

Privacy Regulation and
Targeted Advertising:
Evidence from Apple's App
Tracking Transparency

Guy Aridor, Yeon-Koo Che

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Privacy Regulation and Targeted Advertising: Evidence from Apple's App Tracking Transparency

Abstract

We use a novel dataset of online advertiser performance and product sales to quantify the medium-term economic effects of Apple's App Tracking Transparency Policy (ATT). We find that ATT significantly degraded the ability by Facebook advertisers to target advertisements based on its off-platform data. A within-advertiser comparison reveals that conversion-optimized advertisements, for which such data is crucial for targeting, suffered a 37.1% reduction in click-through rates, compared with clicks-optimized advertisements that depend less on such data, indicating a significant fall in the relevance of the former ads as perceived by users. Although advertisers did appear to substitute away from Facebook for the Google ecosystem, those with higher baseline dependence on Facebook experienced difficulty in customer acquisition, receiving 26.2% fewer orders from new customers.

Keywords: privacy, app tracking transparency, online advertising, privacy regulation, social media.

Guy Aridor
Northwestern University
Kellogg School of Management
Evanston / IL / USA
guy.aridor@kellogg.northwestern.edu

Yeon-Koo Che
Columbia University
New York City / NY / USA
yeonkooche@gmail.com

January 19, 2024

We thank Tobias Salz for his early contributions to this project as well as Sam Goldberg and Silvio Ravaioli for helpful comments. We further thank Maurice Rahmey, Nils Wernerfelt and Michael Zoorob for answering questions about advertising on Facebook, and Chelsea Mitchell and Andrew D'Amico for excellent research assistance. All errors are our own.

1 Introduction

The online advertising market is an important part of the modern economy, accounting for more than \$200 billion in annual revenue in the US alone (IAB, 2023). Many technology companies monetize their services through advertising, and the targeting capabilities enabled by online advertising help sustain small businesses and expand the set of offered product varieties (Baslandze et al., 2023). At the same time, the heavy use of consumer data in online advertising has raised concerns about consumer privacy. In response, new regulations have been put in place, such as the EU’s General Data Protection Regulation (GDPR). Beyond government-based privacy regulation, the firms within this industry have also begun to self-regulate and directly make changes to the underlying technology that affects how consumer data is processed online. These regulations have important benefits by protecting consumers’ data, but they also may have adverse effects in the advertising market and downstream product markets.

In this paper, we empirically examine the consequences of one of the most prominent cases of self-regulation – Apple’s App Tracking Transparency (ATT) iOS 14.5 update. ATT allowed consumers on iOS devices to block their identifiers from being sent to mobile applications, making it impossible for advertisers to track opted-out consumers across apps and websites. This resulted in advertisers neither being able to use consumers’ off-platform data for targeting nor measure whether an advertisement ended in a purchase. One unique aspect of this policy is its potential impact on behaviorally targeted advertising platforms, such as Facebook, given its heavy dependence on off-platform data for targeting, and its reliance on the mobile ecosystem.¹ Thus, we exploit these distinctive aspects of the policy to quantify how it impacted the effectiveness of targeted advertisements and product sales as well as how advertisers substituted across advertising platforms and modalities in response.

To study these questions we rely on novel data from an advertising data analytics company on the advertising campaigns of thousands of online advertisers across channels such as Facebook, Google, and TikTok as well as observed real sales outcomes from a subset of their Shopify dashboards. For each of these advertising platforms, we observe how much money was spent (spend), how many views (impressions) and ad clicks (clicks) were made, and how many consumers completed an off-platform purchase event associated with the campaign (conversions). The companies that use this provider tend to be smaller, direct-to-consumer, and online-focused companies – which is precisely the companies whose existence purportedly hinges on targeted social media advertising (Werner, 2022) and for whom it is important to understand how privacy regulation impacts them.

¹For instance, <https://www.cnbc.com/2021/03/11/why-facebook-is-so-upset-about-apple-idfa-change-insiders-spill.html> explains that Facebook expected the update to harm their ecosystem

We first document the time trends from September 2020 until October 2022 for the relevant measures of Facebook’s conversion-optimized advertisements performance: conversions, cost per conversion, and click-through rate. We find that after the onset of ATT there is a stark drop in the number of conversions and an increase of nearly 50% in the cost per conversion that closely follows the adoption curve of iOS devices with the iOS 14.5 update and which persists through the end of the sample period. Although suggestive of a dramatic impact on the effectiveness of Facebook advertising, both of these quantities are subject to measurement error since Facebook can no longer reliably log conversions after ATT. Thus, we also study an on-platform measure of the relevance of advertisements to consumers: click-through rate, which falls 7.7% after ATT.

We next conduct several analyses quantifying the extent to which Facebook advertisers were impacted by ATT along several dimensions. We find that smaller advertisers, who likely have less on-platform data, are impacted more than large advertisers, and campaigns with targeting criteria more reliant on off-platform data were severely impacted. While this evidence is suggestive of degraded targeting effectiveness, it is not causal. To provide causal evidence, we conduct a within-advertiser difference-in-differences analysis that compares the relative performance of the off-platform objective (sales/conversions) vs. the on-platform objective (link clicks). Such an analysis controls for differences across advertisers and allows for equilibrium adjustments by the platform in targeting over time; importantly, it considers a control group of campaigns relatively unimpacted by the policy.² We find that there is a 37.1% reduction in click-through rates and a similar magnitude increase in the relative cost per click for the conversion-optimized campaigns.

How did ATT impact the advertisers with heavy reliance on Facebook ads? Even with the apparent degradation of Facebook ads quality, these advertisers could in principle mitigate the impact by substituting away from Facebook for other forms of advertising. To answer this question, we study the sales data from Shopify for those heavily dependent on Facebook ads before ATT, and compare them with those less beholden to Facebook ads pre-ATT. Not only is the Shopify data unaffected by measurement issues, but it also contains rich information about sales performances. We consider a difference-in-differences specification where advertisers with high pre-ATT Facebook share are the treatment group and advertisers with low pre-ATT Facebook share are the control group. We find that, even after controlling for advertising spending, there is a 26.2% and 24.7% reduction in the number of orders and total revenue, respectively, from new customers recorded in the Shopify data. We further find a relative increase in the fraction of total orders coming from repeat customers and a null effect on the average order value per customer.

Finally, we find evidence that market share shifts from the Facebook advertising ecosystem to Google and that there is larger growth in Google Display (behaviorally targeted advertising)

²While in principle, advertisers could substitute between these different types of campaigns, we find little evidence for this after ATT.

relative to Google Search (contextual advertising). We also document time trends indicating that prices on Google Search ads increased after ATT. While this provides suggestive evidence for equilibrium adjustments in the advertising market, the results indicate that Facebook-dependent advertisers were still harmed in the medium-term primarily through reduced customer acquisition.

Overall, this paper makes two primary contributions. The first is that we empirically document that Apple’s App Tracking Transparency policy had a clear negative impact on the efficacy of targeted advertising and the effect was felt more significantly by small businesses dependent on Facebook advertising. The second is that it provides suggestive evidence for how such a policy changes the composition of the online advertising ecosystem in terms of the platforms and types of advertising used by advertisers. While our work does not directly measure the welfare consequences of the policy, we view our paper as providing important evidence to help guide the regulation of consumer data in the market for online advertising.

2 Related Work

The current paper is related to the literature that studies the impact of privacy regulation, including public ones such as GDPR, on business investment/creation ([Jia et al., 2021](#); [Janssen et al., 2022](#)), online search ([Zhao et al., 2021](#)), e-commerce revenues and online advertising ([Goldfarb and Tucker, 2011c](#); [Aridor et al., 2023](#); [Goldberg et al., 2023](#)), data storage ([Demirer et al., 2023](#)), and market concentration for web trackers ([Peukert et al., 2022](#); [Johnson et al., 2023](#)). Several other papers also study the impact of Apple’s App Tracking Transparency (ATT) and more broadly changes to the privacy environment on iOS. [Bian et al. \(2021\)](#); [Li and Tsai \(2022\)](#); [Kesler \(2022\)](#); [Cheyre et al. \(2023\)](#) study the market-level impact of iOS policies – highlighting that these policies, including ATT, had negative effects on the incentives to develop new applications. To the best of our knowledge, this paper is the first that provides direct evidence for the degradation of ad efficacy (as measured both by reduced CTR and sales) and characterizes advertiser modality substitution in response to such regulation. Compared with public regulations such as GDPR, ATT provides a cleaner empirical characterization of the effects of opt-in-based privacy regulation, as it neither involves the compliance issues ([Ganglmair et al., 2023](#)) nor suffers from the heterogeneity in the design of opt-in prompts ([Utz et al., 2019](#)) that are endemic to GDPR.

Our paper is also related to the marketing literature that studies online behavioral display advertising and the role played by consumer-level data for its effectiveness ([Bleier and Eisenbeiss, 2015](#); [Johnson et al., 2017](#); [Sahni et al., 2019](#); [Rafieian and Yoganarasimhan, 2021](#); [Wernerfelt et al., 2022](#)). In contrast, our work does not directly study the effectiveness of behaviorally targeted advertising but rather uses ATT as a negative shock to its effectiveness to assess substitution patterns and measure its impact on downstream economic outcomes. Thus, it focuses on the im-

plications of reduced ability to target advertisements. We utilize insights from the experimental design of [Wernerfelt et al. \(2022\)](#), which measures the relative importance of on-platform versus off-platform data, to guide our empirical strategy. We find comparable results to their experiment, even though our analysis inherently accounts for equilibrium effects after the loss of off-platform data and consumer opt-in choices.

Finally, our paper is also related to the broader literature on advertiser competition. In the theoretical literature, the typical modeling approach is to consider a two-sided market (i.e., following [Rochet and Tirole \(2003\)](#)) with advertisers on one side and consumers on the other ([Anderson and Coate, 2005](#); [Wilbur, 2008](#)). A more recent focus has been on advertiser competition and substitution patterns when consumers multi-home in their platform choices ([Ambrus et al., 2016](#); [Athey et al., 2018](#); [Gentzkow et al., 2022](#)). In contrast to these papers, we abstract away from the consumer side and focus on advertiser demand across different *channels*. Thus, instead of focusing on competition on outlets within a given channel, we characterize substitution across channels. We are not the first to focus specifically on channel substitution; [Goldfarb and Tucker \(2011b\)](#) show that search advertising prices increase for lawyers when regulation forbids them from contacting via mail and [Goldfarb and Tucker \(2011a\)](#) provide evidence for the substitutability between online and offline advertising. Our paper complements this work by studying the substitution between different online advertising platforms and modalities.


3 Apple’s App Tracking Transparency Update

In this section, we describe the relevant background on Apple’s App Tracking Transparency iOS 14.5 update and its implications for advertisers. From the consumer perspective, the main change is that when users upgrade to iOS 14.5 and open up an application they see the prompt in the left panel of [Figure 1](#), which asks them whether this application can track them. Unlike other privacy regulations such as the GDPR, there were neither compliance issues nor heterogeneity in the design of the opt-in prompt ([Utz et al., 2019](#)). In particular, if an application wants to remain listed on the app store they are required to comply with the policy by including this prompt in their application and the prompt design was held constant across applications by Apple’s requirements. Furthermore, consumers can fully opt out across all applications in the ecosystem.

If a user chooses to opt out on a given application, then both the application and its associated advertisers can no longer observe the user’s identifier for advertisers (IDFA). This means that advertisers cannot track user actions across different applications. Industry reports highlighted that opt-out rates were as high as 66%, indicating substantial usage.³ Furthermore, the right panel of [Figure 1](#) shows that while the adoption date was on April 25th, 2021 – iOS 14.5 adoption

³<https://youappi.com/the-state-of-att-opt-in-rates-in-2023/>

Figure 1: ATT Details and Adoption



amongst users was rather gradual, but with a sudden increase in adoption rate roughly 5 weeks after the policy as Apple began to nudge users to adopt the new OS. It is also important to note that the policy *only* impacts iOS mobile applications and not web browsers on phones or computers.

There are two important implications for behaviorally targeted display advertising platforms, such as Facebook. Facebook uses data from outside of its ecosystem of applications to have additional data points to target individuals as well as to measure the effectiveness of its ads. The latter aspect indicates that, beyond the reduction in data, a crucial element of the feedback loop needed for targeting ads is broken. For instance, suppose a consumer sees an advertisement on Facebook for a Nike shoe, clicks on the ad to go to the Nike application, and subsequently purchases on the application. To target ads effectively, Facebook must be able to learn that the advertisement was a success for some individuals and not for others, but if a consumer opts out through ATT then Facebook is unable to measure whether the consumer bought on Nike’s app or not. This also means that Facebook is limited in its ability to accurately report conversions to advertisers. Indeed, following ATT, Facebook rolled out Aggregated Event Measurement to aid in this measurement issue where they replaced actual observed conversions with “modeled” conversions for users that opted out.⁴ Thus, both the loss in off-platform data and conversion measurement issues can contribute to an overall degradation in targeting and impact the data observed by advertisers.

4 Data

Our paper uses a dataset from an anonymous data provider which provides a granular view of the digital advertising platforms used by primarily e-commerce firms – those most likely to be reliant on targeted advertising. We study a period of two years – September 2020-October 2022

⁴See <https://www.facebook.com/business/help/721422165168355?id=1877298665783613> for additional details on measurement changes on Facebook post-ATT.

– where we observe information about 2,442 advertisers throughout this sample period, including their Facebook, Google, and TikTok advertising expenditures as well as their Shopify e-commerce revenues.

Our data source provides us with weekly aggregates of advertiser performance. Advertisers, whose identities are anonymized, contract with the data provider and share their relevant performance data from various advertising and e-commerce platforms. The data source’s information is continually updated and is backfilled until mid-2020. For each of the advertising platforms, we observe the total amount of dollars (spend), the number of times the advertisements were seen (impressions) and clicked on (clicks) as well as the total number of conversions associated with the advertising campaign (conversions). The measurement of the first three (spend, impressions, clicks) is not affected by ATT; they are measured accurately and consistently before and after ATT. However, the conversion is potentially affected by ATT as it is typically collected through a pixel that the advertiser embeds within its website or application that requires a consistent identifier across the platform of interest and the third-party website/app.⁵ Finally, we compute the implied price per click (CPC), price per 1,000 impressions (CPM), and price per conversion (CPP). The relevant price is different across different advertising forms – for instance, for Facebook, TikTok, and Google Display advertisers typically pay per impression (CPM), whereas for Google Search advertisers pay per click (CPC).

Within each advertising platform, we observe this data at different levels of granularity. For Facebook, we observe each of these measures for each unique pair of campaign objectives and targeting criteria. For the targeting criteria, we observe whether the ad was a retargeting campaign (aimed at customers who already viewed their app/website or are in their database), a prospecting campaign (aimed at automatically finding new relevant customers based on behavioral data), or a lookalike campaign (aimed at finding new relevant customers based on existing customer observables). For the campaign objectives, we observe a wide range of objectives but, apart from studying within-platform readjustment, we primarily focus on campaigns that are optimized for off-platform conversions or on-platform clicks as these are the vast majority of the campaigns. For Google, we observe the breakdown according to the advertising type: Google Search ads, Google Discovery/Video which show up on YouTube, and Google Display which show up on many websites around the Internet. For TikTok, we observe aggregated data across all the advertisers’ campaigns on the application.

For a subset of these advertisers, we observe weekly aggregated data from their Shopify accounts. Shopify is an online platform that provides sellers with a suite of software tools to bundle

⁵See <https://www.facebook.com/business/tools/meta-pixel> for more information on the Meta pixel and <https://ads.tiktok.com/help/article/tiktok-pixel?lang=en> for more information on the TikTok pixel.

all their commerce activities – online and offline – in one dashboard.⁶ Shopify handles the logistics involved in creating an online webstore, payment processing, and many other aspects of selling online. Importantly for our purposes, this data provides us with a complete view of sales for the advertisers. From Shopify, we observe the total weekly revenue, the number of orders, the average order value, and the fraction of orders that come from repeat customers. Unlike the data on conversions from the advertising platforms, the Shopify data has no measurement issues as a result of ATT. Notably, the measurement of repeat customers relies on data unaffected by the changes from ATT since they are typically user-provided email addresses or phone numbers. Since some of our advertisers have low order counts, we aggregate our data to a monthly level to avoid having to make assumptions about when zero means an exit or no orders in a week for this advertiser.

As there is no clear control group that we can consistently utilize across all the outcome variables, we rely on three distinct empirical strategies. We will discuss them in detail when we use them; here, we describe some commonalities across them. To ensure that we are measuring the medium-term effects of the policy, we primarily conduct our analysis over the period from September 2020 to October 2022, though some of our analyses are conducted over tighter windows around the policy due to changes in the composition of Google products. Across each of these empirical specifications, we cluster our standard errors at the advertiser level. Almost all of the outcome variables that we consider are lognormally distributed since there are some relatively large advertisers in our sample, but most of them are small online stores. As such, we consider a log transform for most of the outcome variables that we consider when there are no zeros and the inverse hyperbolic sine transform when there are zeros.

5 Empirical Results

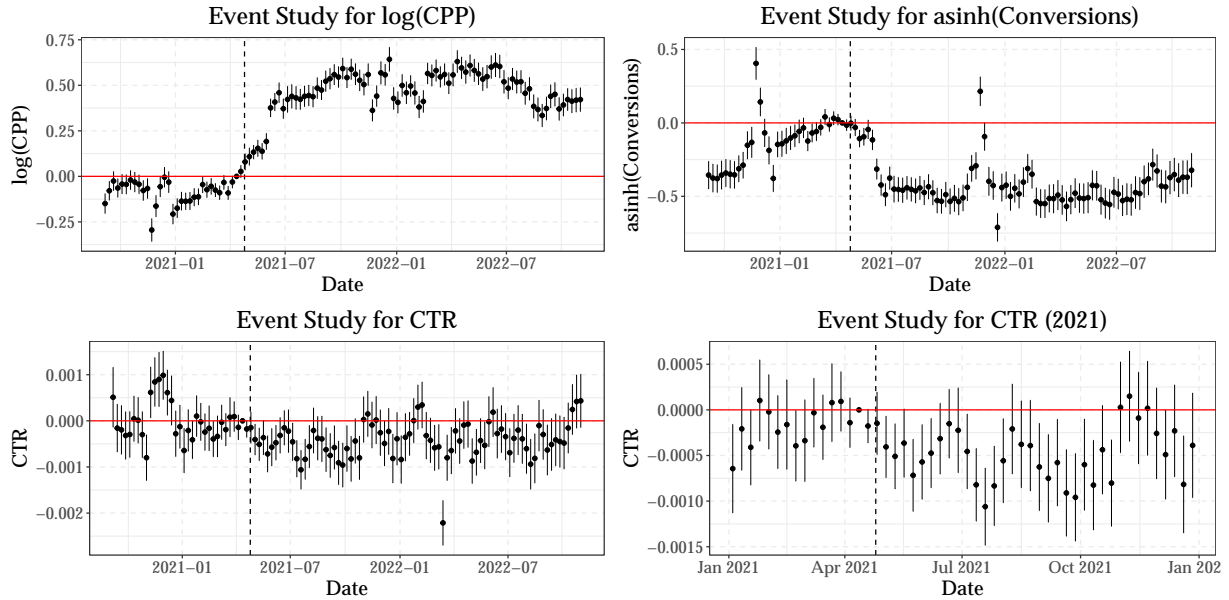
5.1 Descriptive Evidence for Reduced Effectiveness

We first examine the extent to which ATT impacted the acquisition costs, (Facebook-measured) conversions, and click-through rates. To this end, we restrict attention to the advertiser campaigns on Facebook that are optimized explicitly for off-platform conversions and ensure a balanced panel of non-zero spend across each of the considered periods.⁷ With this dataset in hand, we use the

⁶For full details about Shopify, see <https://www.shopify.com/blog/what-is-shopify>.

⁷We winsorize these outcomes at the 1 and 99 percentiles in order to remove several extreme outliers. However, our results are similar without this. Furthermore, our results are similar if we consider a balanced panel of advertisers with at least one registered conversion in every period.

Figure 2: Event Study on Conversions, Acquisition Costs, and CTR



following event study specification to study the time trends for these measures:

$$Y_{it} = \sum_{j=T_0}^T \beta_j date_j(t) + \alpha_i + \epsilon_{it}, \quad (1)$$

where β_j is the coefficient of interest, Y_{it} denotes the outcome of interest in time t for advertiser i , $date_j(t)$ is an indicator for time period t , and α_i denotes advertiser fixed effects. We will consider three variables for Y_{it} : the log of acquisition cost (CPP), the inverse hyperbolic sine transform of the number of conversions, and the click-through rate (CTR).

The results of the event studies for each of these outcome variables are reported in Figure 2. Beyond the spikes around Black Friday and the holiday season, a clear pattern emerges. The rise (resp. fall) of CPP (resp. conversions) coincides with the adoption trends documented in Figure 1. In particular, at the onset of ATT there begins to be an increase in CPP with a nearly discontinuous increase of nearly 25% approximately 6 weeks after ATT takes effect, consistent with the spike in adoption according to Figure 1. However, both of these outcome variables are subject to possible measurement issues as a result of ATT, so it's entirely possible that the drop in conversions and increase in acquisition cost may not reflect the real effect on the efficacy of advertising campaigns.

In contrast, click-through rates are immune from the measurement issue. More importantly, click-through rates directly reflect the relevance of the advertisements shown to consumers, so they are a reasonable proxy for the quality of ad targeting. As the displayed ad becomes less well-targeted, it will be perceived as less relevant by the consumer, so she will be less likely to click the

ad. The changes in the click-through rates are reported in the bottom row of [Figure 2](#) for the full time period (left) as well as zooming into 2021 (right). The results suggest that the effectiveness of ads, mostly likely reflecting the quality of targeting, suffered as a result of ATT, although the magnitude of the fall is not as dramatic as the changes in acquisition costs and conversions.

5.2 Impact on Targeting Effectiveness

In this section, we further explore the evidence for a degradation in the ability to target advertisements following the onset of ATT. To this end, we study how the changes in acquisition costs and click-through rates differ according to the targeting criteria and the size of the advertisers. Finally, we provide causal evidence for the degraded effectiveness of ad targeting following ATT.

We consider two heterogeneity analyses with the following specification:

$$Y_{it} = \beta_j \left(Post_t \times HT_i \right) + \alpha_i + \epsilon_{it}, \quad (2)$$

where $Post_t$ is an indicator function for whether t is past the ATT policy and HT_i is the heterogeneity that we consider. The outcome variables are acquisition costs and click-through rates. The results are reported in [Table 4](#) where columns (1) and (4) show a statistically significant baseline increase of 0.541 for the log(Acquisition Cost) and a decrease of 0.001 for the click-through rate.

We first consider the heterogeneous effects between small and large advertisers where we define small (resp. large) advertisers as those with below (resp. above) mean Facebook spend in the pre-ATT period. Since larger advertisers are likely to have more on-platform and historical data to use for targeting, we would expect that smaller advertisers would be more impacted relative to large advertisers. Column (2) of [Table 4](#) shows that larger advertisers experience a 16.2 percentage points smaller increase in acquisition costs relative to smaller advertisers and column (4) shows that they also have a smaller reduction in click-through rate.

Second, we study heterogeneity across targeting criteria. We consider the change in acquisition costs relative to lookalike campaigns, those least likely to be impacted by ATT since these campaigns rely on finding new customers based on matching existing customer observables. We expect that prospecting and retargeting campaigns will be more impacted relative to lookalike campaigns since they are more dependent on off-platform data relative to lookalike campaigns. Columns (3) and (6) of [Table 4](#) report this heterogeneity analysis, which finds mixed results as retargeting campaigns have a reduced click-through rate and similar acquisition costs, whereas prospecting campaigns have increased acquisition costs and similar click-through rate.

The previous results cannot be interpreted causally as they don't have a proper control group. Therefore, we explore whether we can make a causal claim about the degraded ability to target advertisements. To this end, we exploit the fact that ATT impacted exclusively the advertising

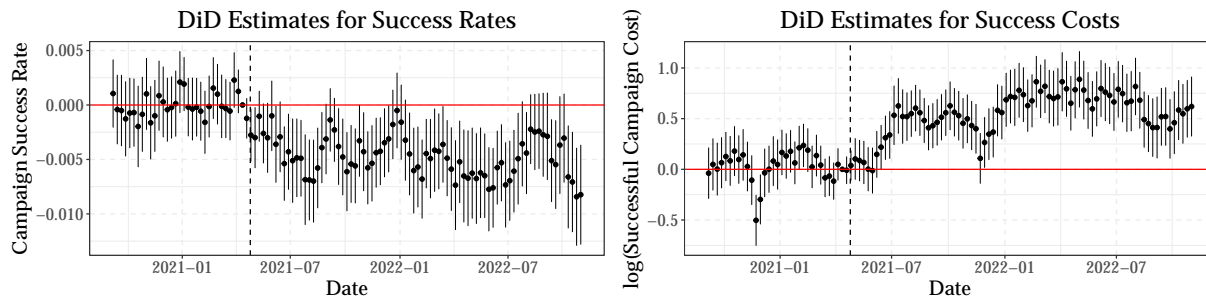
campaigns using off-platform data. Off-platform data is used to directly inform the targeting criteria and it is also used to “close the feedback loop” to accurately measure whether an advertisement concluded with a conversion or not. While the former impacts any advertising campaign run after ATT, the latter only impacts the off-platform sales-optimized campaigns. Thus, we rely on the insight from [Wernerfelt et al. \(2022\)](#) to consider the relative performance of off-platform (sales) optimized campaigns to on-platform (link clicks) optimized campaigns. This has two benefits: maximizing link clicks is the most common on-platform campaign objective and the outcome is measured on the platform without any measurement issues.

We consider a *within-advertiser* difference-in-differences analysis where we compare the performance of off-platform (sales)-optimized campaigns to on-platform (clicks)-optimized campaigns for the same advertiser. This allows us to isolate the role that measurement plays and its relative impact on the ability to target advertising. By comparing within-advertiser performance, it also controls for possible differences across advertisers – for instance, their size or frequency of purchases – that are orthogonal to the treatment effect of interest as well as possible adjustments to the targeting algorithm by Facebook over time. In this case, we consider the following specification for advertiser i , advertiser campaign objective j , and time t :

$$Y_{ijt} = \beta \left(Post_t \times T_j \right) + \alpha_i + \kappa_t + \epsilon_{ijt} \quad (3)$$

where β is the coefficient of interest, T_j is whether the campaign j is in the treated group, $Post_t$ indicates whether t is the past ATT implementation date, α_i denotes advertiser fixed effects, and κ_t denotes time fixed effects. We will also consider a time-varying specification of this where we estimate a time-varying β and interact T_j with $date_t$ instead of $Post_t$. We will use this to assess the reasonable of the parallel trends assumption. We estimate this specification over a sample of advertisers who had non-zero spend across the full time period and used both types of campaigns before ATT.⁸

Figure 3: Difference-in-Differences for Campaign Success

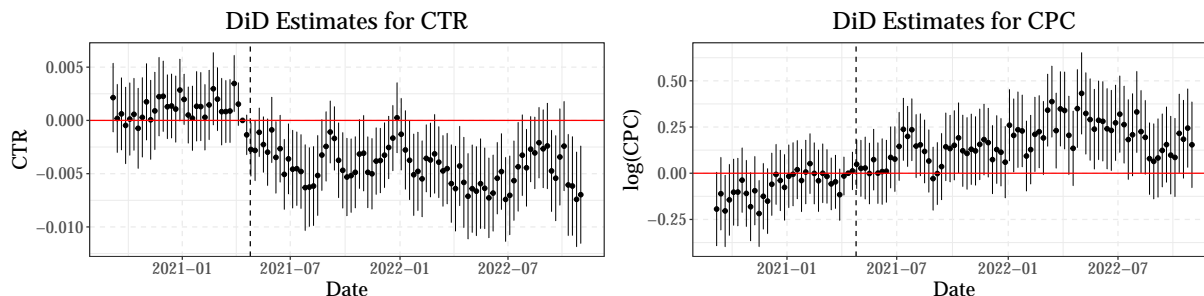


⁸We show using the same within-advertiser specification in Appendix [A.3.1](#) that there is little intensive-margin substitution across on-platform objectives after ATT.

For the outcome variable Y_{ijt} , we first consider the relative success of the respective campaigns in completing their objective. For conversion-optimized campaigns, an off-platform purchase is registered as a success, whereas for link click-optimized campaigns, an on-platform click is registered as a success. Thus, we compare the performance of these two types of campaigns in terms of the cost per measured success and the fraction of impressions that convert to a success.⁹

Even though conversion-optimized campaigns are potentially subject to the measurement issue and the clicks-optimized campaigns are not, the results presented in Figure 3 allow us to rule out that the aforementioned trends are purely time trends. Indeed, Figure 3 shows that there is a clear and sudden drop (resp. increase) in the success rate (resp. cost per success) of the campaigns that are optimized for conversions with strong evidence for parallel trends across the two groups.

Figure 4: Difference-in-Differences for Click-Through Rate and CPC



This result, however, is still subject to the aforementioned measurement issues. Thus, to overcome this we consider the success of both campaigns using only *on-platform* measures. In particular, we look at the impact of both campaigns on clicks. We use the same specification for click-through rate and cost per click with the results in Figure 4. These estimates show a consistent pattern with no evidence for pre-trends for either measure until the onset of ATT and then a sudden decline that results in a steady treatment effect over time, once ATT adoption reaches high enough levels. The overall estimated effect is a 0.0052 reduction in click-through rates, which represents a 37.1% reduction (in light of baseline click-through rate 0.014), for off-platform optimized campaigns and provides causal evidence for a reduction in the targeting effectiveness of off-platform conversion optimized campaigns after ATT.

5.3 Impact on Product Sales using Shopify

The previous results suggest that the effectiveness of Facebook advertisements was negatively impacted by ATT. To draw broader implications about the policy's impact, we need to understand its impact on downstream sales. In principle, advertisers could mitigate the negative shock by

⁹For this analysis drop a small number of erroneous observations where the number of conversions is larger than the number of impressions.

substituting away from Facebook for other platforms and, after some adjustment period, they could re-optimize their marketing strategies to have little impact on their bottom line. However, for some Facebook-reliant advertisers, targeting provided by other platforms may not be as effective, in which case the substitution may be costly.

To study how the advertisers coped with the shock, we analyze the subset of advertisers that also connect their Shopify dashboards, which provides us with a direct measure for online sales and allows us to measure whether more Facebook-dependent advertisers’ sales were more adversely impacted.

We utilize *across-advertiser* differences in reliance on Facebook advertising before the onset of ATT. While advertising allocations are endogenous, one can argue that the different advertisers have optimized their user acquisition strategy pre-ATT so that they are differentially reliant on the targeting capabilities enabled by Facebook, and, depending on the reliance on the Facebook ads, an advertiser may find it costly to shift advertising strategies post-ATT.

Table 1: Estimates on Sales

Estimation Method	<i>Dependent variable:</i>			
	log(Orders)	log(Revenue)	Repeated Order Ratio	log(Average Order Value)
	(1)	(2)	(3)	(4)
DiD	-0.242* (0.129)	-0.230* (0.131)	3.994*** (1.415)	0.015 (0.028)
DiD (w/ asinh(Ad Spending) Control)	-0.171 (0.118)	-0.156 (0.119)	3.563** (1.419)	0.019 (0.027)

*p<0.1; **p<0.05; ***p<0.01

NOTES: The rows present the estimated average treatment effect coefficient using the difference-in-differences specification with and without controls for asinh(total advertising spending). Standard errors are clustered at the advertiser level.

To study how the differential reliance on Facebook advertising translates into different abilities to cope with the shock, we adopt a difference-in-differences specification:

$$Y_{it} = \beta \left(Post_t \times T_i \right) + \alpha_i + \kappa_t + \epsilon_{it}, \quad (4)$$

where the notation is similar to before, except that T_i indicates whether they are a high-share Facebook advertiser. We use a monthly aggregation of the Shopify data and compute each advertiser’s pre-ATT market share of Facebook advertising relative to Google and TikTok based on our primary data. Since the distribution of advertisers in the sample is skewed towards Facebook usage, we consider advertisers with above (resp. below) the mean Facebook share in the treatment (resp. control)

group.¹⁰ This specification also lends itself to considering synthetic controls (Abadie et al., 2010) and synthetic difference-in-differences (Arkhangelsky et al., 2021) as robustness checks. Similar to the within-advertiser specification, we also consider a time-varying version of this in order to assess the reasonableness of the parallel trends assumption. Furthermore, we ensure that we have a balanced panel so that we only consider advertisers with at least one order in each of the months that we consider which leaves us with 222 advertisers.

Table 1 presents the results for the outcome variables that we observe from Shopify: log(Revenue), log(Order Count), Repeat Order Ratio, and log(Average Order Value). The first row presents the results using the difference-in-differences empirical specification (4). The takeaway across each of these is consistent: there is a reduction in orders and revenue of 23-24%, the fraction of total orders coming from repeat customers has increased, and the average order value is unchanged. One important caveat is that the results are on revenues and don't control for possible shifts in costs, such as changing advertising spend. To understand whether this is simply a result of shifting advertising spend, we estimate the difference-in-differences specification controlling for the inverse hyperbolic sine transform of advertising spend. These results are presented in the second row of Table 1. While this reports similar effect sizes for the ratio of repeat customers, we no longer find statistically significant reductions in revenues or orders though we still find sizeable and negative point estimates.

Table 2: Estimates on New vs. Repeat Customers

	<i>Dependent variable:</i>			
	asinh(New Customer Orders) (1)	asinh(New Customer Revenue) (2)	asinh(Repeat Customer Orders) (3)	asinh(Repeat Customer Revenue) (4)
DiD	-0.344*** (0.131)	-0.335** (0.135)	-0.178 (0.149)	-0.354* (0.212)
DiD (w/ asinh(Ad Spending) Control)	-0.262** (0.119)	-0.247** (0.121)	-0.118 (0.141)	-0.283 (0.204)

*p<0.1; **p<0.05; ***p<0.01

NOTES: The first two columns are asinh(orders) and asinh(revenue) coming from new customers. The second two columns are asinh(orders) and asinh(revenue) coming from repeat customers. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification with and without controls for asinh(total advertising spending). Standard errors are clustered at the advertiser level.

These results suggest that firms' abilities to acquire new customers may have been hampered after ATT, as the fraction of their total orders that come from repeat customers increased. Indeed, the pre-ATT baseline for the share of orders coming from repeat customers was 33.93, indicating that the estimates imply a 10.5% increase in the share of orders coming from repeat customers. To further investigate this, we construct the number of orders coming from new and repeat customers

¹⁰The cutoff is a market share of 0.644.

respectively. Furthermore, given that we estimated a precise null effect on the change in average order value, we assume that there is no difference in average order value between new and repeat customers to further impute the revenues from new and repeat customers using the average order value in each period and the number of orders for new and repeat customers.

We then estimate the same set of specifications on the changes in orders for revenue and orders, but broken down into revenues and orders from new and repeat customers respectively. The results are presented in [Table 2](#) and present a clear picture. Columns (1) and (2) show that there is a 26.2% decrease in orders coming from new customers as well as a corresponding decline in revenue of 24.7%, even when controlling for the overall amount of ad spend. Furthermore, columns (3) and (4) show that while there is a negative point estimate on repeat customer orders and revenue, it is more imprecise and less impacted relative to new customers. We assess the reasonableness of parallel trends in the following two ways: we present time-varying estimates for the difference-in-differences specification and we consider the results for synthetic control ([Abadie et al., 2010](#)) for $\log(\text{Orders})$, Repeated Customer Order Ratio, $\text{asinh}(\text{New Customer Orders})$, and $\text{asinh}(\text{New Customer Revenue})$. [Figure 6](#) considers the time-varying difference-in-differences specification with ad spending controls and provides some evidence that the parallel trends assumption seems to reasonably hold. Furthermore, [Figure 7](#), [Table 5](#), and [Table 6](#) finds consistent patterns and effect sizes using synthetic controls and synthetic difference-in-differences respectively.

These results suggest that ATT had an economically meaningful negative impact on firms that are reliant on Facebook advertising. Furthermore, since our sample period considers nearly a year and a half after the onset of ATT this was not a temporary loss due to an adjustment period, but rather a persistent reduction in sales. Several results suggest that the mechanism for the sales reduction comes from a loss in the ability to precisely target advertisements. First, the fraction of orders coming from repeated customers increases and we primarily find effects on the orders from new, not repeated, customers. Second, we find similar effect sizes for revenues and orders as well as a precise null effect on the average order value.

5.4 Substitution for Other Advertising Platforms

The results from [Section 5.3](#) indicate that, despite any substitution we observe, it did not allow Facebook-dependent advertisers to achieve similar levels of profitability as they were able to before ATT. Nevertheless, we can exploit ATT as an exogenous shock (to the efficacy of Facebook advertising) that allows us to investigate how advertisers substituted across different platforms and advertising modalities. This teaches us about which types of advertising are viewed as closer (albeit imperfect) substitutes to Facebook. Ascertaining a clear causal effect of these is difficult as quality, quantities, and prices are adjusting simultaneously in the new equilibrium. Therefore, we

provide mainly descriptive evidence on substitution patterns.

Table 3: Advertising Platform Usage Post-ATT

Estimation	Platform	<i>Dependent variable:</i>		
		Spend Share (1)	Impression Share (2)	Click Share (3)
DiD	Google	0.048*** (0.009)	0.067*** (0.010)	0.057*** (0.009)
DiD	Facebook	-0.044*** (0.009)	-0.061*** (0.010)	-0.057*** (0.009)
DiD	TikTok	-0.004 (0.002)	-0.005 (0.003)	-0.001 (0.003)
DiD	Google Search	0.011 (0.008)	0.012 (0.008)	0.012 (0.008)
DiD	Google Display	0.005*** (0.002)	0.028*** (0.006)	0.013*** (0.004)

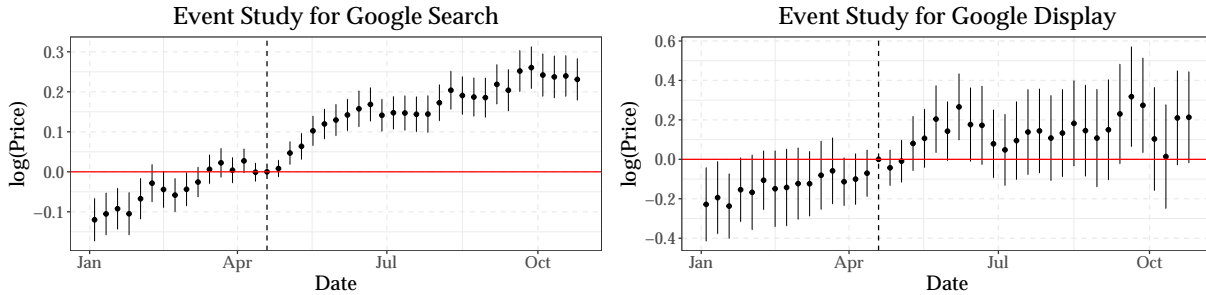
*p<0.1; **p<0.05; ***p<0.01

NOTES: Each cell displays the estimated average treatment effect across each of the specifications. The first three rows present the results of the difference-in-differences specifications for share of Spend, Impressions, and Clicks on overall spending on Facebook, TikTok, and Google. The final two rows present the same dependent variables for Google Search and Google Display products. Standard errors are clustered at the advertiser level.

We focus on the effect of ATT on the relative market shares of advertising platforms.¹¹ We use the across-advertiser difference-in-differences specification (4) with the same definition of treatment vs. control as before based on the mean Facebook share. We estimate this specification to control for possible seasonal trends in different advertising modalities that would not be captured with the within-advertiser specification. We use a weekly balanced panel of advertisers who spend non-zero dollars on any advertising platform throughout the same sample period as before (September 2020-October 2022) and our outcome variables are the relative market share of impressions, clicks, and spend across the different platforms. The first three rows of Table 3 show the results for the relative market shares of Google, Facebook, and TikTok. They suggest that for each of the measures that we consider Google benefited at the expense of Facebook, gaining 0.048-0.067 market share, whereas there was no shift in market share to TikTok. We consider a synthetic difference-in-differences specification as a robustness check with quantitatively similar results reported in Table 8.

¹¹In Appendix A.3.1 we explore substitution within campaign objectives on Facebook itself using the within-advertiser difference-in-differences specification and find little substitution away from off-platform optimized campaigns.

Figure 5: Event Studies for Google Advertising Prices



The Google advertising ecosystem is expansive – including contextual (search) and behaviorally targeted advertising (display). Relative to the Facebook ecosystem, the behaviorally targeted advertising of Google is less reliant on mobile and was expected to be less impacted by ATT than Facebook. This leads us to our next question: which types of Google services gain in market share? We study substitution towards the two most prominent Google advertising services: Google Search and Display. For this analysis, we are impacted by the fact that Google’s newest “smart”, dynamic allocation product, Performance Max, picked up substantial market share within the Google ecosystem following its rollout on November 2nd, 2021.¹² As such, we restrict our analysis to a tight window around ATT from January 1st, 2021 until October 31st, 2021, which also excludes the 2020 holiday period and ensures that we are picking up the substitution effects of ATT, not the Google product rollout.

The fourth and fifth rows of Table 3 display the shift in market share and find positive, but imprecise, increases in market share for Google Search with a more precisely estimated and statistically significant increase across all measures for Google Display. Thus, there appears to be a reallocation of shares mainly towards Google Display but to some extent to Google Search as well, consistent with the view that Facebook’s ads are a closer substitute to the former. Furthermore, there appears to be a discrepancy between the impression/click share and the spend share. As such, we also explore the time trends for the prices (CPC for search, CPM for display) of the two services with under the specification (1) in Figure 5. A challenge is that we only observe prices when advertisers buy advertisements from a particular platform and so to reduce possible selection, we use a balanced panel of advertisers with non-zero spend in the respective Google service. While we cannot rule out that the results are driven by seasonal time trends, they suggest that there is a price increase for Google Search and a smaller price increase for Google Display.¹³

¹²See <https://blog.google/products/ads-commerce/performance-max/> for additional information on the roll-out.

¹³We do not provide similar plots for Facebook since it’s possible to pay per click or per impression on the platform and we cannot directly observe the price paid.

6 Conclusion

In this paper, we use a dataset of online advertising performance and product sales to study the effects of Apple’s App Tracking Transparency Policy. We find that ATT significantly impacted the ability of Facebook advertisers to acquire customers and harmed the product sales of Facebook-dependent advertisers in the medium-term, even though they could, and often did, substitute for other advertising platforms. This finding suggests unique qualities of Facebook-enabled targeted advertising that at least some advertisers found it difficult to substitute away. While our results provide a quantification of the firm-side costs of ATT, a full policy evaluation of ATT would include the welfare effects on consumers. Finally, our results point to equilibrium effects across different advertising platforms and an exciting avenue for future work is to better understand the degree of substitutability of various online advertising modalities.

References

- Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American Statistical Association* 105(490), 493–505.
- Ambrus, A., E. Calvano, and M. Reisinger (2016). Either or both competition: A “two-sided” theory of advertising with overlapping viewerships. *American Economic Journal: Microeconomics* 8(3), 189–222.
- Anderson, S. P. and S. Coate (2005). Market provision of broadcasting: A welfare analysis. *The Review of Economic Studies* 72(4), 947–972.
- Aridor, G., Y.-K. Che, and T. Salz (2023). The effect of privacy regulation on the data industry: Empirical evidence from gdpr. *RAND Journal of Economics*.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review* 111(12), 4088–4118.
- Athey, S., E. Calvano, and J. S. Gans (2018). The impact of consumer multi-homing on advertising markets and media competition. *Management Science* 64(4), 1574–1590.
- Baslandze, S., J. Greenwood, R. Marto, and S. Moreira (2023). The expansion of varieties in the new age of advertising. *Review of Economic Dynamics*.
- Bian, B., X. Ma, and H. Tang (2021). The supply and demand for data privacy: Evidence from mobile apps. *Available at SSRN*.

- Bleier, A. and M. Eisenbeiss (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science* 34(5), 669–688.
- Cheyre, C., B. T. Leyden, S. Baviskar, and A. Acquisti (2023). The impact of apple’s app tracking transparency framework on the app ecosystem. *Available at SSRN 4453463*.
- Demirer, M., D. J. Hernández, D. Li, and S. Peng (2023). Data, privacy laws, and firm production: Evidence from gdpr.
- Ganglmair, B., J. Krämer, and J. Gambato (2023). Regulatory compliance with limited enforceability: Evidence from privacy policies. *Available at SSRN*.
- Gentzkow, M., J. M. Shapiro, F. Yang, and A. Yurukoglu (2022). Pricing power in advertising markets: Theory and evidence. Technical report, National Bureau of Economic Research.
- Goldberg, S. G., G. A. Johnson, and S. K. Shriver (2023). Regulating privacy online: An economic evaluation of the gdpr. *American Economic Journal: Economic Policy*.
- Goldfarb, A. and C. Tucker (2011a). Advertising bans and the substitutability of online and offline advertising. *Journal of Marketing Research* 48(2), 207–227.
- Goldfarb, A. and C. Tucker (2011b). Search engine advertising: Channel substitution when pricing ads to context. *Management Science* 57(3), 458–470.
- Goldfarb, A. and C. E. Tucker (2011c). Privacy regulation and online advertising. *Management Science* 57(1), 57–71.
- IAB (2023). Internet Advertising Revenue Report. Available at https://www.iab.com/wp-content/uploads/2023/04/IAB_PwC_Internet_Advertising_Revenue_Report_2022.pdf.
- Janssen, R., R. Kesler, M. E. Kummer, and J. Waldfogel (2022). Gdpr and the lost generation of innovative apps. Technical report, National Bureau of Economic Research.
- Jia, J., G. Z. Jin, and L. Wagman (2021). The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science* 40(4), 661–684.
- Johnson, G., R. A. Lewis, and E. Nubbemeyer (2017). The online display ad effectiveness funnel & carryover: Lessons from 432 field experiments. *Available at SSRN 2701578*.
- Johnson, G. A., S. K. Shriver, and S. G. Goldberg (2023). Privacy and market concentration: intended and unintended consequences of the gdpr. *Management Science*.

- Kesler, R. (2022). The impact of apple’s app tracking transparency on app monetization. *Available at SSRN 4090786*.
- Li, D. and H.-T. Tsai (2022). Mobile apps and targeted advertising: Competitive effects of data exchange. *Available at SSRN 4088166*.
- Peukert, C., S. Bechtold, M. Batikas, and T. Kretschmer (2022). Regulatory spillovers and data governance: Evidence from the gdpr. *Marketing Science 41(4)*, 746–768.
- Rafieian, O. and H. Yoganarasimhan (2021). Targeting and privacy in mobile advertising. *Marketing Science 40(2)*, 193–218.
- Rochet, J.-C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association 1(4)*, 990–1029.
- Sahni, N. S., S. Narayanan, and K. Kalyanam (2019). An experimental investigation of the effects of retargeted advertising: The role of frequency and timing. *Journal of Marketing Research 56(3)*, 401–418.
- Utz, C., M. Degeling, S. Fahl, F. Schaub, and T. Holz (2019). (un) informed consent: Studying gdpr consent notices in the field. In *Proceedings of the 2019 acm sigsac conference on computer and communications security*, pp. 973–990.
- Werner, G. (2022). The state of the direct-to-consumer (dte) industry. <https://influencermarketinghub.com/direct-to-consumer-industry/>.
- Wernerfelt, N., A. Tuchman, B. Shapiro, and R. Moakler (2022). Estimating the value of offsite data to advertisers on meta. *University of Chicago, Becker Friedman Institute for Economics Working Paper (114)*.
- Wilbur, K. C. (2008). A two-sided, empirical model of television advertising and viewing markets. *Marketing Science 27(3)*, 356–378.
- Zhao, Y., P. Yildirim, and P. K. Chintagunta (2021). Privacy regulations and online search friction: Evidence from gdpr. *Available at SSRN 3903599*.

Online Appendix

A Additional Figures and Tables

A.1 Omitted Acquisition Cost Figures and Tables

Table 4: Heterogeneous Effects on Acquisition Costs and CTR

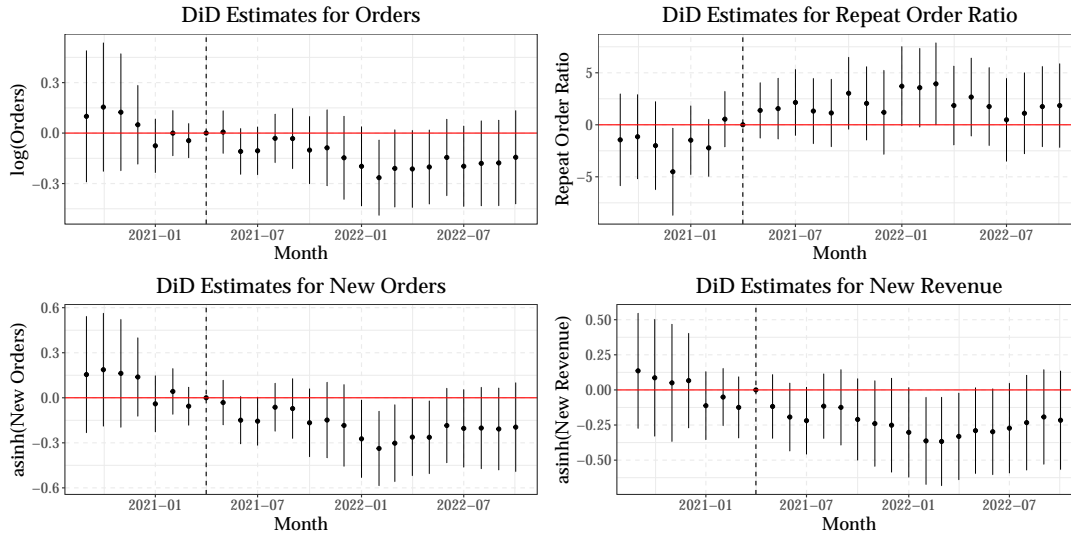
	<i>Dependent variable:</i>					
	log(Acquisition Cost)			CTR		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	0.541*** (0.019)	0.582*** (0.022)	0.510*** (0.023)	-0.001** (0.0003)	-0.001** (0.0003)	-0.0001 (0.0002)
$Post_t \times$ Large Advertiser		-0.162*** (0.039)			0.001** (0.0004)	
$Post_t \times$ Prospecting			0.056*** (0.020)			0.0003 (0.0003)
$Post_t \times$ Retargeting			-0.026 (0.021)			-0.001*** (0.0003)
Prospecting			-0.180*** (0.020)			-0.0004 (0.0004)
Retargeting			-0.584*** (0.023)			-0.001*** (0.0003)
Week FE	No	No	No	No	No	No
Advertiser FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	61,091	61,091	150,677	64,297	64,297	164,689
R ²	0.822	0.823	0.757	0.672	0.672	0.541
Adjusted R ²	0.821	0.822	0.756	0.669	0.669	0.539

*p<0.1; **p<0.05; ***p<0.01

NOTES: The dependent variable for the first three columns is log(Acquisition Cost) and the last three columns is CTR. The first column presents the results of the coefficient for the increase in acquisition cost after ATT. The second column considers heterogeneity by whether the advertiser was large (above median spending pre-ATT) or small. The third column considers heterogeneity by whether the campaign was a lookalike (held out group), prospecting, or retargeting campaign. The fourth through sixth columns represent the same specifications for the CTR outcome variable. The number of observations differs between columns (1)-(3) and (4)-(6) since log(Acquisition Cost) is not defined when the number of conversions is zero. Standard errors for each regression are clustered at the advertiser level.

A.2 Omitted Sales Outcome Figures and Tables

Figure 6: Difference-in-Differences for Sales Outcomes



NOTES: The outcome variables on the top row from left to right: $\log(\text{Orders})$, Repeated Customer Order Ratio. The outcome variables on the bottom row from left to right: $\text{asinh}(\text{New Customer Orders})$, $\text{asinh}(\text{New Customer Revenue})$. Plots are the difference-in-differences specification including controls for the $\text{asinh}(\text{Total Advertising Spend})$.

Table 5: Estimates on Sales

Estimation Method	<i>Dependent variable:</i>			
	$\log(\text{Orders})$	$\log(\text{Revenue})$	Repeated Order Ratio	$\log(\text{Average Order Value})$
	(1)	(2)	(3)	(4)
SynthDiD	-0.182** (0.091)	-0.177** (0.088)	3.12*** (1.313)	0.023 (0.028)

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

NOTES: The rows present the estimated average treatment effect coefficient using the synthetic difference-in-differences specification. Standard errors are clustered at the advertiser level.

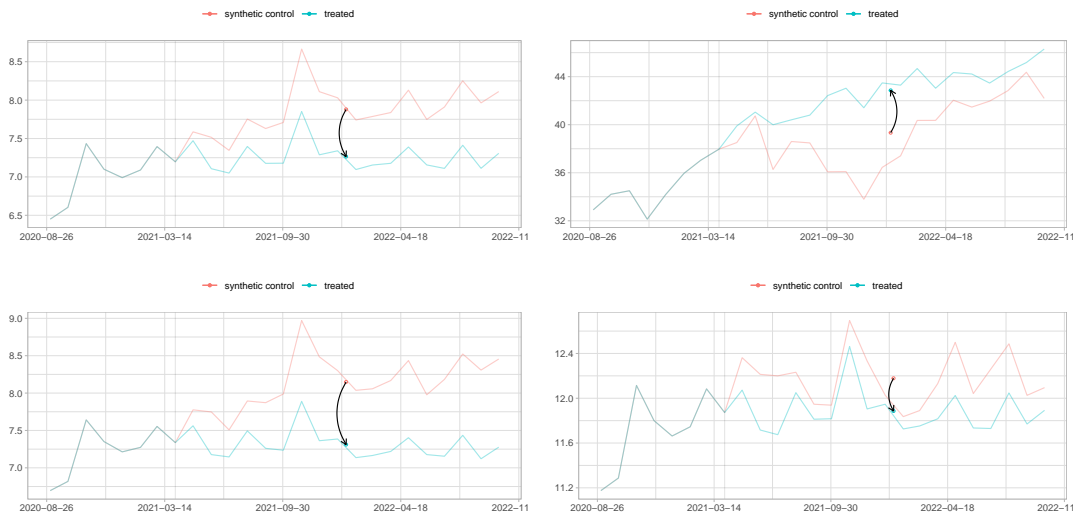
Table 6: Estimates on New vs. Repeat Customers

<i>Dependent variable:</i>				
	asinh(New Customer Orders)	asinh(New Customer Revenue)	asinh(Repeat Customer Orders)	asinh(Repeat Customer Revenue)
	(1)	(2)	(3)	(4)
SynthDiD	-0.248*** (0.100)	-0.235*** (0.097)	-0.172* (0.104)	-0.177 (0.108)

*p<0.1; **p<0.05; ***p<0.01

NOTES: The first two columns are asinh(orders) and asinh(revenue) coming from new customers. The second two columns are asinh(orders) and asinh(revenue) coming from repeat customers. The rows present the estimated average treatment effect coefficient using the synthetic difference-in-differences. Standard errors are clustered at the advertiser level.

Figure 7: Synthetic Control for Sales Outcomes



NOTES: The outcome variables on the top row from left to right: log(Orders), Repeated Order Ratio. The outcome variables on the bottom row from left to right: asinh(New Customer Orders), asinh(New Customer Revenue). All plots are plotting the synthetic control specifications for the respective variables.

A.3 Additional Platform Substitution Results

A.3.1 On-Platform Substitution

Table 7: Facebook Objective Substitution

	<i>Dependent variable:</i>		
	Spend Share	Impression Share	$\mathbb{1}(\text{Spend}_t > 0)$
	(1)	(2)	(3)
$Post_t \times$ On Platform Actions	0.004 (0.007)	0.015 (0.009)	0.028** (0.012)
$Post_t \times$ On Platform Reach	0.001 (0.005)	0.015* (0.008)	0.026*** (0.009)
On Platform Actions	-0.885*** (0.011)	-0.834*** (0.013)	-0.697*** (0.015)
On Platform Reach	-0.921*** (0.007)	-0.868*** (0.011)	-0.844*** (0.012)
Week FE	Yes	Yes	Yes
Advertiser FE	Yes	Yes	Yes
Observations	212,892	212,892	212,892
R ²	0.899	0.821	0.614
Adjusted R ²	0.899	0.820	0.613

*p<0.1; **p<0.05; ***p<0.01

NOTES: The dependent variables are the share of spend (column 1), share of impressions (column 2), and whether there is non-zero spend (column 3). The left out category is off platform actions. Standard errors clustered at the advertiser level.

We study substitution *within* objectives on the Facebook platform. One measure of advertiser substitution would be to re-optimize marketing campaigns towards on-platform as opposed to off-platform actions since they can plausibly serve as proxies for relevant off-platform actions, while retaining the targeting feedback loop by having a direct measure of the on-platform outcome. We characterize the different Facebook objectives into three groups: Off Platform Conversions, On Platform Actions, On Platform Reach. For off platform conversions we consider campaigns with one of the following objectives: Conversions, Sales Outcomes, Product Catalog Sales, App Installs, App Promotion. For On Platform actions we consider: link clicks, store visits, page likes, leads outcome, traffic outcomes, engagement outcomes, and post engagement. For On Platform Reach we consider video views, brand awareness, reach, and awareness outcome.¹⁴ We consider three outcome measures: the campaign objectives' share of spending, share of impressions, and an indicator for whether spend was non-zero. The former two measures provide a measure of intensive margin substitution – to what extent do advertisers shift their share of spending more towards on

¹⁴This includes all campaign objectives except for messages and event responses since it's ambiguous how to categorize them.

platform actions – and the final measures provide a measure of extensive margin substitution – to what extent do advertisers start to run on-platform campaigns.

We consider a balanced panel of Facebook advertisers that have positive spend on off-platform optimized campaigns for the entirety of the sample period and use the within-advertiser difference-in-differences specification (3). The results are reported in Table 7, which show a precise null effect on substitution towards on-platform objectives apart from an economically small degree of substitution in the extensive margin towards objectives optimizing for on-platform reach.

A.3.2 Omitted Figures and Tables

Table 8: Advertising Platform Usage Post-ATT (SynthDiD)

Estimation	Platform	<i>Dependent variable:</i>		
		Spend Share (1)	Impression Share (2)	Click Share (3)
SynthDiD	Google	0.045*** (0.008)	0.063*** (0.009)	0.052*** (0.008)
SynthDiD	Facebook	-0.042*** (0.008)	-0.057*** (0.009)	-0.050*** (0.008)
SynthDiD	TikTok	-0.002 (0.002)	-0.003 (0.003)	0.001 (0.003)
SynthDiD	Google Search	0.006 (0.007)	0.007 (0.007)	0.011 (0.007)
SynthDiD	Google Display	0.004** (0.0015)	0.025*** (0.006)	0.011*** (0.0035)

*p<0.1; **p<0.05; ***p<0.01

NOTES: Each cell displays the estimated average treatment effect across each of the specifications. The first three rows present the results of the synthetic difference-in-differences specifications for share of Spend, Impressions, and Clicks on overall spending on Facebook, TikTok, and Google. The final two rows present the same dependent variables for Google Search and Google Display products. Standard errors are clustered at the advertiser level.