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

Grazia Cecere[†] Sarah Lemaire [‡]

November 2021

PRELIMINARY VERSION

Abstract

We investigate the effect of privacy protection rules on targeting efficiency and ad prices. We use a change in Apple's privacy policy the *App Tracking Transparency* option, as a natural experiment. Introduced in Spring 2021, this new feature of the iOS 14 aims at providing smartphone users with more control over their data. It requires app developers to request explicit permission to track users beyond the app in use. However, this privacy policy is expected to reduce ad effectiveness on mobile devices. To assess the effect of the policy, we use an original database of estimated Facebook ad performances in the US market. We compare the outcomes of ad campaigns targeting iOS users versus Android users. The results suggest a reduction in targeting efficiency and ad prices for ads aimed at iOS users compared to Android users.

[*]We thank Alexandre de Cornière, Pierre Dubois, Isis Durrmeyer, Daniel Ershov, Renato Gomes, Ulrich Hege, Christian Hellwig, Doh-Shin Jeon, Nour Meddahi, Charles Pebereau, Wilfried Sand-Zantman for their helpful comments. The financial support of the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 670494) is also gratefully acknowledged.

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

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

1 Introduction

Firms collect huge amounts of individual data. These data can be used to enable targeted advertising and offer subsidized or free consumer services on the internet. Some of the currently most successful companies have adopted this business model. However, data collection has raised significant privacy concerns and resulted in regulatory initiatives. These privacy regulations aim at protecting consumers. Yet, there is a concern that they may also have an adverse effect in the long run. In particular, limiting data collection prevents platforms from making money with targeted advertising. In this case, regulations may hinder the development of qualitative internet services. This paper examines one side of this trade-off. More precisely, it intends to shed light on the impact of a change to privacy regulations on targeting efficiency and ad prices.

Regulators are not the only ones to engage in privacy protection. Digital companies are increasingly trying to attract consumers with privacy-friendly products or services. In April 2021, Apple introduced a feature on its new Operating System (OS), iOS 14.5 offering more control to consumers: the App Tracking Transparency. The scheme requires all apps to collect users' explicit consent before tracking them, with a simple yes-or-no form. In the same way, Google plans to ban third-party cookies on Chrome by 2023. By setting privacy as the default option, these initiatives are likely to affect the market for ads much more. Preventing the systematic aggregation of data should limit the precision of users' information. It may, thus, drastically reduce the profitability and effectiveness of targeted advertising. This is likely to threaten the survival of many firms and block the provision of some free services. In the app market, many apps are free and rely on their ability to collect high-frequency data. The welfare effects of a massive reduction in user tracking are not obvious. In particular, earlier regulations have not had the intended widespread impact. Aridor et al. (2020) indeed find that the introduction of the GDPR resulted in only a 12.5% drop in total cookies. Johnson et al. (2020) find that only 0.23% of ad impressions arise from consumers who opted out of behavioral advertising through the self-regulatory AdChoice program. This paper investigates the effect of Apple's App Tracking Transparency scheme, which could affect almost 50%[1] of mobile users in the U.S. It addresses questions related to the welfare effects of the policy such as:

[1]According to Kantar, iOS had a market share of 46.8% in the U.S. in September 2021 (`https://www.kantar.com/campaigns/smartphone-os-market-share`).

What is the effect of a change in tracking ability on ad efficiency? How is this effect reflected in relative prices? Are all firms likely to be affected in the same way?

To measure the impact of Apple's new privacy policy on targeting efficiency and prices, we use a database of Facebook ad performances. Our empirical design exploits the fact that the privacy option is offered to iOS users but not to Android users. As a consequence, the ability of firms to profile Android users should not change. Hence, we use a difference-in-difference design, using ad targeting Android users as the control and ad targeting iOS users as treatment. We interpret any relative evolution in ad effectiveness or ad price as the consequence of an equilibrium effect. This equilibrium effect is a combination of the effect of the policy on data quality and a strategic effect. We use data collected through Facebook's Marketing API (Application Programming Interface) which offers daily estimates of ad performances, depending on a specified audience. To build audiences, we use demographic criteria associated with an operating system namely iOS or Android.

We find that the introduction of Apple's privacy scheme has decreased the effectiveness and prices of ads targeted at iOS users. Ads targeted at iOS users triggered 8% fewer actions per impression after the introduction of the App Tracking Transparency policy. In the same way, their price per impression fell by 10%. We find that the effect strengthens over time, which is in line with the users' increasing o the perating system updates. The effect also intensifies for audiences that are harder to identify.

Our article contributes to two streams of literature. First, it contributes to the empirical literature documenting the effects of privacy regulations on ad effectiveness. Goldfarb and Tucker (2011) study the effects of the implementation of the EU ePrivacy Directive. They measure ad effectiveness with stated purchase intent after exposure to an ad. The paper provides empirical evidence that the policy has reduced ad-effectiveness by 65%. Our work complements this analysis as we measure ad effectiveness with a user action-per-impression rate. This rate may not generate revenues as directly as a purchase would; however, it has the advantage of counting effective actions and not intended actions.

Second, our article contributes to the literature measuring the value of ad tracking. Ravichandran and Korula (2019) use Google Ad Managers to estimate the effect of disabling third-party cookies. They conducted a randomized field experiment with the top 500 global publishers. In this experiment, they disable access to third-party cookies for users in the treatment group. They find that average publisher revenues decrease by 52% when deprived of cookies. In the same way, Johnson et al. (2020) find a strong negative effect of privacy restrictions. They study the impact of the Ad Choices self-regulatory initiative. They find that the revenue generated by opt-out consumers drops by 52%. Alcobendas et al. (2021) also find a rather strong negative effect of self-regulation. Their article uses bidding data from Yahoo with a structural model. They assess the impact of Google's announced plan to block third-party cookies on the internet browser Chrome in 2022 showing that the ban would reduce publisher revenue by 30%. Conversely, Marotta et al. (2019) show a small effect of cookies on ad revenues. With data from a large media company, they find that when users' cookies are available, a publisher's revenue increases by only 4%. All these papers measure the effect of privacy restrictions on ad revenues, while we intend to measure the joint effect on ad effectiveness and ad prices.

To the best of our knowledge, our article is the first empirical analysis to estimate the effect of the *App Tracking Transparency* policy. This may be due to the difficulty in obtaining suiteable data. This article is also the first empirical investigation studying a privacy policy affecting ads on apps. Moreover, our paper contributes in three ways to the literature on the effect of privacy policies. First, *App Tracking Transparency* policy provides a good opportunity to investigate the effects of an opt-in policy, while a large part of the literature looks at opt-out privacy policies. It is likely to affect the market for ads differently. Indeed, opt-in policies offer privacy by default to consumers and do not request them to make an extra effort to protect their privacy. The share of consumers likely to be effectively affected by the change will probably be higher than in other initiatives. Second, as we use estimated ad performances on Facebook, we measure the effect of a privacy policy on the largest player in the display market. Many previous studies focus on minor players. This article, hence, complements existing studying providing a perspective on how firms with different market powers may be affected by privacy policies. Finally, our data allow us to see how the effect on ad effectiveness translates into price adjustments; thus, giving some insights on the market value of targeting efficiency.

The paper is organized as follows. Section 2 provides some key facts on the digital advertising sector and on Apple's App Tracking Transparency policy. Section 3 presents the data. Section 4 describes our empirical strategy used to measure the impact on targeting efficiency and the price of a change in data quality. Section

5 provides empirical evidence of the impact of the new privacy policy. Section 6 explores the heterogeneity effect, depending on the target audience. Section 7 concludes.

2 Context

2.1 The Market for Digital Advertising

2.1.1 Three main segments in the market

According to the UK Competition and Markets Authority (CMA) 2020 report on online platforms and digital advertising, the market for digital advertising consists of three main segments: *search advertising* (the advertiser pays for prominence in user searches, which typically applies to search engines), *classified* (the advertiser pays for a listing on a comparison or e-commerce website, for instance) and *display advertising* (the advertiser pays for its ad to be displayed on a content website, for instance). In this paper, we are interested in the last segment. *Display advertising* can take two forms. It can be an *open* channel where websites propose -their *inventories* to advertisers- i.e., the number of advertising spots they are offering, with the help of several layers of intermediaries and several layers of auctions. Google is the main player in this market and owns all the largest intermediaries operating in every layer of the open display advertising market. However, the display channel can also be *owned and operated*. In this case, everything from the inventory to the advertiser interface is managed by the same player. This applies to most social media platforms, for instance, and in this part of the segment, Facebook is a dominant player. Taking into account both display channels, in the UK, Facebook generates more than 50% display advertising revenue.[2]

2.1.2 Pertinence Analyzing Facebook Data in the Display Segment

Our paper uses data on Facebook ad prices and thus focuses on the display segment. Since Facebook is a strong player in both the open display channel through its *Audience Network* and the owned and operated display channel through *Facebook* and *Instagram* social networks, the ad price data should react to any significant event

[2]Source: CMA Online platforms and digital advertising market study, July 2020.
`https://www.gov.uk/cma-cases/online-platforms-and-digital-advertising-market-study`

affecting the display channel and should be representative of a large part of the display market. This suggests that our findings should have some external validity for the rest of the market.

Additionally, studying a major actor in the fully integrated display channel has a significant advantage. Since all the steps involved in ad placement optimization are managed by the same firm, we can abstract from the complex competition process and multiple auction layers that occur in the open display channel. This allows us to consider Facebook as a firm offering its inventory to advertisers at an (average) price over which it has full control.

2.2 Apple's App Tracking Transparency

2.2.1 Giving Users a Choice over Tracking

App Tracking Transparency feature was included in the iOS 14.5 version in April 26th 2021. [3] The policy requires app developers to ask users for their consent to access Apple's *Identifier For Advertisers* (IDFA) which is a unique code assigned to mobile devices. This identifier is extremely valuable to data brokers and advertisers, as it allows them to track users across apps. Tracking has several purposes. It allows a firm to aggregate data on a specific user. These aggregated data can be used to offer targeted ads, but also to attribute conversion rate to a given ad to evaluate the effectiveness of advertising.[4]

If consumers updated their OS, they are now able to opt in or out of tracking based on a simple yes-or-no form (see figure 3 in the Appendix). Although this policy may raise the value provided to consumers by enabling increased privacy, at the same time it could severely destroy value. It may indeed deprive digital firms especially small-and-medium sized companies from vital cash inflows and from consumers acquired though targeted advertising. The consequence on the industry should of course depend partly on adoption rates of the new iOS versions and on the opt-in rate for tracking. There are unfortunately no official sources disclosing these two rates. Statistics provided by practitioners however indicate an opt-in rate

[3] In September 2020, Apple announced that one of the features of its new iOS 14 was that it would provide consumers with more transparency about privacy. In the face of resistance from several digital advertising actors, especially Facebook, the introduction of this new feature was postponed to April 26th 2021.

[4] See CMA's *Online platforms and digital advertising market study*(2020), Appendix G p.G1 & G2.

of between 12% and 25%.[5]

2.2.2 Expected Effect on Facebook's Ad Networks

Despite its already prolific data generating consumer interface, Facebook (like Google) is among the companies with the most extensive networks of third-party data providers.[6] These data improve Facebook ad targeting services in two ways. They allow to build more precise consumer profiles and they contribute to the attribution of conversions. If the change in the iOS transparency policy induces many users to choose *no-tracking*, the social network's ability to efficiently target consumers should be negatively affected. The strength of the policy's effect on Facebook is not entirely straightforward, however. Several factors may mitigate it. First, the App Tracking Transparency policy will not deprive Facebook of all its external data. Data collected through *Facebook* log-ins should not be affected. Second, as described in Jeon (2021), in open display advertising large DSPs (Demand Side Platforms) and SSPs (Supply Side Platforms) engage in cookie syncing (or cookie matching) to link third-party data coming from different sources. Facebook could engage in similar practices to compensate for the impossibility of using the IDFA.

3 Data

3.1 Ad Campaign Delivery Estimate

We use Facebook *Ad Campaign Delivery Estimate* data, collected using Facebook's Marketing API (Application Programming Interface). This API is a tool developed by Facebook to help advertisers plan massive ad campaigns. It works through requests and allows advertisers to obtain an estimate of ad performance, depending on a specified target audience. A request gives estimates on the number of *impressions*, the number of *people reached* by an ad, the number of *actions* (e.g., clicks, app installs) taken by people exposed to the ad. More precisely, each request returns three times 15 points forming three full supply functions – called the Ad Set Delivery Estimate –

[5] `https://9to5mac.com/2021/07/23/app-tracking-transparency-opt-in-snap/` last retrieved September 23, 2021

[6] See 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`

which includes estimated *impressions*, *reach* and *actions* for any possible daily budget (See a sample request output on Figure 4 in Appendix).[7] Requests also give the size of the target audience on the day or month of the request (*DAU: Daily Average Users* and *MAU: Monthly Average Users*). These estimates depend on the specified audience's demographics on which platform the audience is reached (e.g. Facebook, Instagram, Audience Network) and on what the ad placement should maximize (e.g. views, clicks, app installs). The estimates provided by the API are generated with data from Facebook Ad Manager (Facebook's Demand Side Platform) and data from ads of thousands of outside publishers (apps and websites).

3.2 Main Outcomes of Interest

With the data provided by Facebook Delivery Estimates, we build two indicators that measure *targeting efficiency* and *ad prices*. To be able to compare delivery estimate outcomes across audiences, we arbitrarily fix our daily budget to 100 € and compute *impressions* and *actions* for each of the audiences according to the returned supply curves. We then compute the CR (Conversion Rate) and CPM (Cost-per-Mille) according to:

$$CR_{100€} = \frac{Actions_{100€} \times 100}{Impressions_{100€}}$$

and

$$CPM_{100€} = \frac{100€ \times 1,000}{Impressions_{100€}}$$

The conversion rate denotes the proportion of *impressions* that trigger an *action* by consumers. We multiply it by 100 such that it could be expressed in percentage points. This variable measures targeting efficiency in our analysis. Indeed, we expect that a better targeted ad will trigger more actions per impression than a poorly targeted ad. The cost-per-mille 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.

We will use the CR and CPM as our main outcome variables to respectively measure the effect of Apple's new privacy policy on targeting efficiency and on the price of ads. Given the claims that privacy protection prevents firms from offering relevant

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

advertising, we expect the App Tracking Transparency scheme to have a negative impact on conversion rates. The effect on prices, however, may be subject to different market effects.

3.3 Target Audiences

Facebook Ad Manager provides a wide range of audiences that can be more or less precise to allow advertisers to identify the most appropriate audience. More specifically, the advertiser can choose to more or less criteria to define his target audience. These include several sociodemographic characteristics such as age, gender, education, income, and user location. Facebook Ad Manager provides also individual interests from a list of several hundred predetermined interests, such as basketball, food, cats, health or religion. Additionally, the advertiser can select users according to their behavior on the social network. For example, advertisers can target people who are friends of soccer fans, who have a friend whose birthday is in January or who are Facebook page admins. Finally, the advertiser can choose to target users depending on technical criteria such as the device they use to access Facebook, their web browser or their operating system.

3.4 Data Collection and Descriptive Statistics

We collected daily data on estimated ad prices for different demographic groups on Facebook from mid-March 2021 to mid-July 2021. We collect daily requests for delivery estimates for consumers located in the U.S.

We made requests using a combination of the audience criteria in Table 1, using one criterion per column:

Table 1: Targeting Criteria in Collected Audiences

Goal	Placement	Age	Gender	Education	Income	Interest/Behavior	OS
App Install	Facebook (Fcbk)	18-34	Male	Bachelor degree	Top 10%	Games	Fcbk access (mobile): iOS devices
Link click	Fcbk + Aud. Network	35-49	Female	N/A	Top 10 to 25%	Fcbk page admin	Fcbk access (mobile): Android devices
		50+	N/A		Top 25 to 50%	Religion	user OS: iOS 8 to 13.7
					N/A	N/A	user OS: iOS 14 to 14.4
							user OS: iOS 14.5 and above
							user OS: Android

We collect delivery estimates for ads with two possible goals: link clicks and app install. We also collect data for ads published on Facebook only and both on

Facebook and on the Audience Network (Facebook's open display channel advertising service). We gather data on ads targeting individuals with respect to several demographic characteristics: age, gender, education and income.

We obtained three supply functions at a given budget (one for *impressions*, one for *reach*, one for *actions*) per day and per state for each of these combinations. The age, gender, education and income criteria allow us to isolate the potential effects of demographic characteristics of the target population and to obtain a wide variety of audiences. Both *Goal* and *Placement* criteria are mandatory for any request. This data collection strategy allows us to measure the variation in ad performance measures across different audience types.

There are several ways to obtain a delivery estimate for ads targeting different OS. We can select either the general user OS or the OS through which the user accesses Facebook. The latter is a subsample of the former since not all accessible display consumers have Facebook accounts.[8] We collected the ad prices associated with different iOS versions, such as users of iOS 14.4 and above and users of iOS 14.5. Having both measures allows some robustness checks. Finally, we identified two Interest/Behavior categories: Games and Facebook Page Admins. Identifying consumers interested in games should depend on their activity within the Facebook social network but also on their activity off-Facebook. This information could come from other apps (the web browser or gaming apps) installed on the consumer's device. Identifying consumers who are Facebook Page Admins is possible only for Facebook, thus the information is exclusive to the ad network.

To avoid biases in our estimates, we remove missing observations such that we obtain exactly the same set of audience-date couples in our treatment group (iOS users) and control group (Android users). To ensure also that no audience is overrepresented in the before-period compared to the after period (and vice versa), we remove outliers, i.e., audiences that are observed significantly less than average in the before period or the after period.

In Table 2, we summarize the key variables presented in our sample. We have 211,342 observations. We compute the variables that measure ad performances with an ad set budget fixed at 100 euros. The average prices per action in our sample are rather high compared to the average for all industries on Facebook. [9] Figure 5 and

[8]Ads can be displayed on the Audience Network, thus, they can target only app users.

[9]`https://instapage.com/blog/facebook-advertising-benchmarks`, last retrieved 26 October 2021. Our conversion rate corresponds to a CTR on this page.

6 (in the Appendix) respectively show the average CR and CPM over time for the subsample of audiences using app install and link click optimization goals.

Table 2: Summary Statistics of the Overall Sample

Variable	Mean	Std. Dev.	Min	Max
iOS	0.5	0.5	0	1
After	0.586	0.493	0	1
App Install	0.501	0.5	0	1
log (DAU)	13.679	1.069	10.694	16.89
CR_{100}	0.584	0.349	0.055	2.815
$Impressions_{100}$	22486.63	9321.742	5202.683	76319.76
$NbAction_{100}$	144.966	131.205	11.467	848.646
CPM_{100}	5.223	2.121	1.31	19.221

4 Empirical Strategy

4.1 A Difference-in-Differences Design

Our empirical strategy aims first to estimate the effect of a change in the ability of an ad network to aggregate consumer data on ad performances (or targeting efficiency) and ad prices. To do this, we use the natural experiment provided by Apple's *App Tracking Transparency* policy. This policy should result in a negative exogenous shock to Facebook's ability to efficiently target consumers on iOS 14.5 and its later versions. To obtain the causal effect of the *App Tracking Transparency* 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.

In the first analysis, our *treatment* group consists of all iOS users, and our *control* group consists of Android users. Since a proportion of iOS users adopted the latest versions of the iOS and are hence exposed to the new privacy policy, this should allow us to identify a lower bound for the treatment effect. Our data then allow us to perform a second-grained finer analysis, comparing outcomes *post*-change among iOS versions and hence comparing outcomes for iOS 14.5 versus Android.

4.2 Defining the Before and After Periods

The introduction of the *App Tracking Transparency* scheme took place at the same time as the introduction of Apple's iOS 14.5 on April 26^{th}, 2021. However, Apple did not push the adoption of its new OS right away, making this adoption slower than usual. To ensure that a sufficently large share of users are affected by the *App Tracking Transparency* scheme in our *after* period, we follow the number of daily average users (DAU) on iOS 14.5+ measured by Facebook. To do this, we take the mean DAU of all audiences we associated with *user OS: iOS 14.5 and above*. Figure 1 shows the evolution of the mean DAU of audiences associated with iOS 14.5 and above. The y-axis indicates the mean DAU in millions in the US market. This graph shows the adoption measured by Facebook picks up around the 12^{th} of May 2021. As a consequence, we classify in the before-period all the data before April 26^{th} and in the after-period all the data after May 12^{th}. We do not include the data between the two dates in the main estimates, but we include them in the robustness checks. The graph also shows that adoption accelerates in June. This is consistent with the fact that Apple only started pushing for OS update at the release of the 14.6 version on May 26^{th} which pushed the adoption of 14.5+ versions quickly and drastically increased.

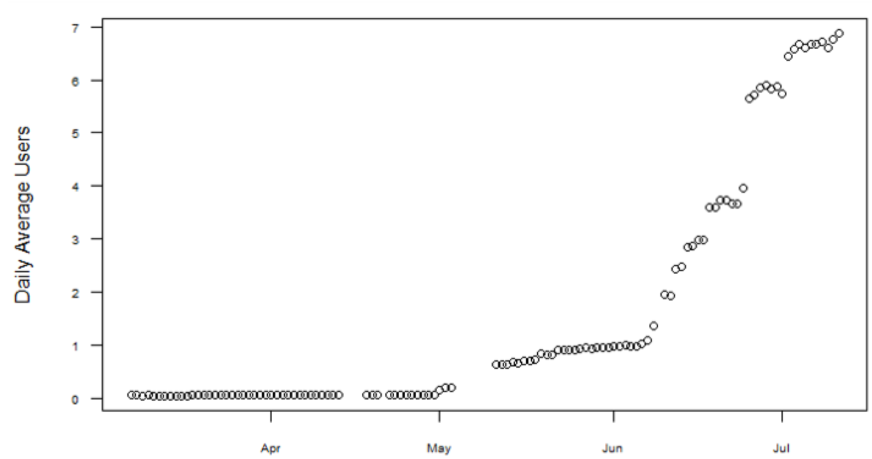

Figure 1: Average DAU (in millions) of Audiences Targeting iOS 14.5+ Users in the US

4.3 Main Empirical Strategy

The baseline equation estimates the change in our two main outcomes of interest: targeting efficiency and ad prices. We exploit variations in outcomes before and after the policy change. 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 observation date:

$$\mathbf{y_{oat}} = \alpha_{oat} + \mathbb{1}_{\text{After}} + \mathbb{1}_{\text{iOS}} + \delta(\mathbb{1}_{\text{iOS}} \times \mathbb{1}_{\text{After}}) + \epsilon_{oat} \tag{1}$$

$\mathbb{1}_{\text{iOS}}$ denotes iOS users who are likely to be exposed to the policy change as opposed to Android users and $\mathbb{1}_{\text{After}}$ measures the time period after the introduction of the App Tracking Transparency policy in version 14.5 of the iOS. 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 fixed effect, a day fixed effect, and a request hour fixed effect[10], and the log of the average daily audience size, denoted $log(DAU)$.

5 Main Estimates

5.1 Effect of App Tracking Transparency on Targeting Efficiency and Prices

The results of the main estimation with pooled iOS users are presented in Table 3.

Columns (1) and (2) report the estimates of the targeting efficiency proxied by the conversation rate *CR*. The estimates indicate that the impact of the privacy policy change on targeting efficiency is negative. In other words, preventing an ad platform from aggregating its own data on consumers with third party data seems to have a negative impact on its ability to propose relevant ads to consumers. More precisely, comparing the shift in conversion rate to the average conversion rate in the before-period on iOS indicates a decrease of 8% in Facebook's ability to transform impressions into actions. However, as we observe a pooled sample of all iOS users, this effect is likely to be a lower bound for the true treatment effect. Namely, the

[10] For technical reasons, all requests may not be made at the same time of the day.

Table 3: Difference-in-Differences Estimates for the Effect of the *App Tracking Transparency* on Targeting Efficiency and Prices

	Dependent variable:			
	CR (in %)		CPM (in €)	
	(1)	(2)	(3)	(4)
DiD	-0.051^{***}	-0.051^{***}	-0.543^{***}	-0.542^{***}
	(−0.059, −0.043)	(−0.059, −0.043)	(−0.578, −0.508)	(−0.577, −0.507)
Audience FE	No	Yes	No	Yes
Date FE	No	Yes	No	Yes
Hour FE	No	Yes	No	Yes
log (DAU)	No	Yes	No	Yes
Mean iOS before	0.665	0.665	5.106	5.106
Observations	211,342	211,342	211,342	211,342
Adjusted R^2	0.026	0.887	0.014	0.943
RSE	0.345	0.118	2.106	0.506

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

The robust standard errors are clustered by audience.

Treatment group: Facebook Access iOS device.

Control group: Facebook Access Android device.

decrease in the ability of Facebook to trigger conversions with iOS 14.5+ users should be even more important. In Column (2) of Table 3, we add covariates and fixed effects that greatly improve the explanatory power of the model but do not substantially change our estimated treatment effect, indicating some robustness in our treatment effect estimate.

Columns (3) and (4) explore the effect of a change in the ability of an ad network to aggregate consumer data on relative ad prices, measured by CPM. The results suggest that the introduction of the *App Tracking Transparency* scheme had a negative effect on the relative price of ads targeted at iOS users, compared to that targeted at Android users. Comparing the effect to the average CPM on iOS before the change shows that the magnitude of the decrease in price due to the new privacy policy is of 10%. Our results hence suggest that the decrease in targeting efficiency on one of the OSs resulted in a decrease in the relative price of ads targeted at users of the concerned OS.

Table 8 in the Appendix allows us to decompose the effect on price and targeting efficiency. This table reports the effect of Apple's policy on impressions and actions offered for a budget of 100 euros. It seems that the policy had a positive effect both on the number of impressions and on the number of actions. This suggests that Facebook's response was to increase impression supply, triggering hence more actions. The action-per-impression ratio however, did not turn out to be as good as before the *App Tracking Transparency*.

5.2 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 a time trend. More formally, we perform:

$$y_{iOS,at} - y_{Android,at} = \beta t + \epsilon_{at} \tag{2}$$

In the estimation, we add an audience fixed-effect as well as a control for the log of the daily audience size $log(DAU)$. Graph 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. There is on average a sharp difference between the outcomes before and after. More specifically, we see that the targeting ability of Facebook on iOS devices is less than that on Android devices. This translates into a difference in prices that also increases over time. There are at least two possible effects at play to rationalize such dynamics. First, as shown in Graph 1, the adoption of versions 14.5 and above increases over the time-period. As a result, the proportion of effectively treated people in the pool of iOS users increases over time. Hence, our time-trend captures the increased intensity of treatment in the intended-to-treat group. As expected, this translates into a stronger effect as more people adopt the privacy-friendly OS. Second, the negative trend in price difference could also capture the expectations of advertisers regarding the negative effect of the App Tracking Transparency on targeting efficiency. Advertisers willingness-to-pay could decrease over time as a consequence of experimenting with reduced ad performance on iOS devices after the introduction of the App Tracking Transparency policy.

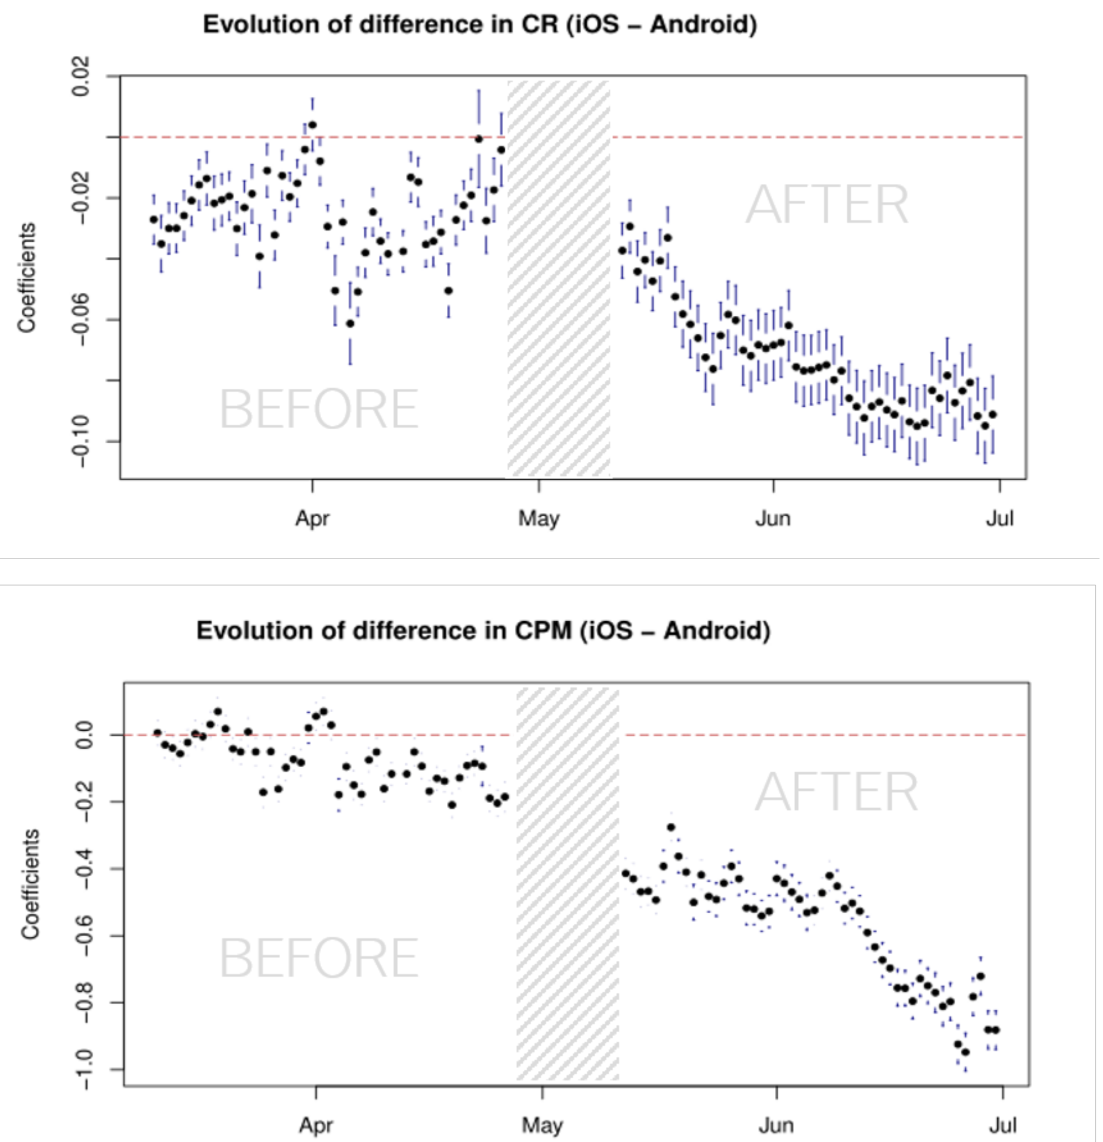

Figure 2: Evolution of the difference (iOS -Android) in CR and CPM. Estimates include audience fixed-effects and $log(DAU)$

5.3 The Effect of the Privacy Policy Change is confirmed in the Subsample of Audiences Associated with Precise iOS Versions

To test the robustness of our result, we estimate the main equation with a smaller database that includes less diverse audiences but more precise iOS versions. The data allow us to differentiate between audiences associated with iOS versions between 14 and 14.4 versus audiences associated with iOS version 14.5 and above.

Again, we consider audiences associated with Android as our control group. However, we now define our treatment group as the audiences associated with versions 14 to 14.4 of the iOS before the introduction of the *App Tracking Transparency* scheme and to versions 14.5 and above after. The results of the difference-in-differences method on our two outcomes of interest are displayed in Table 4.

Table 4: Difference-in-differences Estimates Considering as Treatment iOS 14 to 14.4 before the Introduction of the *App Tracking Transparency* and iOS 14.5 and above After

	Dependent variable:			
	CR (in %)		CPM (in €)	
	(1)	(2)	(3)	(4)
DiD	−0.221***	−0.076***	−0.192**	−0.923***
	(−0.244, −0.197)	(−0.104, −0.047)	(−0.378, −0.005)	(−1.089, −0.756)
Audience FE	No	Yes	No	Yes
Date FE	No	Yes	No	Yes
Hour FE	No	Yes	No	Yes
log (DAU) Control	No	Yes	No	Yes
Mean iOS before	0.695	0.695	2.984	2.984
Observations	13,858	13,858	13,858	13,858
Adjusted R^2	0.019	0.941	0.091	0.820
RSE	0.410	0.101	1.140	0.507

Note: *p<0.1; **p<0.05; ***p<0.01
The robust standard errors are clustered at the audience level.
Treatment group: user OS iOS 14 to 14.4 before, iOS 14.5 and above after.
Control group: user OS Android.

This additional analysis confirms the results found in the previous sections.[11] The introduction of the *App Tracking Transparency* scheme degraded the ability of Facebook to distribute targeted ads. This results in an average decrease in the prices of ads placed on iOS devices. As the set of audiences considered is not the same, we cannot directly compare the effects in Table 3 to those in Table 4. However, we can compare the magnitudes of the treatment effects. In Table 4, the effect on conversion rates represents a decrease of 11% in targeting efficiency, which is as expected somewhat higher than our baseline estimate. The effect on prices however, represents a drop in price of 30% , which is much higher than our baseline estimate.

[11]Note, however, that as we use a smaller sample of audiences, the result is much more sensitive to the addition of covariates, and our confidence intervals are much larger.

6 Intensity of the Effect on Harder to Reach Audiences

In this section, we explore how the change in the ability to aggregate consumer data may have affected different audiences differently. More precisely, we focus on three audience characteristics that could a priori have a strong influence on the intensity of the effect of a change in privacy policy. First, we look at the heterogeneity of the effect with respect to how refined and precise the target audience is. Second, we wonder whether the change in privacy policy affects ads with goals easier or harder to achieve differently. Finally, we investigate whether the policy affects ad effectiveness differently with large easy-to-reach audiences than with small narrow audiences. Indeed, one could expect that the inability of an ad network to aggregate data across sources would have a stronger impact on targeting efficiency for very precise, very narrow audiences and for ad campaigns aiming at triggering a more demanding action from consumers namely install an app versus click on a link.

6.1 Are Narrow Audiences more Affected by Apple's new Privacy Policy?

In this section, we more precisely explore the effect of the change in privacy policy for more or less refined audiences, i.e., for audiences constructed with more or less targeting criteria. As the *App Tracking Transparency* hinders the ability of ad networks to aggregate data from different sources, we expect it to be more difficult to efficiently identify and target very precise audiences. For example, it could become harder to identify educated males who are interested in religion, as this may require complementary external data. In contrast, targeting males only may not be much more difficult after the change in privacy policy. For this analysis, we construct a factor variable indicating the number of targeting characteristics (i.e. Demographic characteristics as well as interests or behaviors) used to create each audience. The variable ranges from *1-criterion*, which indicates that the audience is constructed with only one criterion, to *5-criteria*, which indicates that the audience includes five criteria. We interact each variable with the treatment to measure whether ad efficiency is likely to decrease with the increase in the number of criteria. Table 5 displays the main estimates. As expected, the conversion rate when targeting more refined audiences is more affected by the change in the ad network's ability to aggregate users' data. Once again, this translates into a gradual reduction in ad prices, which are stronger for more precise target audiences. One explanation for this effect on price could be that advertisers anticipate the greater shift in targeting efficiency for very precise audiences and lower their willingness-to-pay for these types of narrow audiences. Our results have important implications for platform management, narrow audiences are potentially targeted by small businesses that are interested in targeting a specific group of users in a small market.

Table 5: Targeting Precision and Treatment effect on CR and CPM

	Dependent variable:	
	CR	CPM
	(1)	(2)
DiD	0.189***	−0.363***
	(0.099, 0.278)	(−0.527, −0.200)
DiD*2-criteria	−0.150***	−0.113
	(−0.243, −0.058)	(−0.285, 0.060)
DiD*3-criteria	−0.233***	−0.166**
	(−0.323, −0.142)	(−0.332, −0.001)
DiD*4-criteria	−0.253***	−0.169**
	(−0.344, −0.163)	(−0.333, −0.004)
DiD*5-criteria	−0.274***	−0.259***
	(−0.364, −0.184)	(−0.431, −0.086)
Audience FE	Yes	Yes
Date FE	Yes	Yes
Hour FE	Yes	Yes
log (DAU)	Yes	Yes
Mean iOS before	0.665	5.106
Observations	211,342	211,342
R^2	0.890	0.941
Adjusted R^2	0.889	0.941
RSE	0.116	0.516
Note:		*p<0.1; **p<0.05; ***p<0.01
		The robust standard errors are clustered by audiences.

6.2 Are Ads Differently Affected by Apple's new Privacy Policy, Depending on their Objective?

We collected in our data prices for ads with different *optimization goals*, i.e., different objectives that Facebook should try to achieve when placing impressions. For some ads, Facebook is asked to maximize *link clicks*, while for others, it should maximize *app installs*. All audiences include one or the other optimization goal. These two goals may not be as easy to achieve for an ad network. Pushing a consumer to download an app may require much more data on his off-Facebook activity on his mobile phone than having him click on a link. As a consequence, we expect a certain heterogeneity in treatment effects across ad optimization criteria. Table 6 presents the estimates of the difference-in-difference changes where the baseline is the treatment effect for the goal *link click*. The second line displays the marginal treatment effect on ads promoting *app installs*.

First, we observe that the change in Facebook's ability to aggregate data had a negative effect on its ad effectiveness both for ads aiming at triggering link clicks and for ads aiming at triggering app installs. However, the results display significant heterogeneity between the two goals. The ad effectiveness in column (1) of ads looking for link clicks decreases by only 0.6% after the introduction of the *App Tracking Transparency* policy, more than ten times less than our main estimate of -8%, indicating almost no effect of the new policy on these types of ads. In contrast, the ad effectiveness of ads promoting *app installs* decreased sharply, almost twice as much as the main estimate, as the conversion rate decreased by 14.9% with respect to the mean conversion rate of all iOS audiences before the change. In summary, the limitation in the ability of ad networks to aggregate data across apps seems to have had a stronger impact on ads promoting an action that is more difficult trigger. This is probably because identifying consumers willing to download an app is more data-intense and relies more on having access to different sources of data.

Column (2) indicates that the effect of the *App Tracking Transparency* policy on the price of ads targeting iOS is also negative regardless of the optimization goal. However, the difference between the effect on CPM for ads aiming at triggering clicks and ads aiming at triggering app installs is much less striking than the difference between the effect on conversion rates. This could have two explanations (which are not mutually exclusive). It could be that advertisers do not anticipate well the change in ad effectiveness for both types of ads and expect the difference in performances to be much smaller. This would lead to a negative shock in demand of similar magnitude for ads aiming at link clicks or app installs, irrespective of the real difference in ad effectiveness. It could also be that Facebook strategically decided to lower the price of both types of ads, to induce a positive demand shock for both optimization goals.

Table 6: Differentiated Effect of the *App Tracking Transparency* on optimization goals

	Dependent variable:	
	CR	CPM
	(1)	(2)
DiD	−0.004	-0.509^{***}
	(−0.014, 0.006)	(−0.552, −0.467)
DiD*App Install	-0.095^{***}	-0.066^{**}
	(−0.106, −0.084)	(−0.117, −0.015)
Audience FE	Yes	Yes
Date FE	Yes	Yes
Hour FE	Yes	Yes
log (DAU)	Yes	Yes
Mean iOS before	0.665	5.106
Observations	211,342	211,342
Adjusted R^2	0.902	0.943
RSE	0.109	0.506

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
The robust standard errors are clustered by audiences.

6.3 Are Smaller Audiences more Affected by Apple's New Privacy Policy?

For each audience we observe our data and we obtain an estimated daily audience size : the *Daily Average Users*. In this section, we explore whether the introduction of the *App Tracking Transparency* policy had a differentiated impact on targeting efficiency and prices for ads targeting different audience sizes. We indeed expect that smaller audiences are harder to reach and require more data aggregation to be targeted efficiently. To this end, we interact our difference-in-difference coefficient with a factor variable classifying the audience in four quartiles according to their size. The first quartile contains the most narrow audiences while the fourth quartile contains the most wide audiences. We include date and hour-of-the-day fixed-effects as well as for audience fixed effects.

Table (7) displays our estimates for the effect on targeting efficiency and price. The results are expressed taking the first quartile as the baseline. We find a negative total effect both on targeting efficiency and price for all quartiles, which is consistent with our main estimate. We also find, as expected, that the effect of the privacy policy on ad effectiveness and price is stronger for the smallest audiences than for the largest audiences. However, the effect is not uniformly decreasing in the audience size, as would have been expected. Thus, it appears that the audience size is less important to understand the distribution of the effect than the audience precision, studied in section 6.1.

Table 7: Treatment effect on CR and CPM, depending on DAU quartile

	Dependent variable:	
	CR	CPM
	(1)	(2)
DiD	-0.061^{***}	-0.717^{***}
	$(-0.071, -0.051)$	$(-0.785, -0.648)$
DiD*quartile-2	0.003	0.258^{***}
	$(-0.007, 0.013)$	(0.183, 0.333)
DiD*quartile-3	0.003	0.196^{***}
	$(-0.009, 0.014)$	(0.124, 0.268)
DiD*quartile-4	0.032^{***}	0.191^{***}
	(0.017, 0.047)	(0.120, 0.262)
Audience FE	Yes	Yes
Date FE	Yes	Yes
Hour FE	Yes	Yes
Mean iOS before	0.665	5.106
Observations	211,342	211,342
R^2	0.899	0.943
Adjusted R^2	0.899	0.943
RSE	0.111	0.508

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
The standard errors are clustered by audiences.

6.4 Discussion

There is a limited empirical stream of literature that have 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 reduction in ad effectiveness of 65% after the introduction of the EU Privacy directive in 2002. This is six to ten times higher than the effect we find in our article. The difference in magnitude between the two can have several (potentially combined) rationales. First, part of the difference in results may be explained by the difference in empirical design. Our conversion rate is not as demanding as "purchases". As a result, least efficiently targeted consumers could keep interacting with the ad but not necessarily go on with a purchase. Additionally, our treatment groups do not include 100% of treated people. Table 3 shows that the treatment group encompasses all iOS users, including users of older versions. In Table 4, the treatment group is composed only of users with the treated version. However, some of them may have agreed to be track[12]. In this respect, the effect we measure may be somewhat underestimated.

However, as discussed earlier in the paper, the difference in magnitude could also be attributed to Facebook's dominance in the market. There are at least three reasons why Facebook could be less affected than other actors of the display advertising market. First, Facebook 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, Facebook is more able than other actors of the industry to find alternative profiling methods. Finally, our empirical design constrains us to 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, Facebook's user data is likely be of sufficient quality, to not be completely obsolete in such a short time period. And ad effectiveness may keep on decreasing as time goes by. For these three reasons, the effect we measure on Facebook may be a lower bound for the effect the App Tracking Transparency may have on other actors of the industry.

[12]Between 12 and 25% according to industry estimates (`https://9to5mac.com/2021/07/23/app-tracking-transparency-opt-in-snap/` last retrieved September 23, 2021)

7 Conclusion

This paper contributes to understanding the impact of privacy policies on the price and effectiveness of ad targeting on a social media platform. We study as a natural experiment the introduction by Apple of the *App Tracking Transparency* scheme on its newest versions of the iOS: 14.5 and later versions. The policy aims at giving consumers a choice over their being tracked by apps. We measure the impact of this policy on targeting efficiency and ad prices thanks to a novel dataset collected through Facebook's Marketing API. The tool provides us with price and targeting efficiency estimates for a large set of audiences. Our results are in line with previous literature studying the effect of privacy policy restriction on ad effectiveness.

We find that the introduction of the *App Tracking Transparency* scheme by Apple had a negative impact both on Facebook's ability to efficiently target consumers and trigger conversions and on the price of ads. This suggests that the loss in targeting efficiency is followed by a decrease in advertisers' demand or willingness-to-pay and/or by a loss for Facebook of 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. We find that the effect on ad effectiveness and ad prices strengthens with the precision of the target audience, and is stronger for ads promoting app installs compared to ads promoting link clicks. This suggests that while simple advertising with easy-to-get audiences is not much threatened, Apple's new scheme may strongly affect the ability of firms to develop more sophisticated ad campaigns with audiences that are harder to reach or actions by consumers that are harder to trigger. The effect measured for ads on Facebook is likely to be a lower bound for the industry, as the social network generates many proprietary data that are not affected by the policy.

There are limitations to this research. First, our data measures only the effect of privacy policy change on one platform which is one of the biggest ad platform worldwide. Second, our estimates measure an immediate effects of the policy changes and do not consider long-term effects. Third, we analyze ad effectiveness and ad prices but we are not able to measure the total effect on revenues for advertisers.

References

Alcobendas, Miguel, Shunto Kobayashi, and Matthew Shum, "The Impact of Privacy Measures on Online Advertising Markets," *Available at SSRN 3782889*, 2021.

Aridor, Guy, Yeon-Koo Che, and Tobias Salz, "The Economic Consequences of Data Privacy Regulation: Empirical Evidence from GDPR," Technical Report, National Bureau of Economic Research 2020.

Goldfarb, Avi and Catherine E Tucker, "Privacy regulation and online advertising," *Management science*, 2011, *57* (1), 57–71.

Jeon, Doh-Shin, "Market power and transparency in open display advertising–a case study," 2021.

Johnson, Garrett A, Scott K Shriver, and Shaoyin Du, "Consumer Privacy Choice in Online Advertising: Who Opts Out and at what Cost to Industry?," *Marketing Science*, 2020, *39* (1), 33–51.

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 third-party cookies on publisher revenue," *Google Report*, 2019.

Appendix

A Figures and Tables

A.1 App Tracking Transparency screenshots

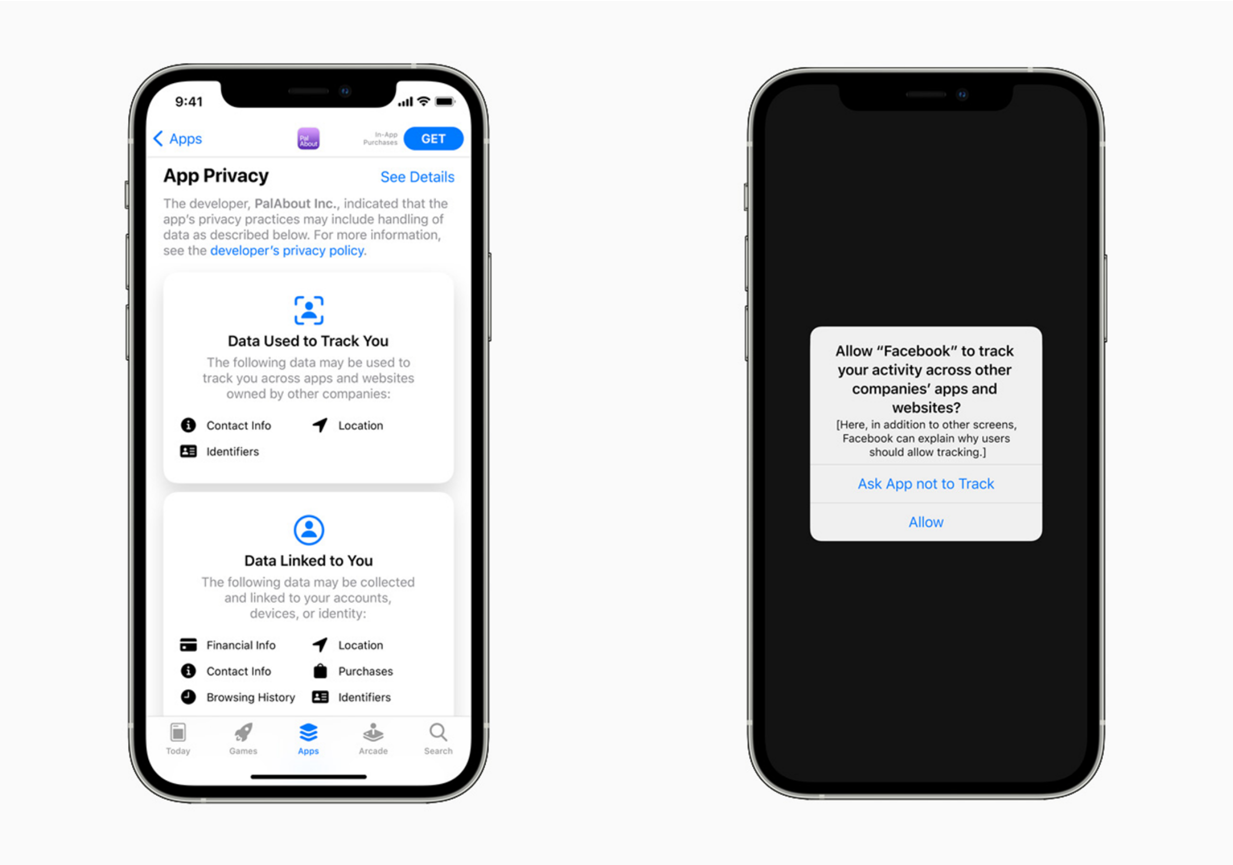

Figure 3: Apple App Tracking Transparency new features - Screenshots

A.2 Sample Delivery Estimate output

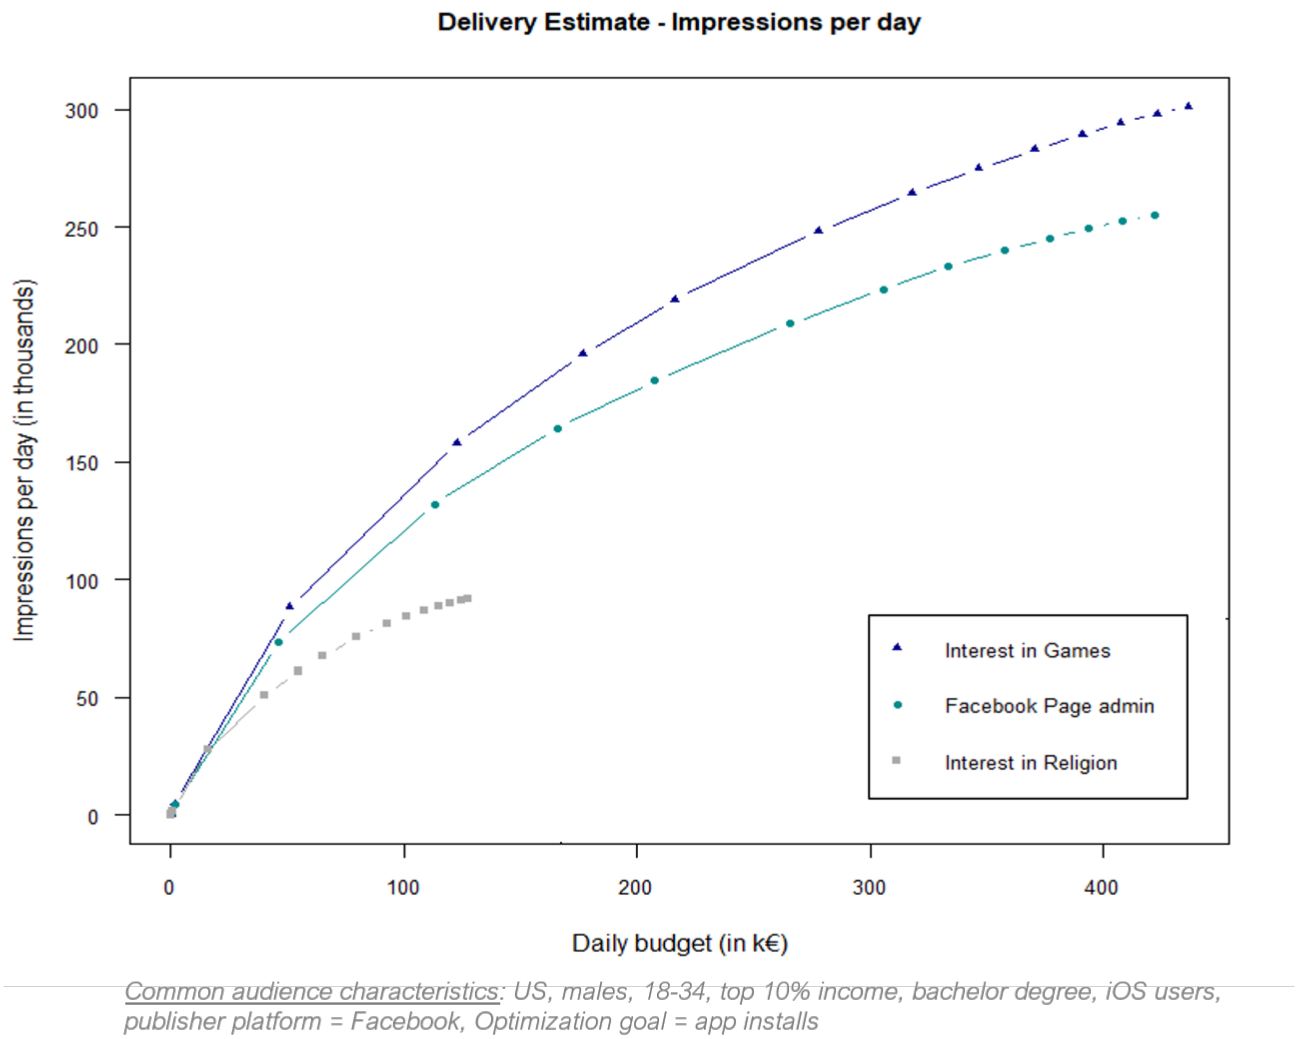

Figure 4: Example of Delivery Estimate output for impressions, with three different audiences.

A.3 Effect on Impressions and Actions

Table 8: Effect of the *App Tracking Transparency* on Impressions and Actions – Daily budget of 100 €

	Dependent variable:	
	Impressions	Actions
	(1)	(2)
DiD	2,127.577***	3.261***
	(1,998.458, 2,256.695)	(1.876, 4.647)
Audience FE	Yes	Yes
Date FE	Yes	Yes
Hour FE	Yes	Yes
DAU Control	Yes	Yes
Observations	211,342	211,342
Adjusted R^2	0.944	0.955
Residual Std. Error	2,208.420	27.744
Note:		*p<0.1; **p<0.05; ***p<0.01
		The standard errors are clustered at the audience level.

B Descriptive Statistics: Graphical Evidences

Figure 5a and 5b respectively show the average CR for audiences with optimization goal *App Installs* and *Link Clicks*. Descriptive evidences suggest that the conversion rate is likely to decrease over time after the introduction of the *App Tracking Transparency* option for ads targeting *App Installs*. Figure 5b suggests that the privacy policy has almost no effect on ad effectiveness for ads with a *Link Clicks* optimization goal.

Figure 6a and 6b respectively show the average CPM for audiences with optimization goal *App Installs* and *Link Clicks*. Overall, the descriptive statistics suggest that the CPM of ads targeting iOS users is likely to decrease after the policy change compared to ads targeting Android users.

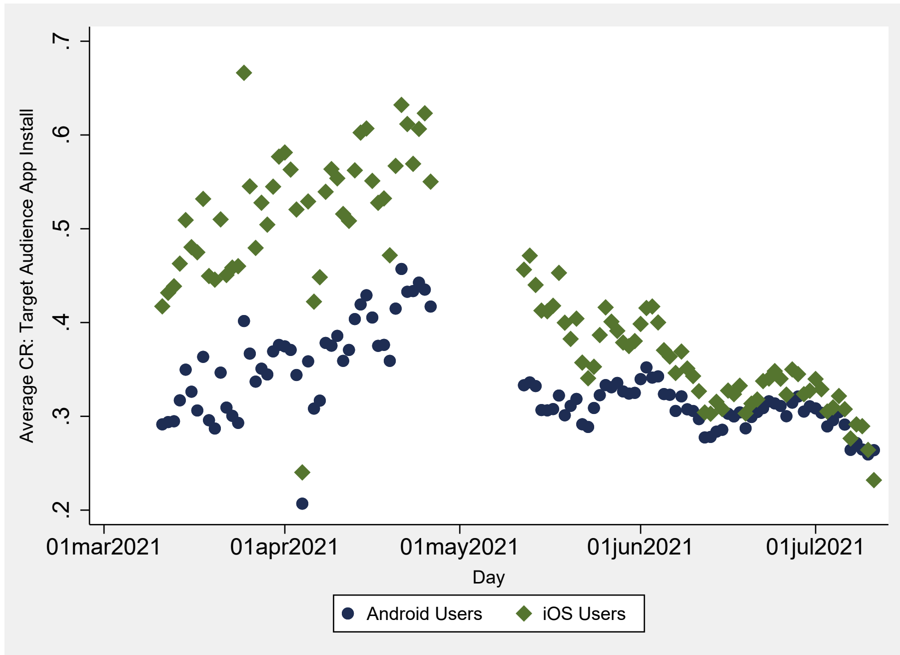

(a) App Installs

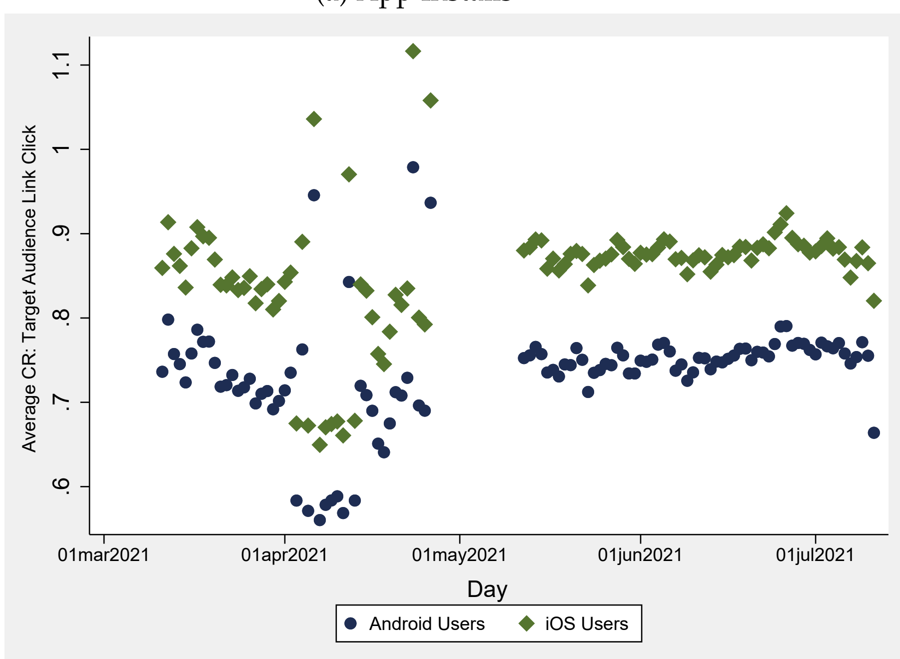

(b) Link Clicks

Figure 5: Descriptive Evidences: Average CR

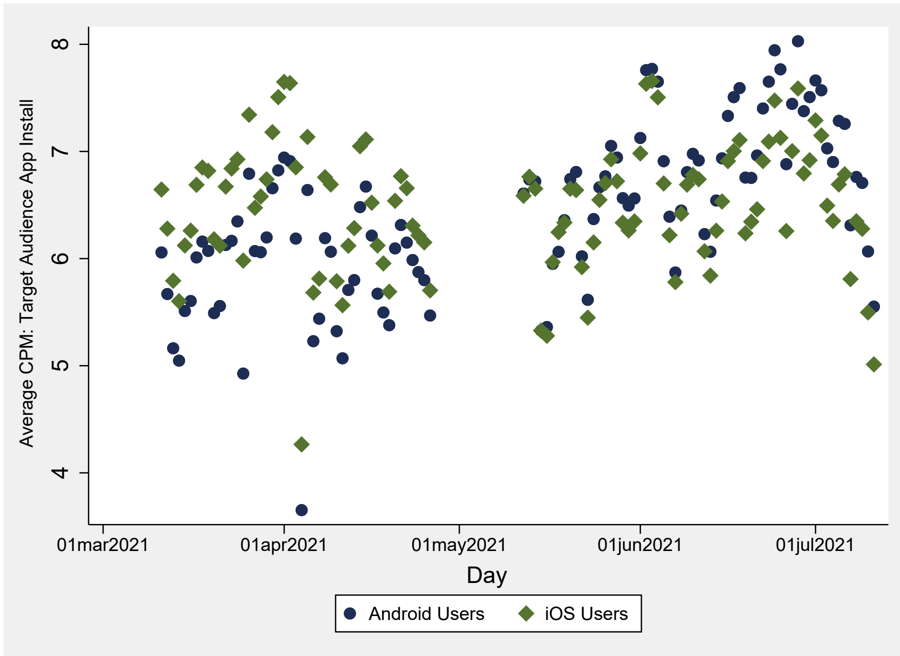

(a) App Installs

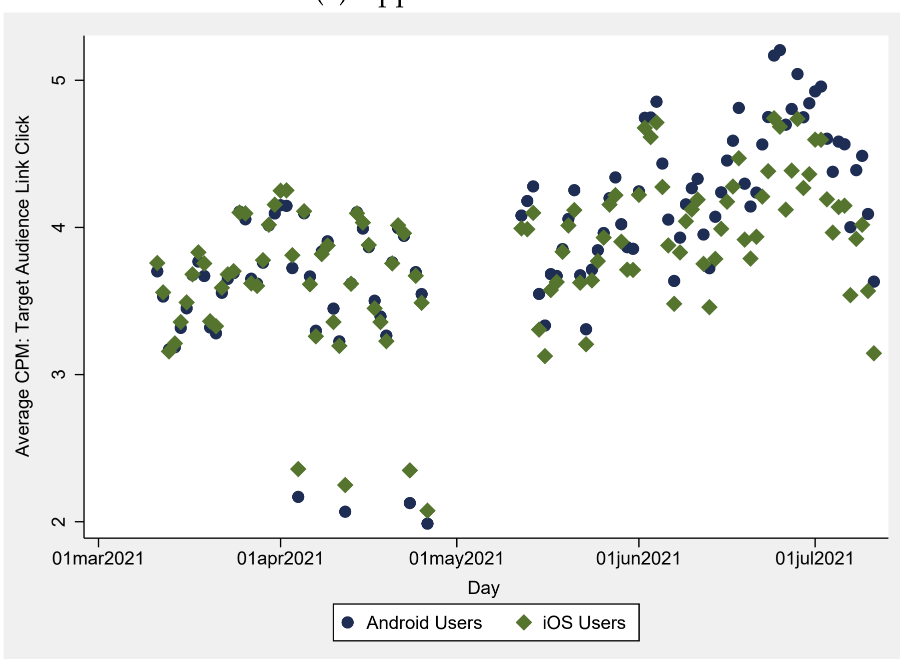

(b) Link Clicks

Figure 6: Descriptive Evidences: Average CPM

C Parallel trends

Table 9 presents the estimates when we construct a placebo treatment period considering the period between February 10th to April 15th to verify the parallel trend hypothesis. The result shows that the placebo treatment has no effect in the main outcomes.

Table 9: Effect of placebo period on Conversion Rates

	Dependent variable:	
	CR	
	(1)	(2)
DiD	0.009	0.010
	(−0.003, 0.021)	(−0.002, 0.022)
Audience FE	No	Yes
Date FE	No	Yes
Hour FE	No	Yes
log(DAU)	No	Yes
Observations	5,904	5,904
Adjusted R^2	0.019	0.812
RSE	0.376	0.165
Note:		*p<0.1; **p<0.05; ***p<0.01
		Standard errors are clustered at the level of the audience.
		Observation period: February 10th to April 15th
		Before: February 10th to March 5th. After: March 16th to April 15th