not spec.</i>

¹⁶Link clicks is the default option in the API.

¹⁷The Audience Network is META’s open display channel, enabling print ads on third party mobile apps.

2.6 Descriptive Statistics

Our final sample includes 203,992 observations and 1,032 different audiences. We perfectly match a given audience in the treatment and the control group. For example, we compare the audience ad outcomes of an ad targeting a young men between 24 and 35 years using iOS with the same age and gender using Android at given point in time. To avoid biases in our estimates, we remove missing observations so that we obtain exactly the same set of audience-date couples in our treatment group (iOS users) and control group (Android users).¹⁸ Table 3 presents the descriptive statistics of our sample. Figure 1 displays the mean of our two dependent variables over time.

Table 3: Summary Statistics of the Overall Sample

Variable	Mean	Std. Dev.	Min	Max
CR_{100e}	0.58	0.35	0.055	2.82
CPM_{100e}	5.23	2.13	1.31	19.22
$Impressions_{100e}$	22479.78	9318.65	5202.68	76319.76
$NbAction_{100e}$	145.42	131.71	11.47	848.65
iOS	0.5	0.5	0	1
Pub. Platform Fcbk	0.499	0.5	0	1
App Install	0.497	0.5	0	1
DAU (in M)	1.68	2.65	0.044	21.65
MAU (in M)	2.11	3.45	0.001	30

¹⁸We also remove 9 dates in which, for technical reasons, the number of observations is significantly lower than the number of requests made. In the main analysis, 12 dates are removed between the before and the after period (see section 3.1. Our data includes 1,032 audiences multiplied by 2 OS, resulting in (102 days - 9 days that were removed - 12 days), which should account for 210,528 observations. However, as we have no predicted values for certain targets, we drop 6,536 observations in order to achieve a perfectly balanced panel.

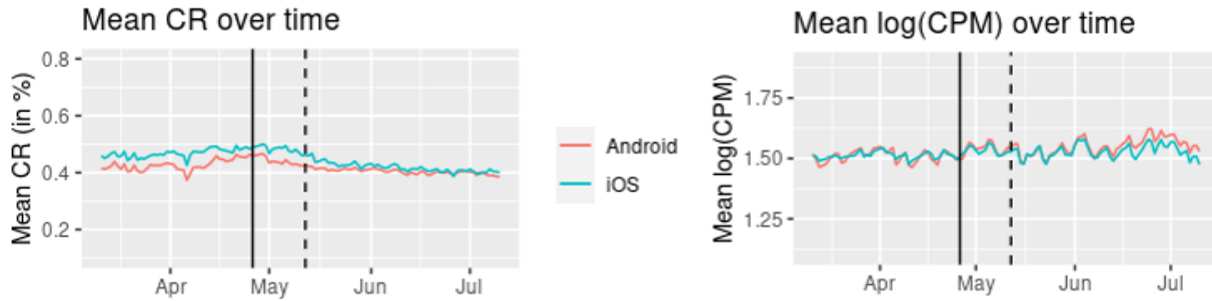


Figure 1: Mean CR and log(CPM) over time

3 Empirical Strategy

3.1 Econometric Identification

Our empirical strategy aims to estimate the effect of a change in the ability of an ad network to track consumers on ad outcomes. We exploit the introduction of Apple’s *ATT* policy. This policy should result in a negative exogenous shock to Facebook’s ability to aggregate data on iOS 14.5 and its later versions without affecting data aggregation on Android users. To obtain the causal effect of the *ATT* policy, we use a difference-in-differences design that exploits the fact that the policy applies to iOS users but is not expected to affect Facebook’s ability to target consumer audiences on Android (see Goldfarb et al. (2022) for methodological issues related to policy change).

Note that we claim identification of the causal effect on CR, as we do not detect any heterogeneous pre-treatment trend that might bias our estimates (Angrist and Pischke (2009), see Appendix B). However, we cannot make this claim for the effect on CPM, as some anticipation is seen in the price. Due to their highly strategic nature, prices are indeed affected by factors other than quality, including market concentration and (anticipatory) demand shifts. Nevertheless, given the relation in time around the introduction of the policy and the coherence between our results

and Facebook’s Q3 earning reports, we are confident that the observed changes are related. More generally, we consider throughout the paper that the output provided by the API concerning *impressions* and *reach* for a given budget are part of the strategy set of META, i.e. a proposition of supply. *Actions* on the contrary, or at least the likelihood that an *impression* will trigger an *action* are exogenously affected by tracking possibilities.

In the main analysis, our *treatment* group consists of all iOS target audience, and our *control* group consists of Android users. Because only a proportion of iOS users adopted the latest versions of the iOS and are thus exposed to the new privacy policy, our analysis measures Intention-To-Treat and we identify a lower bound for the treatment effect.¹⁹ To take into account the lag in adoption of the new OS, we introduce a time lapse of about two weeks between the *before* and *after* periods. The *after* period starts on May 12, on which approximately 13% of iOS users had adopted the iOS 14.5.²⁰

Note that with the *ATT* we do not share the identification concern expressed by other researchers regarding GDPR (Goldberg et al., 2024). Indeed, the simplicity of the Apple’s form and the 0-cost of a pro-privacy choice makes it very unlikely that exposition *per-se* to the choice set would modify consumers’ usage of the apps.

3.2 Baseline Empirical Strategy

The baseline equation estimates the change in our two main outcomes of interest: targeting efficiency and ad prices. In the main analysis, we compare the simplest specification to one including a set of fixed effects and controls. We use the following baseline specification. We denote o the OS, a denotes the audience, and t denotes the

¹⁹Note however that when varying the proportion of treated people (Appendix C.3), the magnitude of the effect does not change.

²⁰To check the robustness of this choice, see Appendix A.6.

observation date:

$$\mathbf{y}_{oat} = \alpha + \beta_1 \mathbb{1}_{\text{After}} + \beta_2 \mathbb{1}_{\text{iOS}} + \delta(\mathbb{1}_{\text{iOS}} \times \mathbb{1}_{\text{After}}) + \epsilon_{oat} \quad (1)$$

$\mathbb{1}_{\text{iOS}}$ is a dummy equal to 1 if the ad targets iOS users and 0 if it targets Android users. $\mathbb{1}_{\text{After}}$ is equal to 1 for requests made after May 12. Our robust standard errors are clustered at the audience level. We add to the baseline specification a set of control variables that include an audience-OS fixed effect, a day fixed effect, and the log of the average daily and monthly audience size, denoted respectively $\log(DAU)$ and $\log(MAU)$.

4 Baseline Results

4.1 Effect of the App Tracking Transparency

The results of the main estimates are presented in Table 4 in columns (1), (2), (5) and (6). Columns (1) and (2) report the estimates of the change in targeting efficiency proxied by the CR . The estimates suggest that the privacy policy change had a negative impact on targeting efficiency. More precisely, comparing the shift in conversion rate to the average conversion rate in the before-period on iOS indicates a decrease of 7.5% in Facebook’s ability to transform impressions into actions. In column (2), we include date and audience-OS fixed effects, as well as a control for the audience size. The R^2 improves greatly; the difference-in-differences coefficient, however, does not vary, suggesting strong robustness in our estimates.

Columns (5) and (6) display the evolution of relative ad prices— measured by the CPM—after the introduction of ATT . The result suggests that the introduction of the ATT was followed by a 10% drop of the relative price of ads targeted at iOS

users, compared to those targeted at Android users.

Note that the coefficients displayed in column (5), while not as robust as those of the two-way fixed-effect model, suggest—through the combination of the *After*, *iOS*, and *After*iOS* coefficients—that, even if the relative price of iOS decreased, there was a general increase in prices both for iOS and Android after the introduction of the *ATT*.²¹ The negative coefficient is therefore probably due to a transfer of demand from ads targeted at iOS users to ads targeted at Android users.²²

Table 6 in the Appendix allows us to decompose the effect on ad prices and targeting efficiency. This table reports the effect of Apple’s policy on impressions and actions offered for a budget of 100 €, mechanically increasing the CPM. We observe an increase in the number of impressions, in the number of people reached and in the number of actions. The result suggests the following dynamic: an upward shift in the supply curve of impressions, hence an increase in price, triggering an upward shift in actions, but of lower magnitude. It is interesting to see that the increased supply of impressions is not accompanied by an intensification of persuasion with a constant number of people reached, but by a diversification of the audience receiving the ad.

In column, (3), (4), (7) and (8), we split the sample between ad targeting using app install and link click. This allows us to study whether the change in privacy policy affects differently ads with optimization goals that are harder to achieve, namely installing an app versus clicking on a link. The results display significant heterogeneity between the two goals. We observe that the change in META’s ability to aggregate data has a negative effect on its ad effectiveness, mainly for ads aiming to

²¹Note that this is coherent in sign and magnitude—if we consider that ad space is relatively constant on Facebook—with the 12.3% growth in ad revenues in the US&Canada between Q1 & Q2-2021 (see Facebook’s earnings report mentioned in a previous footnote).

²²This seems to be corroborated by practitioners: <https://metrictheory.com/blog/apple-ios-14-5-update-facebook-ads-shifting-toward-android-devices/>, November 2023.

trigger *app installs*. The effectiveness of ads promoting *app installs* decreases sharply, by around 21%, while the ad effectiveness of ads looking for *link clicks* increases slightly by 1%, which seems close to insignificant economically speaking.

The limitation in the ability to aggregate data across apps seems to have had an impact only on ads promoting an action that is “more difficult” to trigger. This is probably because identifying consumers willing to download an app is more data-intensive and relies more on having access to different sources of data. Column (7) and (8) of Table 5 suggests a drop in the price of ads with both goals after the introduction of the *ATT*. It also indicates a greater price drop for ads aiming to trigger app installs. However, the heterogeneity in effect is lower than with the *CR*, all types of ads being negatively affected. This introduces one interesting fact: although we have previously seen a correlation between the drop in prices and the drop in targeting quality, it seems here that prices are not mainly driven by relative quality shifts on the platform.

Table 4: Difference-in-Differences Estimates for Targeting Efficiency and Prices

	CR				log(CPM)			
	(1) (All)	(2) (All)	(3) (App Install)	(4) (Link Click)	(5) (All)	(6) (All)	(7) (App Install)	(8) (Link Click)
After*iOS	-0.05*** [-0.06,-0.04]	-0.05*** [-0.06,-0.04]	-0.11*** [-0.13,-0.10]	0.01*** [0.01,0.02]	-0.10*** [-0.11,-0.09]	-0.10*** [-0.11,-0.09]	-0.13*** [-0.15,-0.12]	-0.06*** [-0.07,-0.06]
After	0.00 [-0.00,0.01]				0.15*** [0.15,0.16]			
iOS	0.14*** [0.11,0.16]				0.05*** [0.02,0.09]			
Constant	0.53*** [0.51,0.55]	0.58 [-0.24,1.40]	3.53*** [2.22,4.84]	-2.24*** [-2.62,-1.86]	1.49*** [1.46,1.51]	6.08*** [5.43,6.73]	8.53*** [7.49,9.57]	3.76*** [3.15,4.37]
Clustered SE	Y	Y	Y	Y	Y	Y	Y	Y
Audience-OS FE		Y	Y	Y		Y	Y	Y
Date FE		Y	Y	Y		Y	Y	Y
log(MAU)		Y	Y	Y		Y	Y	Y
log(DAU)		Y	Y	Y		Y	Y	Y
Mean on iOS before	0.67	0.67	0.52	0.81	/	/	/	/
Observations	203,992	203,992	101,318	102,674	203,992	203,992	101,318	102,674
Adjusted R2		2063.000	1031.000	1031.000		2063.000	1031.000	1031.000
R2	0.025	0.914	0.864	0.960	0.019	0.965	0.938	0.962

Note:

*p<0.05; **p<0.01; ***p<0.001; 5% Confidence-intervals
The robust standard errors are clustered at the Audience-OS level.
Treatment group: Facebook Access iOS device.
Control group: Facebook Access Android device.
CR are expressed in % and CPM in €.

4.2 The Negative Effect on Targeting Efficiency and Price Strengthens over Time

To check the robustness and better understand the dynamics of our effect, we regress the difference in outcomes between iOS and Android on the set of date dummies. More formally, we perform:

$$CR_{iOS,at} - CR_{Android,at} = \beta \mathbb{1}_t + \epsilon_{at} \quad (2)$$

$$\log(CPM)_{iOS,at} - \log(CPM)_{Android,at} = \beta \mathbb{1}_t + \epsilon_{at} \quad (3)$$

In the estimation, we add an audience-OS fixed-effect as well as controls for the audience sizes ($\log(DAU)$ and $\log(MAU)$) Figure 2 plots the resulting time trend coefficients. At first glance, we see that the results are in line with the aggregate effect found in the previous analysis. After the introduction of the policy, the targeting ability of Facebook on iOS sharply decreases compared to that on Android. This translates into a difference in CPM that also increases over time, but in several successive phases. There seems to be some anticipation regarding the evolution of relative prices. However, there is, as for targeting efficiency, a visible difference between the average coefficients before and after. At least two possible effects are at play to rationalize such dynamics. First, the adoption of versions 14.5 and above increases over the time-period (see adoption rate in the Appendix A.2). Although this may be somewhat mitigated by the increase in the tracking opt-in rate,²³ the proportion of effectively treated people in the pool of iOS users increases over time. Hence, our time dummies capture the increased intensity of treatment in the intended-to-treat group. As expected, this translates into a stronger effect as more

²³ATT opt-in rate, Flurry Monthly updates. <https://www.flurry.com/blog/att-opt-in-rate-monthly-updates/>, March 2023.

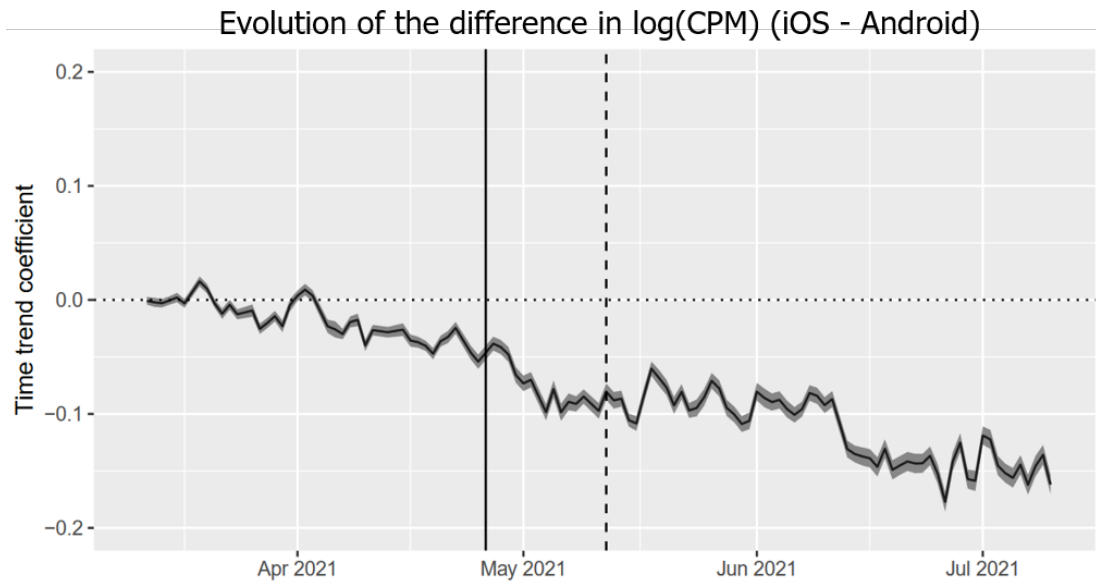
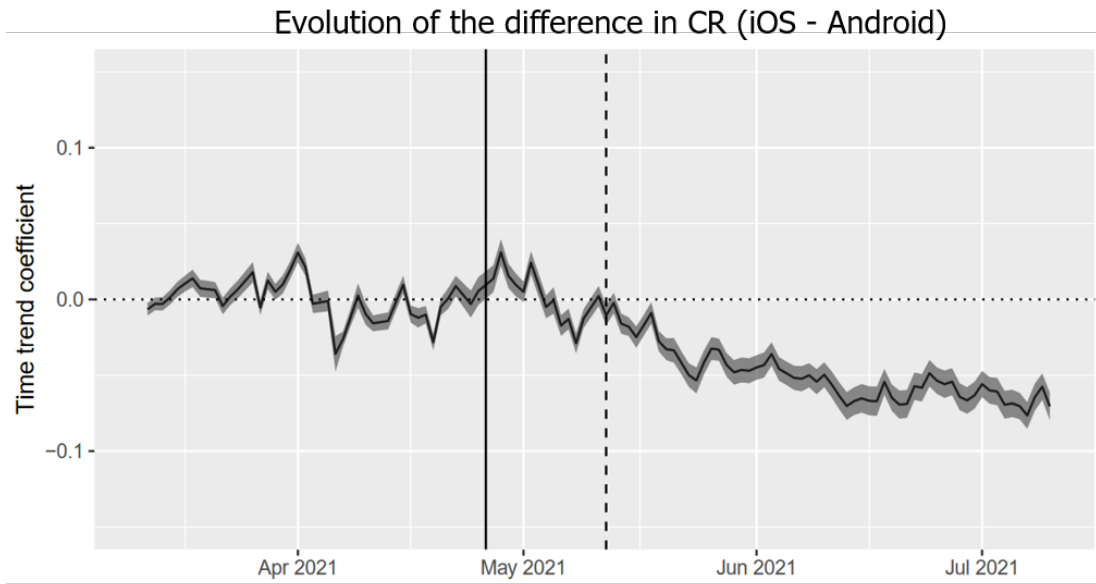


Figure 2: Evolution of the difference (iOS-Android) in CR and $\log(CPM)$. Estimates include OS-audience fixed-effects

people adopt the privacy-friendly OS. Second, the negative trend in price difference also captures the demand and expectations of advertisers regarding the negative effect of the *ATT* on targeting efficiency. The trend uncovers some elasticity of demand, most probably between ads targeted at Android versus ads targeted at

iOS.

5 Heterogeneity of the Effect

In this section, we explore how the change in the ability to aggregate consumer data may have affected audiences differently. First, thanks to the heterogeneity in audience sizes, we investigate whether the effect observed may be attributed only to measurement issues or captures some real effect. Second, we look at the heterogeneity of the effect with respect to the ad placement, comparing ads displayed on Facebook only versus ads displayed both on Facebook Social Network and on the Facebook Audience Network. It can display ads in mobile apps using the same targeting and measurement tools as those used on Facebook. Table 5 shows all estimates.

5.1 Real Effect or Measurement Issue?

One of the concerns associated to the *ATT* is that the impossibility to track users would not only impair consumer profiling but also the ability of ad networks to technically track conversions. We look at the heterogeneity of the effect with respect to audience sizes. We believe that there is no reason for the ability to observe conversions to depend on the size of the target audience. As a consequence, any heterogeneity measured confirms the fact that we capture, at least in parts, a real effect on the CR.

In columns (1) and (3) of Table 5, we interact our difference-in-differences coefficient with a factor variable classifying the audience into five quantiles, according to their size *DAU*). The first quantile contains the smallest audiences, whereas the fifth quantile contains the largest audiences. Column (1) shows that the effect on the

conversion rate is negative overall, whatever the size of the audience but displays some heterogeneity. This strengthens our claim that we capture some real change in the CR.

Surprisingly though, the introduction of *ATT* seems to have had a more intense effect on targeting efficiency towards large audiences than on targeting efficiency towards small audiences. A potential explanation is that despite the *ATT*, it is still possible to target a precise audience through contextual advertising, where the ad content is based on the displayed content. Indeed, while *ATT* affects tracking-induced targeting, it is not supposed to affect contextual targeting.²⁴ On the other hand, ads with broad audiences can hardly be matched with the right individuals via contextual ads. Indeed, a wide audience will contain a much lower density of relevant users for a given ad. One complementary conjecture is that third party data is used, not only to identify the audience requested, but also to identify who, in this audience, is more likely to generate a conversion. The broader the target audience, the greater the demand for data aggregation. Column (2) and (4) of Table 5 shows that the decrease in ad prices follows the same logic. It is more important for audiences that seem to have a larger decrease in quality.

5.2 Are ads Distributed on Facebook only more or less Affected?

In our data, ads can be distributed either only on Facebook (Fcbk) or both on Facebook and Audience Network²⁵ (Fcbk&AN), META's open display ad network. In both cases, only first-party data remains accessible if users opt out of tracking. However, the nature of this first-party data is very different on the two platforms. For

²⁴Practitioners provide insights into this intuition. <https://www.adexchanger.com/mobile/how-facebook-is-overhauling-its-attribution-standards-to-deal-with-apples-att/>, November 2023.

²⁵The option to display an ad only on the Audience Network was not available on META's API during the data collection process.

ads on Facebook, logged-in user data can still be used for targeting. For ads on the Audience Network, only contextual data is easily available if users opt out of tracking. This may translate into heterogeneous effects, depending on which type of data remains available. Columns (3) and (6) of Table 5 introduce this variable.

Note that the API estimates for ads published on Fcbk&AN do not disclose the splitting rule between the two publisher platforms. Hence, these estimates have to be interpreted as weighted averages of outcomes on each of the two platforms separately. Consequently, the comparison in treatment effect between ads published on Fcbk versus on Fcbk&AN is hard to interpret. Only the sign is indicative, as the magnitude depends on the unobserved share of ads displayed on the Audience Network in the joint Fcbk&AN sample. Moreover, not observing any difference in the treatment effect is hard to interpret as well. It can reflect a true homogeneity in effects, or/and a small weight put on the Audience Network in the joint outcomes.

In Column (3) of Table 5, we interact the difference-in-differences term with a dummy equal to 0 for ads on Fcbk, and 1 if the predictions requested are for ads displayed on Fcbk&AN. We observe no difference in the effect on targeting efficiency between the two display options. However, Column (4) shows that the price of ads on Fcbk&AN seems to be more affected than the price of ads on Fcbk. This second non-zero difference in effect also indicates that the Audience Network is sufficiently represented in the joint Fcbk&AN statistic provided by the API. Observing no difference in targeting efficiency between Fcbk and Fcbk&AN can be surprising. Indeed, this result suggests that ads on the Audience Network are not more affected than ads displayed on Facebook. At least two rationales explain this. First, ads on Facebook may be more affected than we think. Indeed, despite its already prolific data-generating social networks, META is among the companies collecting

the most third-party data.²⁶ Interestingly, our result seems to indicate that this complementary third-party data plays an important role in the social network's ability to build relevant advertising profiles. Second, ads on the Audience Network may be less affected. Indeed, the use of contextual advertising and alternative targeting techniques may replace tracking better than expected.

²⁶See CMA's *Online platforms and digital advertising market study* (2020), Appendix G https://assets.publishing.service.gov.uk/media/5fe49554e90e0711ffe07d05/Appendix_G_-_Tracking_and_PETS_v.16_non-confidential_WEB.pdf, March 2023

Table 5: Results of triple Difference-in-Differences

	CR		log(CPM)	
	(1)	(2)	(3)	(4)
After*iOS	-0.02** (-0.03, -0.01)	-0.05*** (-0.07, -0.04)	-0.08*** (-0.09, -0.07)	-0.09*** (-0.09, -0.08)
After*iOS*quantile-2 <i>(smallest)</i>	-0.01 (-0.03, 0.01)		-0.02* (-0.04, -0.001)	
After*iOS*quantile-3	-0.03** (-0.05, -0.01)		-0.03** (-0.05, -0.01)	
After*iOS*quantile-4	-0.03** (-0.06, -0.01)		-0.03** (-0.05, -0.01)	
After*iOS*quantile-5 <i>(largest)</i>	-0.08*** (-0.10, -0.05)		-0.05*** (-0.07, -0.03)	
After*iOS*Fcbk&AN		0.01 (-0.01, 0.02)		-0.03*** (-0.04, -0.01)
Clustered SE	Y	Y	Y	Y
Audience-OS FE	Y	Y	Y	Y
Date FE	Y	Y	Y	Y
log(MAU)	Y	Y	Y	Y
log(DAU)	Y	Y	Y	Y
Observations	203,992	203,992	203,992	203,992
Adjusted R ²	0.92	0.91	0.97	0.97

Note:

*p<0.05; **p<0.01; ***p<0.001

The robust standard errors are clustered by audience.

Treatment group: Facebook Access iOS device.

Control group: Facebook Access Android device.

5.3 Discussion

Only a limited empirical stream of the literature has measured the impact of privacy regulations on targeting efficiency. This makes it hard to assess the economic significance of our results. Goldfarb and Tucker (2011) find a 65% reduction in ad effectiveness after the introduction of the EU Privacy Directive in 2002. Compared to our estimates, this coefficient is much larger.

Several characteristics of our study and our data contribute to explaining the low effect we find on targeting efficiency. First, we consider as conversions actions such as link click or app installs, which are not as demanding as a purchase for example. As a result, least efficiently targeted consumers could keep interacting with the ad but not necessarily go on with a purchase. Complementary studies try to assess the impact of *ATT* on purchases, like that of Aridor and Che (2024). Second, our treatment groups do not include 100% of treated people. Some iOS users may have agreed to be tracked.²⁷ Also, our data pools different iOS target audiences, including users of older versions; thus, formally, our main estimates are lower bounds for the true effect. To evaluate how the effect evolves with a higher proportion of treated users in the sample, we conduct a complementary analysis in Appendix C.3 Table 7. We conduct a robustness check by fixing our *after*-cutoff to June 26, when the combined adoption of iOS 14.5 and 14.6 reached approximately 70%. Although the result is mitigated by a limited amount of *after*-data, we find that the magnitudes of coefficients do not vary. The negative effect on targeting efficiency is 10.5%, while the relative drop in price reaches 14%. However, in absolute value, both prices on iOS and Android seem to have increased. Finally, the low magnitude could also be attributed to META's dominance in the market. This idea is reinforced by the global

²⁷Between 12 and 25%, according to industry estimates (<https://9to5mac.com/2021/07/23/app-tracking-transparency-opt-in-snap/>, September 2021.)

increase in CPM that we observe. Indeed, even if hit in terms of targeting quality, META managed to keep a competitive advantage on the ad market and somehow reinforced it in the short run.

There are at least three reasons for which META could be less negatively affected than other display networks, or even positively affected in the short run by the *ATT*. First, it generates (through its own user interface) a substantial amount of data that is linked to a user profile. This profiling has no reason to be affected by Apple's privacy policy. Second, cookie syncing, which could make up for the inability to access a user's IDFA, is subject to large network effects (Jeon, 2021). As a consequence, META is more able than other actors of the industry to find alternative profiling methods. Third, we consider a short time period around the change in policy. This allows us to not include effects arising from other events that may happen during the year. However, META's user data is likely to be of sufficient quality such that it will not be completely obsolete in such a short time period. And ad effectiveness may keep decreasing as time goes by. This idea is reinforced by the different announcements META has made in 2022 regarding lost opportunities because of the *ATT*.²⁸ To summarize, the effects we measure should be a lower bound for the effect that *ATT* may have on other actors in the industry. Other studies could investigate the change in market power of different actors in the industry in the long run, especially looking at whether the initiative reinforced the oligopolistic situation of the biggest players in the ad industry.

²⁸<https://www.businessinsider.com/facebook-blames-apple-10-billion-loss-ad-privacy-warning-2022-2?op=1&r=US&IR=T>, November 2023.

6 Conclusion

This paper contributes to understanding the impact of privacy policies on ad effectiveness and ad prices on a social media platform. We study the introduction of the *ATT* on the iOS14.5+, which aims to give consumers a choice about being tracked by apps. We use a novel dataset collected through Facebook’s Marketing API. We find that the introduction of the *ATT* scheme by Apple has had a negative impact on Facebook’s ability to efficiently target consumers and trigger conversions. This change was concurrent with a decline in the relative price of ads but with an increase in the overall price. This suggests that the loss in targeting efficiency is followed by a readjustment in advertisers’ demand or willingness-to-pay. However, despite the loss in targeting efficiency, the global increase in ad prices suggests that Facebook has kept part of its competitive advantage. The impact of the policy strengthens with time for ads targeted to the pooled iOS group, in line with an increased proportion of users adopting the latest versions of the iOS. Surprisingly however, we find that the effect on ad effectiveness and the following adjustment of ad prices weakens with the size of the target audience. This leads us to believe that contextual advertising compensates better at the end of tracking for very well defined audiences. The effect is also stronger for ads promoting app installs compared to ads promoting link clicks. This suggests that although simple advertising with easy-to-get conversions is not much threatened, Apple’s new scheme may strongly affect the ability of firms to develop more sophisticated ad campaigns with actions from consumers that are harder to trigger.

There are limitations to this research. First, our data measures only the effect of privacy policy change of one of the biggest ad platform worldwide. It should be complemented with studies on smaller actors. Second, our estimates measure the

immediate effects of this policy change and do not consider long-term effects.

Funding and Competing Interests

The financial support of the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 670494 & grant agreement No. 759733 - PLATFORM) and the financial support of Carnot TSN are also gratefully acknowledged. Disclaimer: Any opinions and conclusions expressed on this articles are those of the authors and do not necessarily represent the views of their institutions.

References

- Ahmadi, Iman, Nadia Abou Nabout, Bernd Skiera, Elham Maleki, and Johannes Fladenhofer**, "Overwhelming targeting options: Selecting audience segments for online advertising," *International Journal of Research in Marketing*, 2023.
- Alcobendas, Miguel, Shunto Kobayashi, and Matthew Shum**, "The Impact of Privacy Measures on Online Advertising Markets," *Available at SSRN 3782889*, 2021.
- Angrist, Joshua D and Jörn-Steffen Pischke**, *Mostly Harmless Econometrics: An empiricist's companion*, Princeton university press, 2009.
- Aridor, Guy and Yeon-Koo Che**, "Privacy Regulation and Targeted Advertising: Evidence from Apple's App Tracking Transparency," 2024.
- , —, and Tobias Salz**, "The Economic Consequences of Data Privacy Regulation: Empirical Evidence from GDPR," Technical Report, National Bureau of Economic Research 2020.
- Bian, Bo, Xinchun Ma, and Huan Tang**, "The Supply and Demand for Data Privacy: Evidence from Mobile Apps," *Available at SSRN*, 2021.
- Cheyre, Cristobal, Benjamin T Leyden, Sagar Baviskar, and Alessandro Acquisti**, "The Impact of Apple's App Tracking Transparency Framework on the App Ecosystem," *Available at SSRN 4453463*, 2023.
- Choi, W Jason, Kinshuk Jerath et al.**, "Privacy and Consumer Empowerment in Online Advertising," *Foundations and Trends in Marketing*, 2022, 15 (3), 153–212.

- Goldberg, Samuel G, Garrett A Johnson, and Scott K Shriver**, “Regulating privacy online: An economic evaluation of the GDPR,” *American Economic Journal: Economic Policy*, 2024, 16 (1), 325–358.
- Goldfarb, Avi and Catherine E Tucker**, “Privacy Regulation and Online Advertising,” *Management Science*, 2011, 57 (1), 57–71.
- **and Verina F Que**, “The Economics of Digital Privacy,” 2023.
- **, Catherine Tucker, and Yanwen Wang**, “Conducting Research in Marketing with Quasi-Experiments,” *Journal of Marketing*, 2022, 86 (3), 1–20.
- Grow, André, Daniela Perrotta, Emanuele Del Fava, Jorge Cimentada, Francesco Rampazzo, Sofia Gil-Clavel, Emilio Zagheni, René D Flores, Ilana Ventura, and Ingmar Weber**, “Is Facebook’s advertising data accurate enough for use in social science research? Insights from a cross-national online survey,” *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2022, 185 (Supplement 2), S343–S363.
- Jeon, Doh-Shin**, “Market Power and Transparency in Open Display Advertising—A Case Study,” 2021.
- Johnson, Garrett**, “Economic Research on Privacy Regulation: Lessons from the GDPR and Beyond,” 2022.
- Johnson, Garrett A, Scott K Shriver, and Shaoyin Du**, “Consumer Privacy Choice in Online Advertising: Who Optes Out and at what Cost to Industry?,” *Marketing Science*, 2020, 39 (1), 33–51.
- Kesler, Reinhold**, “The Impact of Apple’s App Tracking Transparency on App Monetization,” *Available at SSRN 4090786*, 2022.

Kraft, Lennart, Bern Skiera, and Tim Koschella, “Economic impact of opt-in versus optout requirements for personal data Usage: The case of Apple’s app tracking transparency (ATT),” Technical Report, Working Paper 2023.

Laub, René, Klaus M Miller, and Bernd Skiera, “The Economic Value of User Tracking for Publishers,” *Available at SSRN 4251233*, 2022.

Li, Ding and Hsin-Tien Tsai, “Mobile Apps and Targeted Advertising: Competitive Effects of Data Exchange,” *Available at SSRN 4088166*, 2022.

Marotta, Veronica, Vibhanshu Abhishek, and Alessandro Acquisti, “Online Tracking and Publishers’ Revenues: An Empirical Analysis,” in “Workshop on the Economics of Information Security” 2019.

Ravichandran, Deepak and Nitish Korula, “Effect of Disabling Phird-Party Cookies on Publisher Revenue,” *Google Report*, 2019.

Wernerfelt, Nils, Anna Tuchman, Bradley Shapiro, and Robert Moakler, “Estimating the Value of Offsite Data to Advertisers on Meta,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2022, (114).

Appendix

A Setting and Data

A.1 App Tracking Transparency screenshots

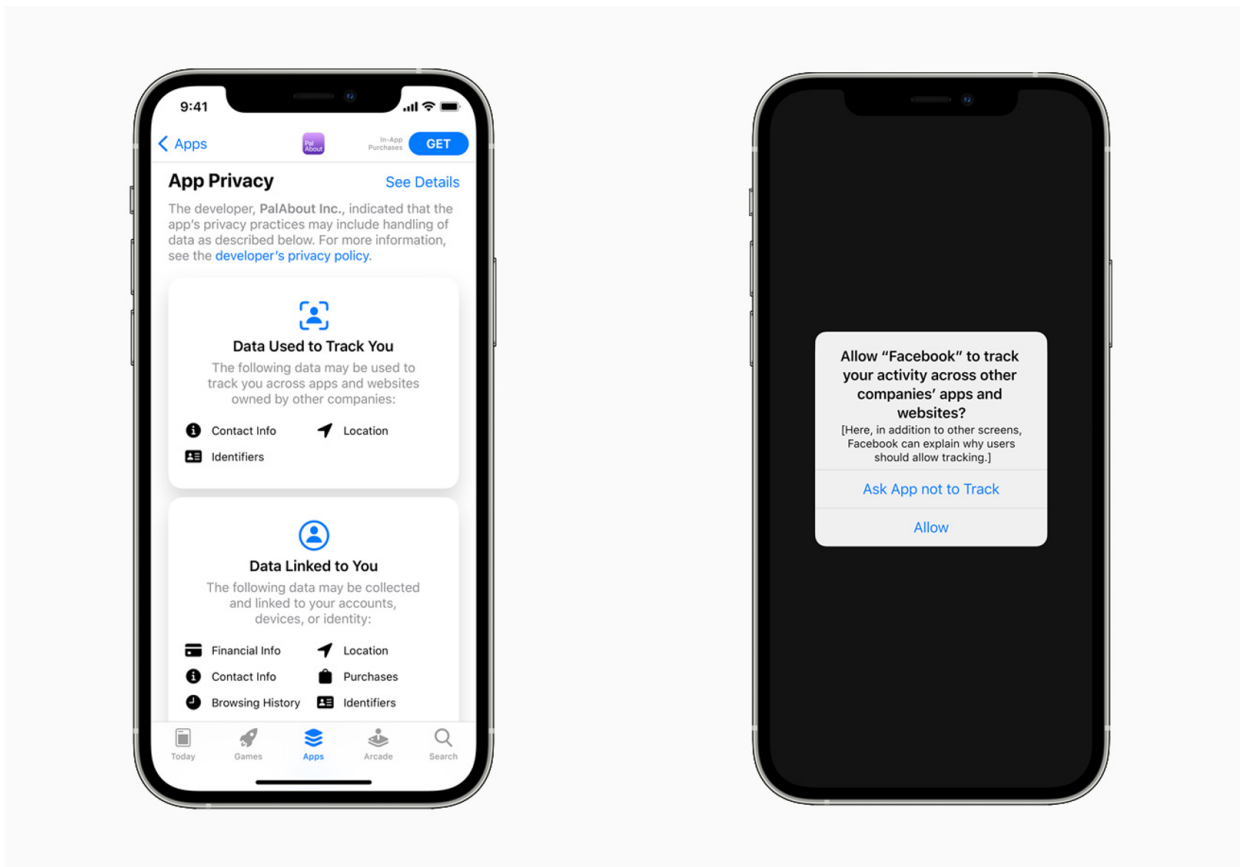


Figure 3: Apple ATT new features - Screenshots

A.2 Adoption of iOS14.5+

Figure 4 shows the evolution of number of iOS users in the US from March 2021 to July 2021.


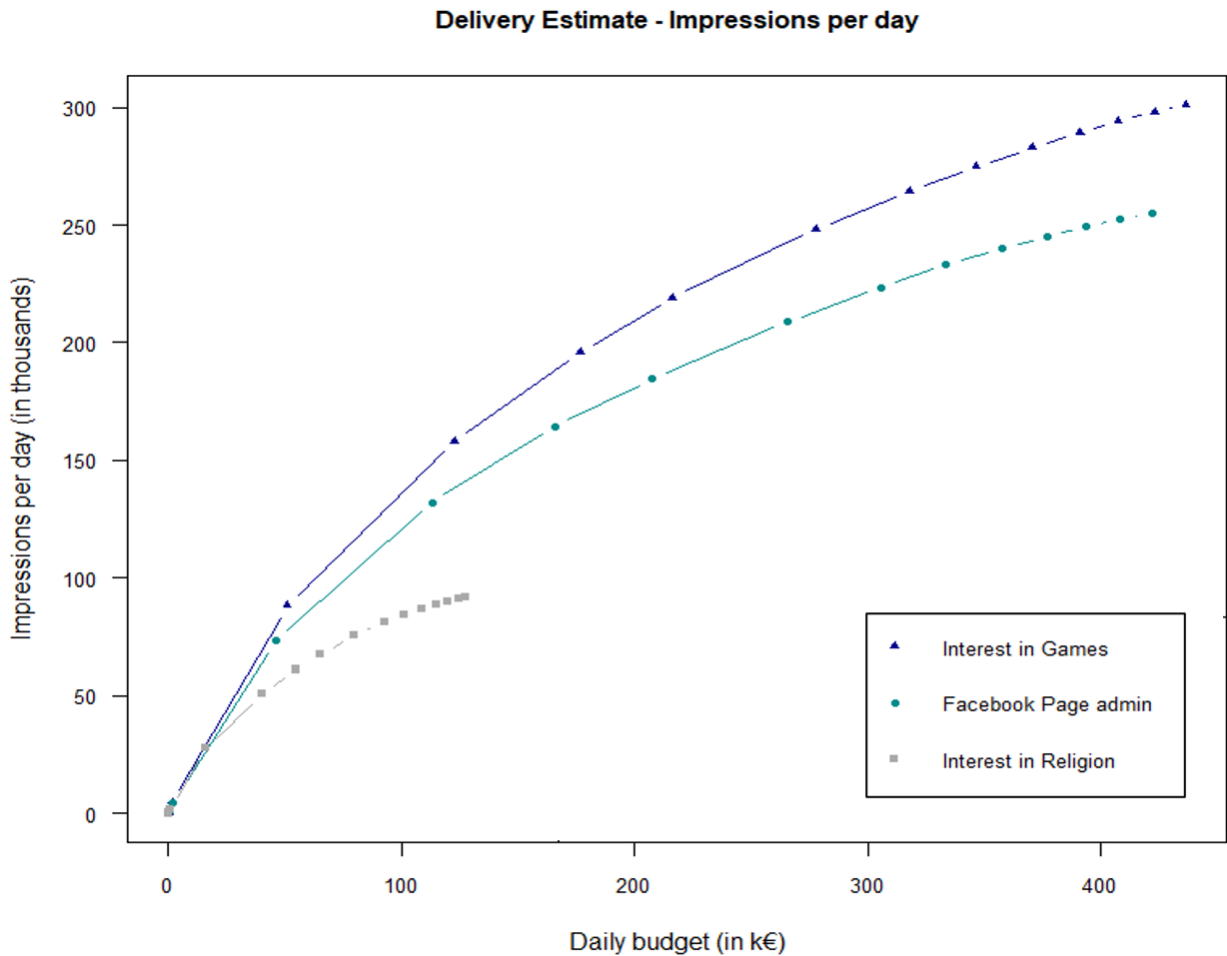


Figure 4: Proportion of Versions 14.5 and 14.6 among iOS Users during the Studied Time-Period

A.3 Sample Delivery Estimate Output

Figure 5 illustrates the curves of the delivery estimate output based on impressions for a given set of audience with a given daily budget.



Common audience characteristics: US, males, 18-34, top 10% income, bachelor degree, iOS users, publisher platform = Facebook, Optimization goal = app installs

Figure 5: Example of Delivery Estimate Output for Impressions, with three different Audiences.

B Parallel trends

The difference-in-differences methodology relies on the assumption that there are non-significant differences in pre-treatment trends. Figure 6 illustrates the trends in CR ratios, indicating that the two lines follow a parallel trend before the implementation of the *ATT*. We run the parallel trend test on a subsample to get a perfectly balanced panel. This subsample consists of precisely the same audience in both the treatment and control groups every day. The assumption holds as the statistical test does not reject the null hypothesis of parallel trends. There are no significant differences in the pre-treatment period by p-values > 0.1 (F-Stat: Parallel trends, 0.47).

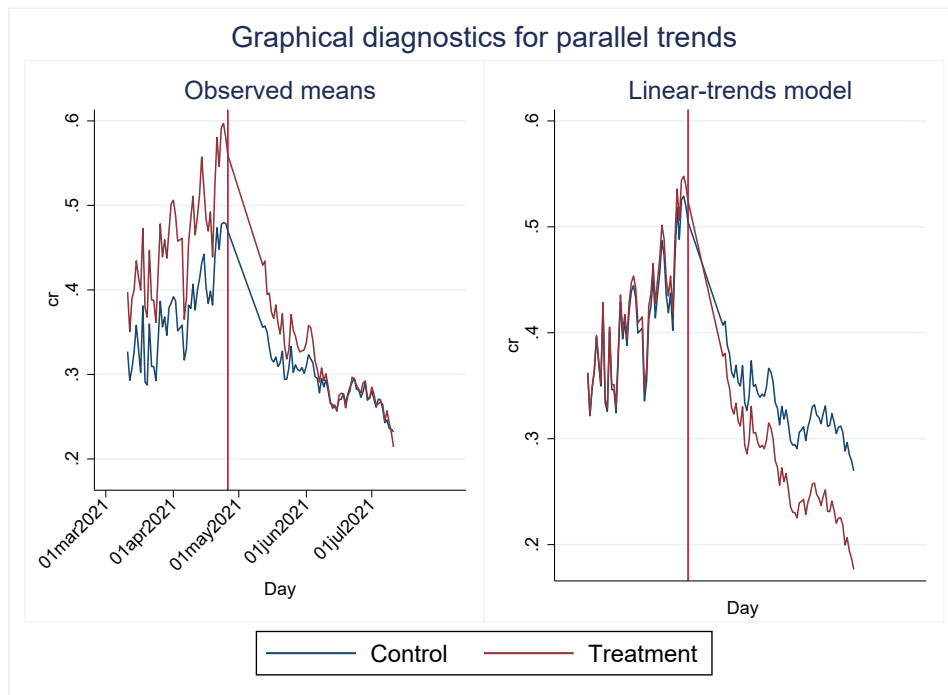


Figure 6: Graphical Diagnostic for Parallel Trends

C Complementary Analysis

C.1 Effect on Impressions and Actions

Table 6 presents estimates where we use alternative dependent variables, namely *Impressions*, *Actions*, and *Reach*.

Table 6: Result on Impressions, Actions and Reach

	<i>Dependent variable:</i>		
	Impressions (1)	Actions (2)	Reach (3)
After*iOS	2,110.16*** [1,892.63; 2,327.69]	2.95** [1.17; 4.73]	718.38*** [619.80; 816.96]
Clustered SE	Y	Y	Y
Audience-OS FE	Y	Y	Y
Date FE	Y	Y	Y
log(MAU) Control	Y	Y	Y
log(DAU) Control	Y	Y	Y
Mean on iOS before	24,634	144	9,213
Observations	203,992	203,992	203,992
Adjusted R ²	0.95	0.97	0.97

Note:

*p<0.05; **p<0.01; ***p<0.001

C.2 Robustness of *After* date

We check that the chosen *After* date does not make a significant difference on the baseline results. We run a series of regressions in which the *After* cut-off varies between May 8 and May 16. We plot the resulting difference-in-differences coefficients in Figure 7 and observe that they remain nearly unchanged as the *After* cut-off shifts.

Difference-in-Differences coefficient, depending on *After* cut-off



Figure 7: Results of Multiple Regressions with a different *After* Cut-off points

C.3 Analysis with High Adoption Rate of iOS 14.5-14.6

To assess how the average effect measured depends on the rate of new OS adoption in the pool of iOS users, we conduct a robustness check by fixing our *after*-cutoff to June 26, when the combined adoption of iOS 14.5 and 14.6 reached approximately 70%.²⁹ Columns (1) and (2) of Table 7 shows the estimates. Figure 8 shows that on this date, the income and gender compositions of the different versions of the iOS are similar. The coherence of the results alongside widespread adoption and consistent audience composition helps alleviate concerns regarding potential selection biases.

Table 7: Regression on CR and log(CPM), after June 26

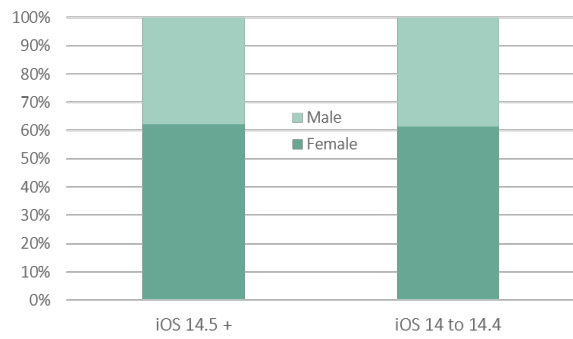
	<i>Dependent variable:</i>			
	(1)	CR (2)	log(CPM) (3)	(4)
iOS*After _{.26th}	-0.07*** (-0.08, -0.06)	-0.07*** (-0.08, -0.06)	-0.14*** (-0.14, -0.13)	-0.14*** (-0.15, -0.13)
iOS	0.14*** (0.13, 0.15)		0.05*** (0.05, 0.06)	
After _{.26th}	0.004 (-0.004, 0.01)		0.21*** (0.20, 0.22)	
Clustered SE	Y	Y	Y	Y
Audience-OS FE		Y		Y
Date FE		Y		Y
log(DAU) control		Y		Y
log(MAU) control		Y		Y
Observations	112,122	112,122	112,122	112,122
Adjusted R ²	0.03	0.90	0.03	0.96
Residual Std. Error	0.34	0.11	0.40	0.08

Note:

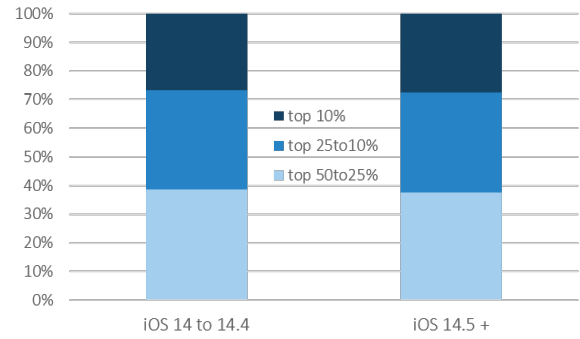
The standard errors are clustered at the audience level.

*p<0.05; **p<0.01; ***p<0.001

²⁹It goes up to more than 75% at the end of our time-period. <https://mediamattersww.com/uncategorized/the-impact-of-ios14-5-and-app-tracking-transparency-att/>, November 2023.



(a) Gender Composition



(b) Income Composition

Figure 8: Sample composition iOS 14 to 14.4 vs. iOS 14.5+ on June 26, 2021, built with Facebook's DAU.