Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency *

Guy Aridor[†] Yeon-Koo Che[‡] Brett Hollenbeck[§]

Maximilian Kaiser [¶] Daniel McCarthy [∥]

April 2025

Abstract

Assembling novel datasets on online advertiser spending, performance, and revenue, we quantify the economic effects of Apple's App Tracking Transparency (ATT) privacy policy on e-commerce firms. We find that conversion-optimized Meta advertisements, affected most by ATT, saw a 37% reduction in click-through rates after ATT. While firms responded by shifting ad spend from Meta to the Google ecosystem, firms with higher baseline Meta dependence nevertheless experienced a substantial decline in firm-wide revenue relative to firms with lower baseline Meta dependence. We quantify these effects using a variety of methods, finding revenue decreases in the range between 8-40% relative to less exposed firms. These declines were primarily borne by smaller e-commerce firms, raising questions about the trade-offs between consumer privacy and the ability of smaller e-commerce and direct-to-consumer firms to succeed in the product market.

^{*}Authors listed in alphabetical order. An earlier version of this paper was circulated as "Privacy Regulation and Targeted Advertising: Evidence from Apple's App Tracking Transparency". We thank the Marketing Science Institute, the UCLA Price Center for Entrepreneurship and Innovation, and the Law & Economics Center Program on Economics & Privacy for generous funding. We thank Tobias Salz for his early contributions to this project as well as Randy Bucklin, Sam Goldberg, Garrett Johnson, and Silvio Ravaioli for helpful comments. We thank audiences at the Cornell Johnson, Digital Economics Paris Seminar, Econometric Society Interdisciplinary Frontiers Economics and AI/ML Meeting, Google Economics, GMU Program on Economics & Privacy Empirical Research Workshop, IIOC, Marketing Strategy Meets Wall Street Conference, Platform Competition Workshop (USC), Toulouse Economics of Platforms Seminar, UC Boulder Leeds, and UCLA Anderson for helpful comments. We further thank Maurice Rahmey, Nils Wernerfelt and Michael Zoorob for answering questions about advertising on Meta, as well as Andrew D'Amico, Chelsea Mitchell, Sangwook Suh, and Kevin Wang for excellent research assistance. Furthermore, we thank Carol Doyle for her help with procuring the Kantar Vivvix dataset. All errors are our own.

[†]Northwestern University, Kellogg School of Management; guy.aridor@kellogg.northwestern.edu

[‡]Columbia University; yeonkooche@gmail.com.

[§]UCLA Anderson School of Management; brett.hollenbeck@anderson.ucla.edu

¹Hamburg University, Grips Intelligence; maximilian.kaiser@gripsintelligence.com

^IUniversity of Maryland College Park, Robert H. Smith School of Business; dmccar@umd.edu

1 Introduction

Digital advertising comprises the largest share of advertising spending at U.S. firms, surpassing both TV and print advertising in 2019 and reaching \$506 billion in total spending worldwide in 2021.¹ Its growth is driven by two factors: precise targeting using consumer data and real-time performance measurement.² These capabilities reduced customer acquisition costs for the e-commerce industry and fostered the rise of direct-to-consumer (DTC) firms by efficiently matching these firms to their audiences. This technology, therefore, has potentially large positive welfare ramifications by enabling the existence of these firms.

However, many consumers and privacy advocates have raised concerns about the tracking of user behavior by online platforms and advertising intermediaries, including third parties not explicitly authorized by users. In response, privacy protection measures such as the EU's General Data Protection Regulation (GDPR) and Apple's App Tracking Transparency (ATT) have sought to limit firms' access to consumer data. These policies may harm firms by reducing their ability to target consumers, potentially leading to welfare losses for both firms and customers.

In this paper, we quantify the economic costs of one such privacy protection measure – Apple's App Tracking Transparency (ATT) – which allows Apple iOS users to opt out of data sharing across their apps (third-party data sharing) by prompting users to either allow or disallow data sharing. The vast majority (80–85%) of users opted out when prompted (Baviskar et al., 2024, Chen, 2021, Laziuk, 2021), disrupting platforms' ability to measure and target ads effectively. Our analysis focuses on three key questions. First, how did ATT impact advertising effectiveness across Meta and Google?³ In principle, other mobile app advertising platforms such as Snapchat and TikTok are also impacted by the ATT policy change, but as Meta and Google are the dominant players in this industry, we focus

¹https://www.emarketer.com/content/digital-ad-spend-worldwide-pass-600-billion-this-year/

 $^{^{2}}$ Unlike general brand advertising, which builds equity over time without direct response metrics (Borkovsky et al., 2017).

³Here and throughout, we use the term Meta advertising to refer to advertising done on Facebook, Instagram, and the Meta Audience Network as we do not distinguish between these platforms in our analysis.

on them. Second, how did firms reallocate advertising spending between these platforms? Finally, how did these changes affect firm revenues? By quantifying revenue impacts and understanding the corresponding effects in the advertising market that contribute to it, we provide a comprehensive assessment of ATT's economic consequences.

We combine two unique sources of data on firm advertising performance and revenue to answer these questions. One comes from an anonymous data provider that enables a granular view of advertising spending and performance across Meta, Google, and TikTok for 1,221 firms, which we denote throughout as the *advertising* dataset. The other comes from Grips Intelligence,⁴ a leading data analytics and market intelligence firm, providing transaction and revenue data for 773 firms, which we denote throughout as the *revenue* dataset.

Regarding the effects of ATT on advertising performance, we show that sales conversions observed by Meta drop in line with the gradual adoption of the iOS version that includes ATT. We then perform a within-firm analysis to estimate the causal effect of ATT on forms of advertising that are reliant on off-platform data. We compare Meta campaigns optimized for off-platform conversions (which were impacted by ATT) versus on-platform clicks (which were not) and find a 36.6% relative reduction in click-through rates for conversion-optimized campaigns (95% CI: 18.2% to 54.5%). Additionally, Meta's online advertising spending share declined by 4.4%, with the majority of this shift benefiting Google, which was less affected by ATT. Together, these analyses show that, for an important class of e-commerce firms, the performance of conversion-optimized Meta advertising was significantly degraded due to ATT and that there was some equilibrium adjustment as a result.

We then explore the downstream implications of this on firm revenues using the revenue dataset. While measuring this impact is important, we face several significant empirical challenges. First, ATT impacts all firms simultaneously. There is no staggered rollout across advertisers or set of fully-exempted firms that could be used as a control group. Second, there is substantial variation in revenue across e-commerce firms as well as within these firms

⁴https://gripsintelligence.com/

over time, making even clean designs subject to noisy estimates. And third, the extent to which firms were exposed to the policy shock is determined by the extent to which they rely on targeted advertising and iOS consumers, both of which are measured with different degrees of measurement error.

We approach this challenge by making a set of comparisons of revenue before and after ATT by more vs less exposed firms. Rather than attempting to identify a single point estimate, we use different estimation approaches and two different measures of ATT-exposure to construct a set of bounds for the plausible range of the average treatment effect. We construct two measures of treatment for each firm based on exposure to ATT. The first is their pre-ATT reliance on Meta, measured as the average share of revenue attributable to Meta advertising in the year before ATT. The second is the average share of their revenue that comes from iOS users in the pre-period. Both iOS and Meta reliance expose firms more to the ATT policy change, but as these measures are not highly correlated, they capture ATT exposure in different ways. We begin by stratifying firms based on median splits of these treatment measures and conducting difference-in-differences analyses. We find that after ATT went into effect, firms that were more Meta-dependent saw a decrease in overall revenue by 37.1% relative to less Meta-dependent firms (95% CI: 12.4% to 55.1%) and firms that were more iOS-dependent saw a decrease in overall revenue by 40.1% relative to less iOS-dependent firms (95% CI: 18.0% to 57.1%).⁵

Next, we consider a specification that more explicitly addresses the concern that virtually all firms have at least some exposure to ATT by implementing the heterogeneous adoption design estimator of de Chaisemartin et al. (2024), which is specifically developed for settings where all units are treated but with different intensities. When using the Meta revenue share as the treatment variable, this method produces estimates that imply that relative to below-

⁵Breaking firms into smaller subgroups based on level of treatment, such as quartiles or terciles, and performing the same difference-in-differences exercise produce estimates with similar effect sizes that are generally increasing in the degree of treatment but in some cases suggest a non-linear relationship between treatment and effect magnitude. We attribute these instances to measurement error in the treatment variable rather than true non-linear relationship.

median treated firms, above-median treatment is associated with a post-ATT decrease in revenue of 8.1% (95% CI: -11.2% to -5.1%).⁶ When using iOS-dependence as the treatment variable, the HAD coefficients imply an 18.5% relative revenue reduction for the more ATT-exposed firms (95% CI: -26.9% to -10.1%). Because it is difficult to precisely attribute revenue to specific types of advertising (or mobile device), these treatment variables are subject to measurement error which may result in substantial attenuation bias pushing estimates towards zero.⁷ While not formally bounding the true effects, we believe these sets of estimates provide a reasonable range for the magnitude of ATT's impact on firm revenues.

Our results have several important policy and managerial implications. The large and negative impact on revenues indicates that opt-in privacy protection measures have a significant economic cost for firms that rely on targeted advertising for revenue generation, especially for smaller firms. The magnitude of the revenue reductions suggests that privacy protection measures can threaten the viability of business models, such as those of DTC firms that rely on targeted advertising, and that the cost of starting up such a business is now substantially higher because of ATT. While recognizing the potential welfare gains associated with added privacy protection, our results suggest there may be a countervailing effect on consumer welfare through this change in the composition of firms that can succeed in the product market. In addition, our results on the value of user data for targeted advertising have implications for the potential costs of the European Commission's prosecution of Meta's "pay or OK" practices for consumer data and lower-funnel tracking restrictions by Meta for health and wellness brands which went into effect in January 2025. Finally, while we do not directly observe Apple's advertising platform in our data, our results also

⁶We derive this from multiplying the coefficients by the average treatment levels among the treated and control units, and choose this comparison to be roughly equivalent to the effect size estimated in the median split difference-in-differences estimates described above.

⁷An additional implication is that these estimates show a decrease in revenue that is typically larger than the share of revenue attributed to Meta advertising and iOS users. This could result from the use of last-touch attribution, which significantly underestimates true revenue associated with Meta advertising, as well as the long-term loss in revenue from lack of customer acquisition. In addition, there may be spillover effects whereby the loss of a substantial amount of data makes targeting models less effective for all user types and advertisers. We discuss these issues in greater detail in section 4.

speak to ongoing antitrust concerns around the potentially anticompetitive impacts of ATT (CMA, 2022; Sokol & Zhu, 2021) by showing the impact on competing advertising platforms, especially Meta and Google.

Related Literature We contribute to a growing literature studying the economic costs of privacy protection measures (Acquisti et al., 2016; Dubé et al., 2024; Goldfarb & Que, 2023). This includes work on the cost to publishers and advertisers of the EU's General Data Protection Regulation (GDPR) (Aridor et al., 2023; Goldberg et al., 2024; Johnson, 2023; Lefrere et al., Forthcoming), the potential impact of limitations on cookies (Goldfarb & Tucker, 2011; Johnson et al., 2020; Kobayashi et al., 2024; Miller & Skiera, 2023), the iOS privacy nutrition labels (Bian et al., 2021), and the use of ad blockers (Todri, 2022; Yan et al., 2022).

Our work is most closely related to several papers that also study the impact of ATT. Wernerfelt et al. (2022) use internal access to Meta to run large-scale field studies studying the effectiveness of ad targeting in which they compare the performance of "offsite conversionoptimized" ad campaigns utilizing offsite data with the performance of ad campaigns treated with "link-click optimization" that make no use of offsite data. They find that removing the offsite data from targeting decreases targeting effectiveness and increases the median cost per incremental customer by 37%, with large effects for small businesses. We extend and complement these findings by measuring the comprehensive effects of ATT using observational data and thus directly incorporating possible equilibrium adjustments by firms and platforms after ATT. Indeed, we find comparable effect sizes on revenue and, as with several recent papers in the literature on the economic effects of privacy protection measures, similarly find that the negative effects are larger for smaller firms (Korganbekova & Zuber, 2023), as summarized in Dubé et al., 2024.

In contemporaneous work, Cecere and Lemaire (2023) also study the effect of ATT on predicted, aggregated ad outcomes and find that ATT reduced targeting efficiency on Meta. We complement this work by using platform-observed advertising data to similarly find a reduction in targeting efficiency and use our revenue data to quantify the downstream economic costs of reduced targeting efficiency. Another contemporaneous paper is Deisenroth et al. (2024), who study the industry-level effects of ATT, highlighting impacts on market entry/exit and producer prices for industries with higher exposure to ATT. We complement this work by using more granular firm-level data that provides information on advertising spending across all major online advertising platforms, not just Meta, directly observe firm revenue, and focus specifically on e-commerce retailers rather than broader cross-industry comparisons. Several other papers (Cheyre et al., 2023; Kesler, 2022; Kollnig et al., 2022; Kraft et al., 2023; Li & Tsai, 2022) also study the impact of ATT, but largely focus on the supply-side response of iOS applications to the regulation. These papers find that ATT reduced app downloads and the incentives to develop new applications and that some applications shifted from relying on advertising revenues to charging for their apps. We complement these papers by studying the effect on the advertisers themselves – as opposed to the application's advertising revenues.

2 Data and Context

2.1 Background on App Tracking Transparency

Apple announced in late 2020 that its new mobile operating system, iOS 14.5, would be rolled out the following year with a feature prompting users to explicitly consent to tracking by each app. This feature officially launched on April 25, 2021.⁸ Before this update, app publishers had access to an "identifier for advertisers" (IDFA), which was available by default on Apple devices. The update removed default access to this and instead prompted users, "Allow [app name] to track your activity across other companies' apps and websites?" (see Figure 1). For users selecting "Ask App Not To Track," the app can no longer use tracking

⁸https://techcrunch.com/2020/06/22/apple-ios-14-ad-tracking/

to observe what those users did after leaving the app.⁹



Figure 1: ATT Data Sharing Prompt

The IDFA had two primary uses for mobile display advertising via platforms such as Meta. First, it provided a view of consumer activity across applications, which could serve as an input for targeting. Second, it enabled Meta to link conversions to advertisements more easily.¹⁰ If a consumer opts out through ATT, however, Meta is unable to link ad impressions or clicks to purchases. This also means that Meta is limited in its ability to accurately report conversions to firms. Indeed, following ATT, Meta attempted to mitigate the impact by transitioning from deterministic to probabilistic attribution models, such as Aggregated Event Measurement, where they replaced actual observed conversions with "modeled" conversions for users that opted out.¹¹ Thus, both the loss in off-platform data and conversion measurement issues contribute to an overall degradation in targeting by reducing the data observed by firms (Johnson et al., 2022; Runge & Seufert, 2021).

⁹Unlike other privacy protection measures such as the GDPR, there were neither compliance issues (Ganglmair et al., 2023) nor heterogeneity in the design of the opt-in prompt (Utz et al., 2019) as a requirement for remaining on the App Store was to include the prompt provided by Apple.

¹⁰Effective targeting depends not only on the firm's targeting criteria but also on Meta optimizing within those criteria to identify individuals most likely to convert while the campaign is active (see https://www.facebook.com/business/help/950694752295474 for details).

¹¹Campaigns targeting non-impacted operating systems remained unchanged, but, if a campaign targeted iOS users, then Meta would change the recommended setup and targeting for the overall campaign. See https://www.facebook.com/business/help/331612538028890?id=428636648170202 for the full details.

2.2 Data Overview

We use detailed data on advertising and revenues for thousands of firms for our analyses. These data come from two distinct sources, both of which contain granular data from a set of firms that opt into our data providers for the purpose of analytics.

The first data source we denote as the *advertising* dataset, which comes from an anonymous advertising analytics provider, and provides granular data on Meta, Google, and Tik-Tok advertising spending and performance for 1,221 firms. The second data source we denote as the *revenue* dataset, which comes from Grips Intelligence and contains first-party Google Analytics traffic and revenue data for 773 firms across the globe at the firm-device-OS level. In the next two subsections, we provide detailed information on each dataset. We specify which data are used in each analysis in relevant table or figure notes.

2.2.1 Advertising Dataset (Anonymous Analytics Provider)

The advertising dataset contains weekly firm performance data for a separate set of firms. These firms, whose identities are anonymized, contract with the data provider and share their relevant performance data from Meta, Google, and TikTok. 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), and the total number of conversions associated with the advertising campaign (conversions). The measurement of the first three variables (spend, impressions, clicks) is not affected by ATT; they are measured accurately and consistently before and after ATT. However, conversions measurement is potentially affected by ATT as this is typically collected through a pixel that the firm embeds within its website or application that requires a consistent identifier across the platform of interest and the third-party website/app.¹² Within each advertising platform, we observe these data at different levels of granularity. For Meta, we observe performance broken down based on

¹²See 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.

campaign objectives (e.g., off-platform conversions, on-platform clicks). For Google, we observe performance broken down based on Google product (e.g., Google Search or Display). We present a set of summary statistics for the advertising dataset in Table 1, indicating that the mean online advertising spending is \$115,390 per month and that the online advertising share across different platforms heavily skews towards Meta.

2.2.2 Revenue Data (Grips Intelligence Data)

The revenue dataset consists mostly of classical online retailers in fashion, consumer electronics, beauty and cosmetics, and general e-commerce retail. Its data are derived from the firm's Google Analytics tracking, which relies only on first-party data to track relevant metrics. As a result, the measurement methodology for this dataset remains consistent and accurate regardless of ATT implementation. From the available firms, we selected a subset of variables – transactions, sessions, and revenue – aggregated at the device-operating-system-traffic-source-day level. The traffic source is determined using last-touch attribution.¹³

		Percentile			
Dataset	Metric	Mean	25th	50th	75th
Revenue dataset	Revenue (\$1,000)	4,896.95	119.65	359.36	1,349.32
	iOS share	0.25	0.14	0.24	0.34
	Android share	0.21	0.11	0.18	0.29
	Mobile share	0.44	0.29	0.45	0.58
	Meta share	0.04	0.00	0.01	0.04
Advertising dataset	Online ad spend (\$1,000)	115.39	7.56	24.92	93.12
	Meta Share	0.753	0.581	1.0	1.0

Table 1: Dataset Summary Statistics

NOTES: Revenue figures are reported in U.S. dollars and are computed using the revenue dataset over April 2020-April 2021. The revenue row presents the summary statistics across firms, where each firm is a single data point represented by its average monthly revenue. The "share" variables for the revenue dataset each refer to the share of revenue associated with each traffic source. Advertising statistics are computed using the advertising dataset over September 2020-April 2021. The online ad spend row presents the summary statistics across firms, where each firm is a single data point represented by its monthly average online advertising spending. The "Meta share" variable for the advertising data refers to the share of monthly average online advertising spending on Meta from the set of Meta, Google, and TikTok advertising.

We present firm-month level summary statistics for the revenue dataset during the pre-ATT period (April 2020 to April 2021) in Table 1. The distribution of monthly revenue

¹³Last-touch attribution, specifically last-non-direct-touch, assigns conversion credit to the final non-direct interaction before purchase. For example, if a customer clicks a Meta ad and then converts via an email link, the email is recorded as the last-touch source.

exhibits significant right skewness, with a median of \$359,000 and a mean of \$4.9 million. For the median firm, iOS sessions generate 25% of revenue compared to 18% from Android. Meta's revenue share averages 0.04, though the attribution methodology used to calculate this measure understates Meta's true contribution to revenue, as the revenue dataset uses last-touch attribution with a 30-minute window to assign credit for conversions to advertising channels. While this short attribution window affects the absolute magnitude of Meta's revenue contribution, it does not impact measures of relative platform dependence.¹⁴

2.2.3 Data Representativeness

As both datasets contain firms that opt into data sharing, a natural question arises regarding the set of participating firms and the broader population they represent. In Online Appendix E, we provide additional details about the incentives driving firms to opt into both the advertising and revenue datasets. We then benchmark our datasets against three "populationlevel" external sources that are minimally affected by firm-side selection: cohort-level public disclosures from Shopify (a widely used e-commerce platform), data from SimilarWeb (a provider of web traffic and performance metrics), and data from Kantar-Vivvix (an advertising intelligence company). Each of these external benchmarks is constructed to reflect a broad cross-section of e-commerce retailers. We compare the variation in firm size and temporal trends in our datasets to those observed in these external benchmarks.

We find that both datasets include firms spanning a wide range of the e-commerce size spectrum, although the advertising dataset includes relatively more smaller firms, while the revenue dataset includes relatively more larger firms. These differences in sample composition are important for interpretation, which is why we conduct heterogeneous treatment effect analyses in Section 4 that explicitly examine how ATT's impact varies with firm size. These analyses provide transparency about how effects may differ across various segments of the

 $^{^{14}}$ As shown in Table 2 of Gordon et al. (2023), one-hour attribution windows tend to significantly understate the true incremental impact of advertising. The 30-minute window used in our revenue dataset is even shorter, likely leading to more severe underestimation.

e-commerce industry.

Given these datasets and their respective compositions, our analyses proceed as follows. First, we use the advertising dataset to assess the impact of ATT on the efficacy of ad campaigns reliant on off-platform data through a within-firm comparison of off- versus onplatform ad campaign performance while controlling for unrelated factors (e.g., firm size and type). Next, we use the advertising dataset to provide evidence of how firms adapted their strategies in response to these changes. Finally, we analyze the revenue dataset to understand the downstream consequences of these changes on revenue within the e-commerce sector.

The effect sizes from these two sets of analyses are not directly comparable. The first component of the analysis quantifies degradation in the effectiveness of the affected forms of advertising. How these changes translate into firm-level revenue declines depends on the magnitude of each firm's reliance on the affected forms of advertising and their ability to adapt.

3 Impact on Advertising Effectiveness

We use the advertising dataset to investigate the effect of ATT on advertising performance.

Descriptive Evidence on Meta Conversion-Optimized Campaign Performance: We first examine suggestive evidence of ATT's impact on the number of conversions and cost per conversion for conversion-optimized Meta advertisements. We document in Tables OA1 and OA2 that these campaigns make up 95.7% of spending on Meta advertising within our sample before ATT.¹⁵ We restrict attention to a balanced panel of firms with Meta advertising spending from September 2020 until October 2022 and estimate the following specification:

$$Y_{it} = \sum_{t} \beta_t \cdot \text{Week}_t + \alpha_i + \epsilon_{it} \tag{1}$$

 $^{^{15}}$ Here and throughout the rest of the paper, conversion-optimized advertisements refer to campaigns that are optimizing for conversions, product catalog sales, or sales outcomes.



Figure 2: Event Study for Conversions and Cost per Conversion



NOTES: The figures plot the event study coefficients for log(cost per conversion) on the left and log(conversions) on the right using specification (1). We note that both of these variables have measurement issues after ATT. Table A.1 presents the associated aggregate post-ATT estimates. Standard errors are clustered at the firm level. The red dotted line in the left figure represents the estimated percentage of iOS devices that updated to iOS 14.5 over time. The first vertical dotted line represents April 25, 2021, when Apple first introduced iOS 14.5. Source: Gupta Media, https://lookerstudio.google.com/u/0/reporting/ 3d5dda40-37ea-4b9f-bd91-bb8df8e12620/page/aDUJC?s=kTs6iab_AhQ

where α_i denotes the firm fixed effects.

Figure 2 plots the estimated β_t for each week. Leaving aside the spikes around the holiday season, it is clear that after ATT the number of conversions drops dramatically and the cost per conversion increases.¹⁶ As suggestive evidence that this increase was caused by ATT, Figure 2 also plots the fraction of iOS devices that had installed iOS 14.5.¹⁷ The gradual adoption of iOS 14.5 coincides with a gradual increase in the cost per conversion that then nearly discontinuously increases as Apple nudged a large portion of users to adopt iOS 14.5 in early June, resulting in an overall 73.2% increase (95% CI: 66.9% to 79.7%) in cost per Meta-observed conversion.

While these results suggest that ATT had a dramatic effect on ad performance, it is important to note that these outcome variables are subject to measurement issues as a result of ATT. The observed decrease in conversions is a mixture of both real reductions in conversions and the degraded ability to link advertisements to conversions. This highlights

¹⁶In the remainder of the paper, we show results at a monthly frequency, but we prefer the weekly frequency for this plot to show how closely outcomes track the adoption of iOS 14.5.

¹⁷For better visual clarity, we also provide the adoption as a standalone plot in Figure A.1.

the challenge that both firms and Meta face after ATT, as accurately attributing conversions to advertisements plays a key role in measuring performance and learning effective targeting rules by enabling Meta to "close the loop." Another limitation of this event study is that it lacks a control group of unaffected companies, making it difficult to isolate ATT's impact. For us to determine whether there were real degradations in targeting caused by the introduction of ATT, we next exploit the fact that the ability to measure advertising clicks is not affected by ATT, unlike the ability to measure conversions.

Causal Effect on Conversion-Optimized Meta Advertising: We focus on quantifying the reduction in the effectiveness of campaigns that rely on off-platform data. To do so, we conduct a within-firm difference-in-differences analysis, comparing the relative performance of conversion-optimized to click-optimized advertising campaigns. This is the observational analog of the experimental comparison conducted in Wernerfelt et al. (2022). Click-optimized campaigns serve as a reasonable control group because (1) they optimize for the last point in the customer acquisition lifecycle that the platform can reliably measure after ATT, (2) they are the most popular campaign objective which can be reliably measured after the implementation of ATT,¹⁸ and (3) clicks are positively correlated with conversions.¹⁹

By focusing on a within-firm comparison, we isolate the effect of ATT on the affected form of advertising while controlling for differences across firms – for instance, their size or frequency of conversions – that are orthogonal to the treatment effect of interest, as well as possible adjustments to the targeting algorithm by Meta over time. We consider the following specification for firm i, advertising campaign objective j, and month t:

$$Y_{ijt} = \sum_{t} \beta_t \Big(\text{Month}_t \times T_j \Big) + \alpha_{ij} + \kappa_t + \epsilon_{ijt}$$
⁽²⁾

where T_j is an indicator for whether the campaign j is a conversion-optimized campaign, α_{ij}

¹⁸Table OA1 provides a breakdown of the market shares of different campaign objectives before ATT.

¹⁹Table A.2 shows that before ATT clicks and conversions were correlated with each other for both campaign objectives, with the relationship being stronger for conversion-optimized campaigns.

denotes firm-campaign fixed effects, and κ_t denotes month fixed effects.



Figure 3: Time-Varying Treatment Effects for Click-Through Rate

NOTES: This figure shows the relative performance of the click-through rates of conversionoptimized campaigns (which were affected by ATT) compared to click-optimized campaigns (which were not) over time, using specification (2). It uses data from a balanced panel of firms that used both types of campaigns pre-ATT in the advertising dataset. Standard errors are clustered at the firm level. Associated aggregate post-ATT estimates are in Table A.1.

One possible threat to this identification strategy is if firms reallocate their advertising spending between conversion-optimized campaigns and campaigns optimized for on-platform objectives. In Appendix Section C.1 we investigate this possibility and find minimal substitution between campaign types. Specifically, off-platform campaigns continued to dominate advertiser spending on Meta, accounting for 95.7% of total spend before ATT and 95.0% afterward (Table OA2). Our firm-level analysis shows some substitution on the extensive margin (i.e., some firms beginning to use on-platform objectives that were not previously used) but no significant adjustment on the intensive margin (i.e., spending levels across campaign types). Therefore, to ensure the validity of our approach, we focus our primary analysis on the subset of firms that spent on both click-optimized and conversion-optimized campaign objectives before ATT. The stability in campaign mix among these firms provides evidence against Stable Unit Treatment Value Assumption (SUTVA) violations stemming

from cross-objective substitution and supports the validity of our within-firm comparison.²⁰

As such, we estimate specification (2) using the click-through rate for firm *i* and campaign objective j for each month t as Y_{ijt} and on a balanced panel of firms that utilized both click-optimized and conversion-optimized campaigns pre-ATT.²¹ Figure 3 shows identical performance between campaign types before ATT, followed by a sharp decline in clickthrough rates for conversion-optimized campaigns after ATT's introduction. The reduction in click-through rates is 0.004 for conversion-optimized campaigns, representing a 36.6%decrease from the baseline rate of 0.011 (95% CI: 18.2% to 54.5%).²² While we cannot directly characterize the impact of ATT on conversions due to the inability of Meta to reliably measure this after ATT, column (1) of Table A.2 shows that a 1% increase in pre-ATT clicks was associated with a 0.61% increase in pre-ATT conversions, suggesting that the causal reduction in clicks will likely be associated with a decline in conversions. Indeed, the decline in click-through rates may understate the true decline in conversions because ATT could also degrade the quality of clicks – ATT not only reduces the ability to measure conversions but also impacts targeting quality, which could affect the alignment between the clicked ad and consumer purchase intent. This evidence implies that ATT significantly degraded the effectiveness of conversion-optimized advertising on Meta.

3.1 Budget Reallocation

Given that ATT negatively impacted the effectiveness of Meta advertising, it is natural to ask whether and how firms adapted by reallocating their advertising spend, as this could influence the overall effect on revenue. To explore this, we focus on Google, the other

 $^{^{20}}$ This stability in campaign mix suggests that any overall changes in Meta advertising spending represent level shifts across all campaign types rather than a selective reallocation, further supporting the validity of our comparison.

²¹This selection criterion yields a subset of 44.9% of firms from Figure 2, with Table A.4 showing similar event study impacts on conversion-optimized ads.

 $^{^{22}}$ An additional possible concern is that this analysis conditions on advertisers that continued to purchase advertising on Meta. We estimate this same specification instead conditioning on the set of advertisers that have positive advertising spending on Meta only in the pre-ATT period, allowing for possible exit off Meta. We report the results in Figure A.2 showing that the estimates remain quantitatively similar.

prominent online advertising platform observed in our data. Although our measures of advertising performance changes on Google are not as precise as those for Meta, we show in Online Appendix C.2 that conversions across various Google services do not exhibit the same abrupt decline post-ATT as observed in Figure 2. This suggests that firms could potentially mitigate the impact of ATT by shifting their advertising spending to Google.

Figure 4: Event Study of Online Advertising Spending



(a) Event Study for Meta Advertising Spending Share



(b) Event Study for Meta Advertising Spending

(c) Event Study for Google Advertising Spending

NOTES: Panel (a) represents event study estimates for Meta online advertising spending share, defined as spending on Meta advertising as a proportion of advertising spending on Meta, Google, and TikTok, using specification (1). Table A.3 presents the associated aggregate post-ATT estimates. Panels (b) and (c) consider the dependent variable as the log of online advertising spending for Meta and Google, respectively. Results use a balanced panel of firms with non-zero online advertising spending in the advertising dataset. Standard errors are clustered at the firm level.

Measuring the equilibrium effects of ATT on the advertising market is challenging as ATT

induces an exogenous reduction in quality for targeted advertising and thus simultaneously impacts quality, quantity, and prices. As our primary goal is to understand the downstream impact on revenue, we focus primarily on reduced-form reallocations to the online advertising platforms of interest since this may impact downstream outcomes. Nonetheless, we specify a micro-founded model of advertising allocations in Online Appendix F that makes a clear prediction – relative demand for Meta should decrease compared to Google – and also highlights the theoretical ambiguity of other key market outcomes in equilibrium.

To empirically validate this, we compute each firm's online advertising spending share on Meta, calculated as their advertising spending on Meta in a month divided by their total advertising spending on Meta, Google, and TikTok during the same month. We then estimate specification (1) and present the estimates in Figure 4. They show little change in market share before the onset of ATT and a gradual decrease in the share of Meta after ATT. Figures 4b and 4c present the event study estimates for the log of advertising spending on Google and Meta, respectively, which show that this result arises from a mixture of continued increase in Google advertising spending and a drop off in Meta advertising spending.²³ The mean market share for Meta ads was 0.75 in the baseline period. The average decline across the post-treatment period was 0.014 (SE: 0.003) or 1.4 percentage points (95% CI: 0.8 to 2.0 percentage points), while by the end of our sample period this effect grew to approximately 3.3 percentage points (4.4% reduction), as shown in Figure 4.

While other market factors could influence platform-specific advertising spending during this period, several aspects of our analysis mitigate these concerns and suggest that this was a result of ATT. First, the timing of the divergence aligns precisely with ATT implementation. Second, the pre-ATT parallel trends in spending across platforms suggest comparable growth trajectories absent intervention. Third, in Online Appendix C.2, we conduct an across-firm difference-in-differences analysis to show that this reallocation was more pronounced for firms with higher pre-ATT Meta dependence.

²³In Online Appendix C.2 we present additional details and analyses for this, as well as show that a similar pattern holds for two relevant quantity variables: clicks and impressions.

These results indicate a meaningful reallocation of advertising spending, suggesting that to characterize ATT's full impact on these firms, we need to understand the impact on total revenue. We turn to this in the next section.

4 Impact on Firm Revenues

This section contains our main results, in which we estimate the impact of ATT on firm revenues using the revenue dataset. Our primary empirical strategy employs an across-firm difference-in-differences (DiD) design to compare pre- and post-ATT revenue for firms differing in their vulnerability to ATT. We consider two complementary exposure metrics: the firm's pre-ATT revenue share attributable to Meta traffic or to iOS devices. While the former follows naturally from Section 3 (degraded conversion-optimization on Meta), the latter captures vulnerability across all channels on devices directly affected by ATT. Because iOS share is based on device-level tags, it is measured more accurately than Meta attribution, which relies on 30-minute last-touch windows and systematically under-counts Meta's true contribution (see Section 2). These two measures are only weakly correlated,²⁴ allowing us to triangulate ATT's impact from two distinct angles.

To measure these forms of dependence, we calculate the average share of revenue coming from Facebook/Instagram sessions or iOS devices, respectively, over the one-year period before ATT's introduction (April 2020 to April 2021) for each firm. Recall from Table 1 that this share is, for the average firm in our sample, 0.04 for Meta and 0.25 for iOS.

There are two key challenges in measuring the causal effect on revenue: most firms are at least partially treated under either measure of exposure and reported exposure potentially has measurement error. As mentioned above, we expect reported Meta dependence to be systematically lower than true dependence on Meta across firms, while reported iOS exposure

²⁴While iOS dependence and Meta dependence are correlated, the correlation is relatively weak – the ϕ coefficient is 0.17 – as there is substantial variation in which firms are labeled as treated under the two definitions. In all, 29% of firms are considered treated under both definitions, 42% are considered treated under only one, and 29% are considered treated under neither.

is likely to have significantly less systematic bias. Unlike classical measurement error with random variation that typically attenuates treatment effect estimates, the systematic underattribution of Meta exposure is a directional bias.





(a) Meta Treatment

Notes: The estimates present the time-varying treatment effects for log(total revenue) using specification (3), with data from the revenue dataset. The treatment indicator is a dummy variable equal to 1 if the firm-level pre-ATT share of revenue from Meta traffic is above the median, and 0 otherwise. Standard errors are clustered at the firm-level.

Rather than relying on a single estimation approach, we employ two estimation methods,

with different strengths and weaknesses, using the two treatment variables to construct a range of plausible estimates for ATT's impact. First, we use a median split of exposure to classify firms into 'high exposure' (treatment) and 'low exposure' (control) groups. The median split approach provides a straightforward interpretation and shows consistent patterns across treatment variables. Second, we consider an alternative specification following de Chaisemartin et al., 2024, the heterogeneous adoption design (HAD) estimator, which measures the causal effect of an additional unit of measured exposure. The HAD estimator offers more generalizable, policy-relevant interpretations. However, it produces estimates with substantially wider confidence intervals, making precise interpretation more challeng-ing.²⁵

Relative Dependence Treatment We first consider the specification that relies on measuring exposure via relative dependence by classifying treated and control units based on median exposure levels. Using Meta (iOS) attributable revenue as the measure of exposure this results in a treatment group where 8.17% (36.51%) of their pre-ATT revenue is attributed to Meta (iOS) traffic, compared to 0.42% (11.84%) for the control group. To assess this, we estimate the following specification for results in this section:

$$Y_{it} = \sum_{t} \beta_t \left(\text{Month}_t \times T_i \right) + \alpha_i + \kappa_t + \epsilon_{it}, \qquad (3)$$

where T_i indicates whether they are more vulnerable to ATT, α_i denotes firm fixed effects, and κ_t denotes month fixed effects. We also run a robustness check in which we include category-month fixed effects. As before, we cluster our standard errors at the firm level.

Results for the primary specification are shown in Figure 5 with Figures 5a and 5b displaying the time-varying treatment effects for Meta and iOS treatment assignment, respectively.

 $^{^{25}}$ In Monte Carlo simulations designed to mimic the empirical setting of our study, we find that the width of the empirical 95% confidence interval of the HAD estimator is approximately three times larger than that of the median-split estimator under the same data generation process. Simulation results are available upon request.

Under both measures, estimated monthly treatment effects for revenue remain statistically insignificant during the pre-ATT period, confirming parallel trends. Following ATT implementation, we observe a gradual decline in revenue for treated firms, with point estimates becoming statistically significant approximately 4 months after ATT's introduction. This is again consistent with the gradual timing of adoption of iOS 14.5 among consumers documented in Figure 2. These results suggest that the rollout of ATT substantially lowered revenue of the e-commerce firms most exposed to it. Notably, despite the relatively weak correlation between these two treatment measures, we observe remarkably consistent effect estimates over time across measures. This is supportive of the notion that the estimated revenue effects reflect the true impact of ATT and are not artifacts of a particular exposure metric.

Table 2 provides the aggregated coefficient estimates for the Meta and iOS dependence measures. The coefficient estimates in columns (1) and (5) suggest a decrease in revenue of 37.1% (95% CI: 12.4% to 55.1%) for more Meta-dependent firms relative to less Metadependent firms and of 40.6% (95% CI: 18.0% to 57.1%) for more iOS-dependent firms relative to less iOS-dependent firms.²⁶ We focus our discussion on the Meta dependence measure results for the rest of this discussion as the results are consistent across both specifications. Column (2) reveals that the negative revenue effect strengthens over time, with a small and statistically insignificant effect in the initial three months post-implementation (-0.138, or -12.9% in percentage terms), followed by a larger and statistically significant effect in the subsequent months (-0.499, or -39.3% in percentage terms). This pattern aligns with the gradual adoption of iOS 14.5, as shown in Figure 2, where the ATT opt-out rate increased steadily and then jumps in June 2021, when Apple began actively prompting users to update their devices through notifications and automatic update settings. Columns (3) and (4) show that this effect is driven by small firms, which are defined as those with below-median pre-ATT average monthly revenue.²⁷ Additional analyses in Table A.5 in the

²⁶Marginal effects are computed by $\exp(\beta) - 1$.

²⁷We plot time-varying treatment effects separately for small and large firms in Figures A.4a and A.4b,

Online Appendix show that including category-by-month fixed effects yields similar results, with a 32.7% decline for Meta-dependent firms and a 29.7% decline for iOS-dependent firms. We include this as a specification check while noting that category and treatment may be correlated, as some categories are inherently more reliant on targeted digital advertising than others. The same table also shows that the number of transactions declined by approximately 21% under both treatment definitions.

	$Dependent \ variable: \ \log(Revenue)$							
	Meta Treatment			iOS Treatment				
	(1) All firms	(2) All firms	(3) Small firms	(4) Large firms	(5) All firms	(6) All firms	(7) Small firms	(8) Large firms
$\overline{\text{After}_t \times \text{Treated}}$	-0.463^{***} (0.165)		-1.132^{***} (0.302)	0.158 (0.121)	-0.522^{***} (0.165)		-0.999*** (0.290)	0.079 (0.116)
$0-3 \text{ months} \times \text{Treated}$		-0.138 (0.151)				$0.040 \\ (0.103)$		
$4+$ months \times Treated		-0.499^{***} (0.171)				-0.483^{***} (0.159)		
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	24868	24868	12468	12400	24868	24868	12468	12400
\mathbb{R}^2	0.73	0.73	0.58	0.67	0.73	0.73	0.58	0.67
Marginal effects $(\%)$	-37.06%	-12.89% -39.29%	-67.76%	17.12%	-40.67%	4.08% -38.31\%	-63.18%	8.22%
Treatment share treated $(\%)$	8.17%	8.17%	10.22%	5.96%	36.51%	36.51%	35.10%	37.73%
Treatment share not treated $(\%)$	0.42%	0.42%	0.34%	0.49%	11.84%	11.84%	11.35%	12.42%

Table 2: Primary Revenue Estimates

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

NOTES: The treatment indicator is a dummy variable equal to 1 if the firm-level share of revenue from the respective traffic source (Meta or iOS) is above the median, 0 otherwise. All columns use log(revenue) as the dependent variable. Columns present estimated average treatment effect coefficients using variants of specification (3), replacing monthly dynamic treatment effects with three specifications: (i) a post-treatment indicator (After_t × Treated), (ii) an indicator for the first 3 months after treatment (0 - 3 months × Treated), and (iii) an indicator for 4+ months after treatment (4+ months × Treated), with data from the revenue dataset. Marginal effects are computed by $\exp(\beta) - 1$. Standard errors are clustered at the firm level.

A limitation of this analysis is the absence of a true control group, as virtually all firms have some exposure to ATT through iOS users or Meta advertising. To address this concern, we supplement our main analysis with two approaches. First, we implement quartile-based comparisons to examine how effects vary across different levels of exposure intensity. These respectively, and for transactions in Figure A.3. results are shown in Table A.6. For the iOS treatment variable, we find that increasing levels of exposure are associated with increasing effect sizes and that the magnitudes of effects are consistent with the median split implementation. For the Meta treatment, we find large effects for the second through fourth quartile of exposure. The quartile coefficients are somewhat imprecisely estimated but suggest large effects are present beginning with the second quartile of exposure.

Continuous treatment Next, we consider a complementary approach that measures the causal effect of an additional unit of dependence on revenue. To do so, we follow de Chaise-martin et al. (2024), who developed an estimator for settings where all units are treated but with different intensities (which we denote the HAD estimator). This approach uses "quasi-stayers" (units with minimal treatment changes) and provides estimates of treatment effects that are robust to heterogeneous adoption. Additional discussion of this method is provided in Appendix B. Results are summarized in Table 3 for both the iOS and Meta treatment. The coefficients suggest that an additional 1 percentage point increase in revenue attributed to Meta advertising is associated with a roughly 1.05% decrease in revenue attributed to iOS users is associated with a roughly .75% decrease in revenue.

	iOS Treatment	Meta Treatment
$\overline{\text{After} \times \text{Treatment}}$	-0.751^{***} (0.174)	-1.048^{***}

Table 3: HAD Estimates for iOS and Meta Treatments

NOTES: Results show summarized estimates from the de Chaisemartin et al. (2024) HAD estimator that aggregate across all post-periods.

To provide a comparison to our median split difference-in-differences estimates, we can compare the magnitudes implied by these coefficients applied to the difference in treatment levels between the above versus below median firms. The difference in Meta exposure between treatment and control groups is approximately 7.75 percentage points (8.17% vs 0.42%), which would imply an effect of 8.1% using these estimates (95% CI: -11.2% to -5.1%). Using the same calculations under the iOS treatment variable implies an 18.5% revenue reduction (95% CI: -26.9% to -10.1%). These estimates are smaller than the treatment effects under the median split estimator of 37.1% and 40.6% under the Meta and iOS treatment variables, respectively, although the gap is noticeably smaller and statistically insignificant (p > .10) under the iOS treatment variable.

Several factors could explain these results: (1) last-touch attribution with 30-minute windows understates firms' true reliance on Meta advertising, meaning actual exposure differences between treatment and control groups are likely larger than measured; (2) the HAD estimator exhibits lower statistical precision in our setting, introducing variability into its point estimates for a given realization of the data; (3) measurement error may affect the HAD estimator differently than the differences-in-differences estimator.²⁸ We note that the iOS-dependence specification, where we expect less measurement error in the measure of dependence, has a more similar treatment effect magnitude to the median split estimator, with differences in treatment effect estimates that are statistically insignificant from one another. This is consistent with the notion that measurement error in the treatment variable may affect our estimation approaches differently, though the exact mechanisms and magnitudes of these effects are difficult to quantify precisely. We also note that all approaches yield consistent, significantly negative inferred treatment effects, which provides robust evidence that ATT had significantly negative impacts on e-commerce firms' revenues, regardless of the specific methodological approach that is employed.

Discussion These findings indicate that privacy protection measures have substantially harmed some e-commerce firms, with effect estimates across all the considered specifications that are larger than what one might expect from the direct revenue share attributed to Meta

²⁸This influences the estimated coefficient since the systematic under-estimation of exposure means that some units with truly low treatment exposure are spuriously reported to have no exposure and are thus incorrectly classified as control units. This leads to a control group that is a mixture of units with zero exposure and those with minimal exposure, biasing the estimated treatment effects downwards.

advertising or iOS users. Several factors likely contribute to these large effects:

- Reliance on Meta in the revenue dataset is measured using last-touch attribution with a 30-minute attribution window, which, as noted in Gordon et al. (2023), is likely to significantly underestimate firms' true underlying reliance on Meta advertising.
- 2. These losses represent foregone growth rather than absolute revenue declines: Figure OA4 in Online Appendix D.1 shows that more Meta-reliant firms experience slower growth compared to less reliant firms, not actual revenue contraction. Consistent with this, auxiliary analyses of a secondary revenue dataset in Online Appendix D.2 indicates that reduction in new customer orders is the driving force of revenue reductions.
- 3. Losing an important customer acquisition channel not only depresses short-term sales, it also depresses long-term sales through lower subsequent repeat purchases, less word of mouth, and so on.²⁹
- 4. Auxiliary analysis from Online Appendix D.2 suggests a small fraction of the revenue decline may be attributable to relatively more affected firms decreasing total advertising spending in response to ATT.
- 5. Cross-device impacts occur as ATT's effects spill over to all devices, not just iOS. This happens through algorithmic learning spillovers (targeting models becoming less effective when trained on incomplete data) and identity linkage issues (inability to connect user activity across devices). These informational externalities are consistent with recent empirical work by Aridor et al. (2023) and Lin and Misra (2022) as well as a growing theoretical literature (Acemoglu et al., 2022; Bergemann et al., 2022; Choi et al., 2019; Miklós-Thal et al., 2024).

While we identify these specific mechanisms, we acknowledge that additional factors beyond ATT also influence revenue outcomes, contributing to uncertainty in our estimates. Col-

 $^{^{29}}$ We also find negative, but imprecisely estimated, point estimates on orders from repeat customers in the analysis of an auxiliary revenue dataset in Online Appendix D.2.

lectively, these findings demonstrate how privacy protection measures can have far-reaching economic consequences beyond their direct implementation targets.

5 Conclusion

As companies and policymakers consider extending or implementing new privacy policies limiting firms' ability to target consumers online, it is important that they be fully informed about the economic costs to firms that may result from these regulations. In this paper, we show that ATT significantly degraded the performance of Meta advertising and, subsequently, that firms more dependent on Meta experienced a substantial relative reduction in revenue, which was primarily borne by small firms.

This paper has several policy and managerial takeaways. Our estimates suggest large economic costs of opt-in privacy protection measures. While there are positive consumer welfare gains from the added privacy protections, the magnitude of the losses threatens the viability of firms, such as direct-to-consumer firms, that rely on targeted social media advertising as their primary source of customer acquisition. As such, there could be a countervailing force on consumer welfare if the revenue losses are large enough to induce substantial exit and to deter entry of these firms into product markets. These findings highlight the importance of developing balanced approaches to privacy protection that protects consumer rights while also supporting the competitiveness of small businesses in the digital economy.

References

- Acemoglu, D., Makhdoumi, A., Malekian, A., & Ozdaglar, A. (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4), 218–256.
- Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. Journal of Economic Literature, 54(2), 442–492.

- Aridor, G., Che, Y.-K., & Salz, T. (2023). The effect of privacy regulation on the data industry: Empirical evidence from GDPR. RAND Journal of Economics, 54(4), 695– 730.
- Baviskar, S., Chowdhury, I., Deisenroth, D., Li, B., & Sokol, D. D. (2024). Att vs. personalized ads: User's data sharing choices under apple's divergent consent strategies.
- Bergemann, D., Bonatti, A., & Gan, T. (2022). The economics of social data. The RAND Journal of Economics, 53(2), 263–296.
- Bian, B., Ma, X., & Tang, H. (2021). The supply and demand for data privacy: Evidence from mobile apps. Available at SSRN 3987541.
- Borkovsky, R., Goldfarb, A., Haviv, A., & Moorthy, S. (2017). Measuring and understanding brand value in a dynamic model of brand management. *Marketing Science*, 36(4), 471–499.
- Cecere, G., & Lemaire, S. (2023). Have I seen you before? Measuring the value of tracking for digital advertising. Available at SSRN 4659963.
- Chen, B. (2021). The Battle for Digital Privacy Is Reshaping the Internet. New York Times. https://www.nytimes.com/2021/09/16/technology/digital-privacy.html
- Cheyre, C., Leyden, B. T., Baviskar, S., & Acquisti, A. (2023). The impact of Apple's App Tracking Transparency framework on the app ecosystem. *Available at SSRN 4453463*.
- Choi, J. P., Jeon, D.-S., & Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173, 113–124.
- CMA. (2022). Mobile ecosystems market study. https://www.gov.uk/cma-cases/mobileecosystems-market-study
- de Chaisemartin, C., D'Haultfœuille, D. C. X., & Knau, F. (2024). Two-way fixed effects and differences-in-differences estimators in heterogeneous adoption designs. arXiv preprint arXiv:2405.04465.

- Deisenroth, D., Manjeer, U., Sohail, Z., Tadelis, S., & Wernerfelt, N. (2024). Digital advertising and market structure: Implications for privacy regulation. National Bureau of Economic Research.
- Dubé, J.-P., Bergemann, D., Demirer, M., Goldfarb, A., Johnson, G., Lambrecht, A., Lin, T., Tuchman, A., Tucker, C. E., & Lynch, J. G. (2024). The intended and unintended consequences of privacy regulation for consumer marketing: A marketing science institute report. Available at SSRN 4847653.
- Ganglmair, B., Krämer, J., & Gambato, J. (2023). Regulatory compliance with limited enforceability: Evidence from privacy policies. *Available at SSRN*.
- Goldberg, S. G., Johnson, G. A., & Shriver, S. K. (2024). Regulating privacy online: An economic evaluation of the GDPR. American Economic Journal: Economic Policy, 16(1), 325–358.
- Goldfarb, A., & Que, V. F. (2023). The economics of digital privacy. Annual Review of Economics, 15, 267–286.
- Goldfarb, A., & Tucker, C. E. (2011). Privacy regulation and online advertising. Management Science, 57(1), 57–71.
- Gordon, B. R., Moakler, R., & Zettelmeyer, F. (2023). Predictive incrementality by experimentation (PIE) for ad measurement. arXiv preprint arXiv:2304.06828.
- Johnson, G. (2023). Economic Research on Privacy Regulation: Lessons From the GDPR and Beyond. *The Economics of Privacy*. http://dx.doi.org/10.2139/ssrn.4290849
- Johnson, G., Runge, J., & Seufert, E. (2022). Privacy-centric digital advertising: Implications for research. Customer Needs and Solutions, 9, 1–6.
- Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1), 33– 51.
- Kesler, R. (2022). The impact of Apple's App Tracking Transparency on app monetization. Available at SSRN 4090786.

- Kobayashi, S., Johnson, G., & Gu, Z. (2024). Privacy-enhanced versus traditional retargeting:Ad effectiveness in an industry-wide field experiment. Available at SSRN 4972368.
- Kollnig, K., Shuba, A., Van Kleek, M., Binns, R., & Shadbolt, N. (2022). Goodbye tracking? impact of iOS App Tracking Transparency and privacy labels. *Proceedings of the 2022* ACM Conference on Fairness, Accountability, and Transparency, 508–520.
- Korganbekova, M., & Zuber, C. (2023). Balancing user privacy and personalization.
- Kraft, L., Skiera, B., & Koschella, T. (2023). Economic impact of opt-in versus opt-out requirements for personal data usage: The case of apple's app tracking transparency (att). Available at SSRN 4598472.
- Laziuk, E. (2021). iOS 14.5 Opt-in Rate Daily Updates Since Launch. *Flurry Analytics*. https://www.flurry.com/blog/ios-14-5-opt-in-rate-att-restricted-app-tracking-transparency-worldwide-us-daily-latest-update/
- Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (Forthcoming). The Impact of the GDPR on Content Providers: A Longitudinal Analysis. *Management Science*.
- Li, D., & Tsai, H.-T. (2022). Mobile apps and targeted advertising: Competitive effects of data exchange. Available at SSRN 4088166.
- Lin, T., & Misra, S. (2022). Frontiers: The identity fragmentation bias. *Marketing Science*, 41(3), 433–440.
- Miklós-Thal, J., Goldfarb, A., Haviv, A., & Tucker, C. (2024). Frontiers: Digital hermits. Marketing Science, 43(4), 697–708.
- Miller, K. M., & Skiera, B. (2023). Economic consequences of online tracking restrictions: Evidence from cookies. International Journal of Research in Marketing, 41(2), 241– 264.
- Runge, J., & Seufert, E. (2021). Apple is changing how digital ads work. Are advertisers prepared? *Harvard Business Review, digital article*.

- Sokol, D. D., & Zhu, F. (2021). Harming competition and consumers under the guise of protecting privacy: An analysis of Apple's iOS 14 policy updates. *Cornell L. Rev. Online*, 107, 94.
- Todri, V. (2022). Frontiers: The Impact of Ad-Blockers on Online Consumer Behavior. Marketing Science, 41(1), 7–18. https://doi.org/10.1287/mksc.2021.1309
- Utz, C., Degeling, M., Fahl, S., Schaub, F., & Holz, T. (2019). (Un)informed consent: Studying GDPR consent notices in the field. Proceedings of the 2019 ACM Sigsac Conference on Computer and Communications Security, 973–990.
- Wernerfelt, N., Tuchman, A., Shapiro, B., & Moakler, R. (2022). Estimating the value of offsite data to advertisers on Meta. http://dx.doi.org/10.2139/ssrn.4176208
- Yan, S., Miller, K., & Skiera, B. (2022). How does the adoption of ad blockers affect news consumption? Journal of Marketing Research, 002224372210761. https://doi.org/10. 1177/00222437221076160

Appendix A Omitted Tables and Figures



NOTES: Figure represents the estimated percentage of iOS devices that updated to iOS 14.5 over time. The first vertical dotted line represents April 25, 2021, when Apple first introduced iOS 14.5. The second vertical dotted line represents June 1 2021, when Apple began encouraging iOS users to update their operating systems. Source: Gupta Media, https://lookerstudio.google.com/u/0/reporting/ $3d5dda40-37ea-4b9f-bd91-bb8df8e12620/page/aDUJC?s=kTs6iab_AhQ$

	Dependent variable:		
	(1)	(2)	
	Click-Through Rate		
$After_t \times Treated$	-0.004^{***} (0.001)	-0.004^{***} (0.001)	
Month FE Firm-Campaign FE	Yes No	Yes Yes	
Observations R ²	$18,427 \\ 0.666$	$18,427 \\ 0.483$	

Table A.1: Difference-in-Differences Estimates for CTR

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

NOTES: This table shows the relative performance of the click-through rates of conversion-optimized campaigns (which were affected by ATT) compared to click-optimized campaigns (which were not) over time. This is the static analog to specification (2), replacing time-varying treatment effects with a single post-treatment indicator (After_t × Treated) to estimate an average treatment effect over the post-treatment period. It uses data from a balanced panel of firms that used both types of campaigns pre-ATT in the advertising dataset. Standard errors are clustered at the firm level.

Figure A.2: Time-Varying Treatment Effects for Click-Through Rate (Robustness)



NOTES: This figure shows the relative performance of the click-through rates of conversionoptimized campaigns (which were affected by ATT) compared to click-optimized campaigns (which were not) over time, using specification (2). It uses data from firms that used both types of campaigns pre-ATT in the advertising dataset. It only considers a balanced panel of firms with positive Meta advertising spending in the pre-ATT period. Standard errors are clustered at the firm level.

	Dependent variable:	
	(1)	(2)
	$\log(1 + \text{Conversions}_i)$	
$\log(1 + \text{Clicks}_{ijt})$	0.252***	0.195***
	(0.017)	(0.013)
$\log(1 + \text{Clicks}_{ijt}) \times \mathbb{1}(j \text{ is Conversion-Optimized Campaign})$	0.649***	0.413***
	(0.031)	(0.045)
Firm-Campaign FE	No	Yes
Week FE	No	Yes
Observations	7,840	7,840
R ²	0.762	0.951
	*p<0.1; **p	<0.05; ***p<0

Table A.2: Correlational Relationship between Clicks and Conversions on Meta

NOTES: All results use the advertising dataset, using a balanced panel of firms who spend on both clickoptimized and conversion-optimized campaigns on Meta. We only consider pre-ATT time period, due to the measurement issues associated with conversions after ATT. We estimate the following specification: $\log(1 + \operatorname{conversions}_{ijt}) = \beta \left(\log(1 + \operatorname{clicks}_{ijt}) \times \mathbb{1}(j \text{ is Conversion-Optimized Campaign}) \right) + \alpha_{ij} + \kappa_t + \epsilon_{ijt}$. Standard errors are clustered at the firm level.

	Dependent variable:			
	(1)	(2)	(3)	
	$\log(\text{Cost per conversion})$	$\log(\text{Conversions})$	Meta online spend share	
$After_t$	$\begin{array}{c} 0.549^{***} \\ (0.019) \end{array}$	-0.302^{***} (0.032)	-0.014^{***} (0.003)	
Firm FE	Yes	Yes	Yes	
Observations	61,091	61,091	31,746	
\mathbb{R}^2	0.837	0.855	0.931	
Marginal effects	73.15%	-26.07%	-	

Table A.3: Event Study Estimates for Meta Advertising Performance

p < 0.1; p < 0.05; p < 0.01

NOTES: All results use the advertising dataset. After_t is an indicator for whether the time is after ATT's introduction. The specification used is the static analog to specification (1), replacing timevarying treatment effects with a single post-treatment indicator (After_t). The dependent variables are log(cost per conversion), log(conversions), and online spend share for Meta ads, equal to ad spend on Meta as a proportion of ad spend on Meta, TikTok and Google. Columns (1) and (2) are estimated over the sample of firms that have a balanced panel in terms of Meta advertising spend, while column (3) is estimated over the sample of firms with a balanced panel of any online advertising spend. We note that the dependent variables in columns (1) and (2) have measurement issues after ATT. Marginal effects are computed by $\exp(\beta) - 1$. Standard errors are clustered at the firm level.

	Dependent variable:				
	(1)	(2)	(3)	(4)	
	log(Con	versions)	$\log(\text{Cost per conversion})$		
	Both campaigns	Only conversions	Both campaigns	Only conversions	
$After_t$	-0.361^{***}	-0.254^{***}	0.559***	0.539***	
	(0.043)	(0.048)	(0.026)	(0.027)	
Firm FE	Yes	Yes	Yes	Yes	
Observations	31,216	29,875	31,216	29,875	
\mathbb{R}^2	0.839	0.861	0.819	0.848	
Marginal effects	-30.30%	-22.43%	74.89%	71.42%	

Table A.4: Event	Study Estimates	for Meta Advertising	Performance ((Robustness)
------------------	-----------------	----------------------	---------------	--------------

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

NOTES: Results use the advertising dataset and we estimate the event study specification (1). The first two columns consider log(conversions) as the dependent variable, and the last two columns consider log(cost per conversion). The first and third columns are estimated over the sample of firms that use both click-optimized and conversion-optimized campaigns before ATT. The second and fourth columns are estimated over the sample of firms that only use conversion-optimized campaigns. Marginal effects are computed by $\exp(\beta) - 1$. The marginal effects associated with the point estimated presented in column (3) are 4.08% for the first 3 months and -38.31% for 4+ months after treatment. Standard errors are clustered at the firm level.

Table A.5: Additional Revenue Estimates

	Dependent variable:					
	log(Revenue) with (Category \times Month FE	$\log(\text{Transactions})$			
	(1) Meta Treatment	(2) iOS Treatment	(3) Meta Treatment	(4) iOS Treatment		
$\overline{\text{After}_t \times \text{Treated}}$	-0.396^{**} (0.180)	-0.352^{**} (0.167)	-0.241^{**} (0.100)	-0.237^{**} (0.100)		
Firm FE Month FE Category \times Month FE	Yes Yes Yes	Yes Yes Yes	Yes Yes No	Yes Yes No		
Observations	24868	24868	24868	24868		
\mathbb{R}^2	0.74	0.74	0.83	0.83		
Marginal effects (%)	-32.70%	-29.67%	-21.42%	-21.10%		
Treatment share treated $(\%)$	8.17%	36.51%	8.17%	36.51%		
Treatment share not treated (%)	0.42%	11.84%	0.42%	11.84%		

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

NOTES: The treatment indicator is a dummy variable equal to 1 if the firm-level share of revenue from the respective traffic source (Meta or iOS) is above the median, 0 otherwise. Columns (1)-(2) use log(revenue) as the dependent variable and include category-by-month fixed effects, where "category" refers to the firm category labels in the revenue data, such as "Lifestyle," "Home/Garden," and "Health." Columns (3)-(4) use log(transactions) as the dependent variable. Marginal effects are computed by $\exp(\beta) - 1$. Standard errors are clustered at the firm level.





NOTES: Results use the revenue dataset. The estimates present the time-varying treatment effects for log(Transactions) using specification (3). The treatment indicator is a dummy variable equal to 1 if the firm-level pre-ATT share of revenue from Meta traffic is above the median, 0 otherwise. Standard errors are clustered at the firm level.

Figure A.4: Time-Varying Estimates for Small and Large Firms (Meta Share Treatment)



NOTES: Results use the revenue dataset. The estimates present the time-varying treatment effects using specification (3) for log(total revenue) across firms whose pre-ATT revenue was below the median (Panel a) and above the median (Panel b). The treatment indicator is a dummy if the firm-level pre-ATT share of revenue from Meta traffic is above the median of the full sample, including both large and small firms, and 0 otherwise. Standard errors are clustered at the firm level.
	(1) iOS Treatment	(2) FB Treatment
$\overline{\text{After}_t \times 2\text{nd Treatment Quartile}}$	-0.426 (0.279)	-1.190^{***} (0.271)
$After_t \times 3rd$ Treatment Quartile	-0.591^{**} (0.273)	-1.218^{***} (0.273)
$After_t \times 4th$ Treatment Quartile	-0.632^{**} (0.253)	-0.890^{***} (0.299)
Firm FE	Yes	Yes
Month FE	Yes	Yes
$\frac{\text{Observations}}{\text{R}^2}$	$24868 \\ 0.73$	24868 0.73

Table A.6: Revenue Estimates Robustness: Quartile Treatment Split

NOTES: Results show estimated revenue effects using the treatment definitions based on measures of ATT exposure constructed from pre-ATT revenue shares attributed to iOS users and Meta advertising. The excluded category in both columns is the lowest quartile (least exposed to ATT). Standard errors are clustered at the firm level.

Appendix B Revenue Analysis: Continuous Treatment Effects Analysis Using Quasi-Stayers

We follow the approach proposed in de Chaisemartin et al. (2024), which is designed for empirical settings in which all units receive treatment, but with varying intensities. The standard TWFE estimator with continuous treatment can produce misleading results when treatment effects are heterogeneous and all units are treated with different doses. de Chaisemartin et al. (2024) suggest a procedure for designs with "quasi-stayers" (units with minimal treatment exposure) that involves testing for parallel trends in pre-treatment periods, and testing for whether treatment effects are mean-independent of treatment intensity using the Yatchew test. If neither test is rejected, their Heterogeneous Adoption Design (HAD) estimator and the standard TWFE estimator are expected to yield consistent results.

Panel A of Tables B.1 and B.2 present the results of the pre-ATT placebo tests for our Meta and iOS treatment variables, respectively, while the results of the Yatchew tests for mean independence are provided in the notes accompanying these tables. Under both treatment definitions, the pre-trend tests show no significant violations of the parallel trends assumption in any pre-treatment periods, suggesting that the parallel trends assumption is plausible in our setting.

The Yatchew tests for mean independence yield no significant p-values (all above 0.15 using iOS treatment definition and 0.22 using the Meta treatment definition) across all post-treatment periods. This provides no evidence against the assumption that treatment effects are mean-independent of treatment intensity.

Panel A	: Pre-ATT	Place	bo Tests
Period	Estimate	SE	95% CI
Month -12	1.09	1.25	[-2.96, 1.93]
Month -11	1.53	1.31	[-2.16, 5.29]
Month -10	1.42	1.54	[-1.22, 4.82]
Month -9	1.15	0.99	[-1.70, 2.74]
Month -8	1.15	0.99	[-2.57, 1.30]
Month -7	0.49	0.98	[-2.34, 1.50]
Month -6	0.40	0.90	[-1.65, 1.81]
Month -5	0.57	0.84	[-1.52, 1.78]
Month -4	0.56	0.81	[-1.58, 1.61]
Month -3	0.71	0.81	[-1.84, 1.36]
Month -2	0.70	0.81	[-1.34, 1.84]
Month -1	0.70	0.84	[-1.54, 1.76]
Panel B: I	Post-ATT	Effect	Estimates
Period	Estimate	SE	$95\%~{\rm CI}$
Month 1	-0.74	1.29	[-2.16, 2.89]
Month 2	-1.31	1.28	[-3.15, 1.87]
Month 3	-0.71	1.13	[-3.01, 1.43]
Month 4	-0.92	0.84	[-2.49, 0.80]
Month 5	-1.03	0.85	[-2.39, 0.94]
Month 6	-1.23	0.89	[-2.19, 1.29]
Month 7	-0.90	0.79	[-1.66, 1.46]
Month 8	-0.97	0.80	[-1.71, 1.42]
Month 9	-0.96	0.76	[-1.83, 1.15]
Month 10	-1.15	0.71	[-2.08, 0.71]
Month 11	-1.26	0.71	[-2.14, 0.64]
Month 12	-1.17	0.69	[-1.94, 0.77]
Month 13	-1.22	0.69	[-2.07, 0.62]
Month 14	-1.16	0.68	[-2.01, 0.67]
Month 15	-1.08	0.68	[-1.96, 0.69]
Month 16	-1.13	0.69	[-2.03, 0.67]
Month 17	-1.20	0.69	[-2.13, 0.59]
Month 18	-1.12	0.69	[-2.04, 0.67]
Month 19	-0.95	0.70	[-1.78, 0.95]
Month 20	-1.03	0.70	[-1.85, 0.90]

Table B.1: Robustness: HAD Tests andEstimates: Meta Treatment

Notes: This table presents HAD estimates for Meta treatment effects following de Chaisemartin et al. (2024). Panel A shows shows placebo tests for pre-treatment periods. Panel B monthly treatment effect estimates after ATT implementation. All p-values from the Heteroskedasticity-robust Yatchew Test are greater than 0.22, indicating no significant violations of model assumptions. Sample sizes range from 621-626 observations. 95% confidence intervals are shown in brackets.

Panel A	: Pre-ATT	Place	bo Tests
Period	Estimate	SE	95% CI
Month -12	1.35	1.41	[-0.91, 4.63]
Month -11	1.21	1.09	[-0.24, 4.03]
Month -10	1.55	0.87	[-0.09, 3.32]
Month -9	1.52	0.83	[-0.06, 3.30]
Month -8	0.68	0.70	[-0.53, 2.21]
Month -7	0.56	0.68	[-0.51, 2.15]
Month -6	0.79	0.66	[-0.49, 2.09]
Month -5	0.69	0.64	[-0.47, 2.04]
Month -4	0.75	0.64	[-0.35, 2.17]
Month -3	0.77	0.65	[-0.23, 2.31]
Month -2	0.83	0.60	[-0.25, 2.12]
Month -1	0.72	0.60	[-0.38, 1.97]
Panel B: I	Post-ATT	Effect	Estimates
Period	Estimate	SE	95% CI
Month 1	-0.65	0.77	[-2.33, 0.71]
Month 2	-0.88	0.65	[-2.19, 0.35]
Month 3	-0.66	0.60	[-1.92, 0.42]
Month 4	-0.73	0.59	[-1.93, 0.39]
Month 5	-0.79	0.58	[-1.80, 0.49]
Month 6	-0.74	0.60	[-2.05, 0.32]
Month 7	-0.68	0.61	[-1.76, 0.60]
Month 8	-0.65	0.60	[-1.67, 0.67]
Month 9	-0.72	0.58	[-2.01, 0.27]
Month 10	-0.73	0.61	[-2.23, 0.16]
Month 11	-0.75	0.60	[-2.18, 0.17]
Month 12	-0.78	0.59	[-2.05, 0.26]
Month 13	-0.76	0.58	[-2.02, 0.28]
Month 14	-0.73	0.60	[-2.10, 0.26]
Month 15	-0.68	0.60	[-2.13, 0.32]
Month 16	-0.82	0.57	[-1.88, 0.40]
Month 17	-0.86	0.57	[-1.88, 0.36]
Month 18	-0.75	0.61	[-2.13, 0.24]
Month 19	-0.80	0.59	[-1.98, 0.32]
Month 20	-0.84	0.60	[-1.98, 0.36]

Table B.2: Robustness: HAD Tests and Estimates: iOS Treatment

Notes: This table presents HAD estimates for iOS treatment effects following de Chaisemartin et al. (2024). Panel A shows placebo tests for pre-treatment periods. Panel B shows monthly treatment effect estimates after ATT implementation. All p-values from the Heteroskedasticity-robust Yatchew Test are greater than 0.15, indicating no significant violations of model assumptions. Sample sizes range from 661-666 observations. 95% confidence intervals are shown in brackets.

Figure B.1: Dynamic Estimates from de Chaisemartin et al. (2024)



NOTES: Results use the revenue dataset. The estimates present the treatment effects as described in de Chaisemartin et al. (2024) The left panel shows results using the pre-treatment Meta revenue share as the treatment variable and the right panel shows results using the pre-treatment iOS share as the treatment variable.

Since both tests suggest that the key assumptions are satisfied, our standard TWFE approach in the main analysis should not suffer from the biases identified in de Chaisemartin et al. (2024). For completeness, we also provide the results from the HAD estimator in Panel B of Figures Tables B.1 and B.2 for the Meta and iOS treatment variables, respectively. These estimates are then plotted in Figure B.1, which shows evidence of consistently negative effects across post-treatment periods for both treatment variables. While individual monthly estimates from the HAD approach exhibit some imprecision, the aggregated post-treatment effect in Table B.3 shows a highly significant negative effect (p < .01), with coefficients of -0.751 for iOS treatment and -1.048 for Meta treatment.

Table B.3: HAD Estimates for iOS and Meta Treatments

	iOS Treatment	Meta Treatment
$\overline{\text{After} \times \text{Treatment}}$	-0.751^{***} (0.174)	-1.048^{***} (0.202)

NOTES: Results show summarized estimates from the de Chaisemartin et al. (2024) HAD estimator that aggregate across all post-periods.

Online Appendix: For Online Publication Only

Appendix C Additional Analyses for Advertising

C.1 Meta Campaign Objective Substitution

We study reallocation within the Meta advertising ecosystem as a result of ATT. One way that firms might adapt to the reduction in effectiveness of off-platform conversion-optimized campaigns is to reallocate their spending within Meta to campaign objectives that do not rely on off-platform measurements. This shift could help maintain the effectiveness of firms' targeting efforts, as on-platform actions can serve as good indicators for off-platform actions, and they still provide direct feedback for optimization. As such, the main goal of this section is to understand the extent and magnitude of such reallocation.

First, we discuss the various targeting objectives that firms can optimize for on the Meta advertising platform. There are a large number of objectives and, for our purposes, we categorize the different objectives into three groups: off-platform conversions, on-platform actions, and 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.¹ On-platform actions consist of link clicks, store visits, page likes, leads outcome, traffic outcomes, engagement outcomes, and post engagement. On-platform reach consists of video views, brand awareness, reach, and awareness outcomes.²

We present summary statistics across the individual campaign objectives, and summarize the mapping to campaign objective categories, over the pre-ATT period, in Table OA1. There are two main observations to note. First, the conversion-optimized campaigns make up the vast majority of total spending. Second, optimizing for link clicks is the most popular

¹In our main analyses, we define conversion-optimized campaigns as those with conversions, sales outcomes, and product catalog sales objectives. While app install campaigns also use off-platform data, their attribution path differs from product sales (Li & Tsai, 2022). Since our focus is on the impact on firm revenue, we exclude these campaigns, though their inclusion does not significantly alter the results.

²This includes all campaign objectives except for messages and event responses because ambiguity arises when attempting to categorize them.

on-platform campaign objective.³ Table OA2 provides summary statistics for each of the campaign objective groups, and shows that 95% to 96% of Meta advertising spend is on conversion-optimized campaigns, both before and after ATT.

Campaign objective	Categorized objective	Total spend share
Conversions	Off-platform conversions	84.14%
Product catalog sales	Off-platform conversions	9.42%
App installs	Off-platform conversions	2.09%
Link clicks	On-platform actions	1.90%
Reach	On-platform reach	0.91%
Brand awareness	On-platform reach	0.79%

On-platform reach

On-platform actions

On-platform actions

On-platform actions

Video views

Store visits

Page likes

Post engagement

0.55%

0.15%

 $0.02\% \\ 0.02\%$

Table OA1: Spend Share of Meta Advertising Campaign Objectives

NOTES: This table presents the aggregated spend share of different campaign objectives across firms in the pre-ATT period. Spend share is defined to be the proportion of all spending for which a particular campaign objective is the source.

Table OA2: Spend Share of Meta Advertising Before vs After ATT: Categorized Objectives

Categorized campaign objective	Pre-ATT	Post-ATT
Off-platform conversions	95.7%	95.0%
On-platform reach	2.2%	2.5%
On-platform actions	2.1%	2.5%

NOTES: This table presents the aggregated spend share of different categorized campaign objectives, using the categorizations provided in Table OA1, across firms in the pre-ATT period (column 2) and the post-ATT period (column 3). Spend share is defined to be the proportion of all spending for which a particular campaign objective is the source over the relevant time period.

While Table OA2 documents that there was minimal aggregate spending away from offplatform conversions in the post-ATT period, we now turn to a firm-level analysis. Our main goal is to understand whether, at the firm level, firms reallocated spend on Meta away from off-platform conversion-optimized campaigns to on-platform action-optimized campaigns.

³This table does not include every objective listed previously, as Meta grouped and rebranded some of them in December 2021, changes that were rolled out slowly and that make up a small fraction of spending in 2022 (https://bit.ly/4gi074u).

We consider three dependent variables: the campaign objectives' share of spending, their share of impressions, and an indicator for whether spend for that campaign objective was non-zero. The former two metrics provide a measure of intensive margin substitution – to what extent do firms shift their share of spending more to on-platform actions – while the final metric provides a measure of extensive margin substitution – to what extent do firms start to run on-platform campaigns. We focus on shares for the intensive margin since we focus on relative reallocation within Meta.

		Dependent variable	2:
	(1)	(2)	(3)
	Spend share	Impression share	$\mathbb{1}(\mathrm{Spend}_t > 0)$
$After_t \times On-platform actions$	-0.007	0.004	0.015
	(0.007)	(0.009)	(0.011)
$After_t \times On-platform reach$	-0.003	0.008	0.021**
	(0.005)	(0.007)	(0.009)
Week FE	Yes	Yes	Yes
Firm-Campaign FE	Yes	Yes	Yes
Observations	72,228	72,228	72,228
\mathbb{R}^2	0.839	0.750	0.535

Table OA3: Meta Campaign Objective Substitution

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

NOTES: All results use the advertising dataset. We estimate specification (2) on a balanced panel of firms that have non-zero spend on Meta. 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 conversions. Standard errors clustered at the firm level.

We consider a balanced panel of Meta firms that have positive spend on any campaign objective throughout the sample period and estimate the within-firm difference-in-differences specification (2). As with the main analyses, this allows us to control for differences across firms. Table OA3 displays the results that, consistent with the aggregate spending in Table OA2, show a precise null effect on substitution to on-platform objectives on the intensivemargin. We note that there is an economically small, but statistically significant, substitution in the extensive margin to objectives optimizing for on-platform reach. As such, this motivates us to have our main analyses in Figure 3 estimated using a balanced panel of firms that use both click- and conversion-optimized objectives before ATT.

C.2 Additional Analyses for Advertising Platform Substitution

In this section, we consider additional analyses to provide a more complete picture of substitution across advertising platforms. We first show that conversions from Google were not as adversely impacted as conversions from Meta as a result of ATT. Then, we explore absolute trends in spending patterns between Meta and Google to show that they follow similar patterns before ATT, but noticeably diverge after ATT. Finally, we explore whether reallocation was more likely by more Meta-dependent firms.

Figure OA1: Event Study for Google Conversions



(a) Event Study for Google Search Conversions

(b) Event Study for Google Display Conversions



(c) Event Study for Overall Google Conversions

Google Advertising Effectiveness after ATT: Our primary comparison across advertising platforms is between Google and Meta. While in the main text we provide evidence that the quality of advertising targeting was degraded on Meta, here we provide event study

NOTES: Results use a balanced panel of firms using Google advertising from the advertising dataset. The plots represent the estimated event study coefficients from specification (1) with standard errors clustered at the firm level and data aggregated at the weekly level. Panels (a) and (b) restrict to Google Search and Display services respectively, while panel (c) includes the full set of Google advertising services.

estimates for the relative changes in the performance of Google advertising. To do so, we estimate the event study specification (1) for the log of conversions across all Google advertising as well as Google Display ads, which provides the closest targeting in the Google ecosystem relative to that offered by Meta, and Google Search ads, the largest Google advertising service.⁴ We report the results in Figure OA1. In contrast to the sudden and persistent drop-off in logged conversions observed on Meta in Figure 2 there is no discernible negative impact on conversions for the Google ecosystem or either Google Search/Display individually.



Figure OA2: Event Study for Online Advertising Spending

NOTES: Results use a balanced panel of firms with advertising spending from the advertising dataset. The plots represent the estimated event study coefficients from specification (1) with standard errors clustered at the firm level and data aggregated at the weekly level. Panels (a) and (b) consider the dependent variable as the log of online advertising spending for Meta and Google, respectively. Panel (c) considers the dependent variable as the log of online advertising across Meta, Google, and TikTok.

Changes in Total Advertising Spending: Given that we have some evidence that Google

⁴For the individual services, we report results only from January 1 until October 31, 2021 since Google launched its Performance Max product in November 2021, which led to substitution within the Google ecosystem that is orthogonal to ATT.

advertising is less impacted than Meta advertising, we explore whether and to what extent firms changed their online advertising shares towards Google as opposed to Meta. To do so, we focus mainly on changes in total advertising spending. We consider advertising spend, and not quantities of advertising purchased or their price, for the following reasons. First, the quantity variable varies across different types of advertising platforms and even within the same platform. For instance, Google Search is purchased per click, whereas Google Display is purchased per impression and on Meta firms can choose to pay per click or per impression. Furthermore, our data allows us only to observe end outcomes, so that we can only observe the average price of the ads that firms actually purchase.

Figure OA3: Event Study for Meta Online Advertising Share of Clicks and Impressions



NOTES: The plots represent the estimated event study coefficients from specification (1) with standard errors clustered at the firm level and data aggregated at the monthly level. Panels (a) and (b) consider the dependent variable of the share of observed clicks and impressions, respectively, attributed to Meta.

In the main text, we show that the share of online advertising spend on Meta advertising declines. In this section, we estimate the event study specification (1) for total online advertising spending as well Google and Meta advertising spending individually, and we plot the estimates in Figure OA2. Figure OA2 shows that spending is increasing on both Google and Meta advertising before ATT at a similar rate, and that, after ATT, the spending on Google continues on a similar trend, whereas on Meta it slowly declines over time. The increasing time trend for online advertising spending is consistent with the revenue time trend shown in Figure OA4. We cannot disentangle whether this comes from supply-side or demand-side

effects, but we additionally show in Figure OA3 that we get similar relative reductions in quantity of Meta ads using either clicks or impressions as our quantity measure.

		Dependent variable:					
	(1)	(2)	(3)				
Platform	Spend share	Impression share	Click share				
Google	0.048***	0.067***	0.057^{***}				
	(0.009)	(0.010)	(0.009)				
Meta	-0.044^{***}	-0.061^{***}	-0.057^{***}				
	(0.009)	(0.010)	(0.009)				
TikTok	-0.004	-0.005	-0.001				
	(0.002)	(0.003)	(0.003)				
Google Search	0.011	0.012	0.012				
0	(0.008)	(0.008)	(0.008)				
Google Display	0.005***	0.028***	0.013***				
	(0.002)	(0.006)	(0.004)				

Table OA4: Advertising Platform Substitution Estimates

*p<0.1; **p<0.05; ***p<0.01 Notes: Results use the advertising dataset. Each cell displays the estimated average treatment effect using specification (3) with the different specified dependent variables (columns) and subsetted to the relevant platforms (rows). The first three rows present the results of the difference-in-differences specifications for share of spend, impressions, and clicks on overall spending on Meta, TikTok, and Google. The final two rows present the same dependent variables for Google Search and Google Display products. The results for Google Search and Google Display are estimated over the period January to October 2021 as Google launched its popular Performance Max product in November 2021, which led to substitution within the products in the Google ecosystem. Standard errors are clustered at the firm level.

Which firms are more likely to reallocate? We now explore whether firms more dependent on Meta were more likely to reallocate their spending by estimating the across-firm difference-in-differences specification (3). We define the treated group as firms with aboveaverage Meta advertising spend as a proportion of total advertising spend in the advertising dataset.⁵ We consider a balanced panel of firms that spend non-zero dollars on any advertising platform throughout the same sample period as before and the considered dependent variables are the online advertising market share of impressions, clicks, and spending across

⁵The mean is reported in Table 1 as 0.75.

the different platforms. The first three rows of Table OA4 show the results for the online advertising market shares of Google, Meta, and TikTok. They suggest that for each of the measures that we consider Google benefited at the expense of Meta, gaining 4.8 to 6.7 percentage points of market share, whereas there was no shift in market share to TikTok. Furthermore, rows (4) and (5) of Table OA4 show the change in market share across different Google products and there is a greater increase in the share of Google Display relative to Google Search.

Appendix D Additional Analyses for Revenue

D.1 Robustness Checks for Revenue Effects

In this section we consider several robustness exercises to complement our main analyses.

D.1.1 Raw Data Plots

In Figures OA4 and OA5 we plot the demeaned log monthly revenue for the treated and control set of firms, where treatment is defined using the Meta revenue share and the iOS revenue share, respectively. In both cases, in the pre-ATT period we see nearly identical trends in revenue with modest monthly growth over time. Roughly when ATT takes effect, we see that this trend continues nearly linearly for the firms with a low Meta/iOS share of revenue, whereas the upward trend stops for the firms with a high Meta/iOS share of revenue, for which revenues over time flatten out.



Figure OA4: Log Revenue for Low and High Meta Shares

NOTES: Results using the revenue dataset. High and low Meta shares are calculated using a median split of pre-ATT revenue from Meta traffic. Plot shows log(revenue) demeaned using the pre-ATT mean along with 95% confidence intervals.



Figure OA5: Log Revenue for Low and High iOS Shares

NOTES: Results using the revenue dataset. High and low iOS shares are calculated using a median split of pre-ATT revenue from iOS traffic. Plot shows log(revenue) demeaned using the pre-ATT mean along with 95% confidence intervals.

D.2 Reduced New Customer Acquisition as a Mechanism

In this section we characterize the extent to which revenue reductions were due to a decline in new customer acquisition. We do so by analyzing a secondary dataset associated with the advertising dataset. In this secondary dataset, we observe aggregated revenue data from the set of firms in the advertising dataset that provide access to their Shopify account and can directly link this to their advertising spending. Importantly for our purposes, these data provide us with a complete view of revenue for the firms. In particular, we observe the total revenue, the number of orders, and the fraction of orders that come from repeat customers. The Shopify data have 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. Thus, the ability to separately measure new and repeat customers allows us to characterize the effects on new customer acquisition.

We note two significant limitations to this analysis. First, the number of firms that provide access to revenue data and are present through the full sample is relatively small. Second, the firms in the advertising dataset are strongly Meta-dependent and our representativeness exercises highlight that the advertising dataset skews towards smaller firms. As such, the estimates in this section will be relatively imprecisely estimated and underestimate the overall effect on orders and revenues, relative to our analyses in Section 4. Despite these limitations, the fact that we can decompose new and repeat customers provides additional evidence directly linking revenue changes to advertising, which we cannot do in our primary analyses.

For this analysis we estimate the across-firm difference-in-differences specification (3) defining treatment as whether Meta advertising spend was above the mean within the set of firms, and we use the measure of the fraction of orders processed by a merchant that come from repeated customers. We consider monthly sales measures and, after joining with the advertising data, compute each firm's pre-ATT spend share for Meta advertising.

Figure OA5 presents the results for the dependent variables that we observe from Shopify: log(revenue), log(order count), and repeat order ratio. The takeaway across each of these is consistent: there is a reduction in orders and revenue of 20-22% and the fraction of total orders coming from repeat customers has increased. To understand whether this is simply a result of shifting advertising spend, we estimate the difference-in-differences specification controlling for the log of advertising spend. These results are presented in the second row of Table OA5. 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 comparable and economically large negative point estimates. For the repeat order ratio, the pre-ATT baseline for the share of orders coming from repeat customers was 33.93%, implying a 10.5% increase in the share of orders coming from repeat customers.

To characterize the absolute impact on new and repeat orders respectively, we use the repeated order ratio combined with the total number of orders to estimate the effects on the number of orders coming from new and repeat customers, respectively. The results of estimating the same empirical specification (3) using the log of new and repeat orders as the dependent variables are presented in columns (2) and (3) of Table OA6. Column (2) shows a statistically significant 28.5% decrease in orders coming from new customers, but column (3) shows a negative, statistically insignificant, effect on repeat customer orders. The second row of Table OA6 shows that this result is robust to controlling for total online advertising spend. That being said, while the effect on the repeat customer ratio remains consistent when controlling for advertising spend, we find that the coefficients for revenue and orders decrease in magnitude. This suggests that some of the revenue decline can be attributed to changes in firms' total advertising spending patterns after ATT, even though the majority of the effect is due to decreased effectiveness of the advertising that is still being purchased.

In sum, this provides evidence that the revenue reductions are primarily due to weakened new customer acquisition and that there does not appear to be a countervailing effect of increased customer retention. If anything, our results point to reductions in revenues among repeat customers as well.

We conduct several robustness checks to validate the result that the primary reduction in orders comes from new customers. Figure OA6 considers the time-varying difference-indifferences specification with ad spending controls and provides evidence that the parallel trends assumption seems to reasonably hold. We then consider the same set of specifications using a negative binomial regression as an alternative to handling the small fraction of zeros in our data. Table OA7 presents the results for total, new customer, and repeat customer orders respectively. The results and effect sizes are largely consistent with our earlier analyses showing that the reduction in orders from new customers seems to be the driving force for the overall reduction in orders. The estimates for θ in Table OA7 indicate that there is moderate overdispersion in the data, supporting the suitability of the negative binomial

	Dependent variable:						
	(1)	(2)	(3)	(4)	(5)	(6)	
	$\log(O)$	rders)	log(Re	venue)	Repeat of	rder ratio	
$After_t \times Treated$	-0.242^{*} (0.129)	-0.171 (0.118)	-0.230^{*} (0.131)	-0.156 (0.119)	$3.994^{***} \\ (1.415)$	3.563^{**} (1.419)	
Ad spending controls Firm FE Month FE	No Yes Yes	Yes Yes Yes	No Yes Yes	Yes Yes Yes	No Yes Yes	Yes Yes Yes	
$\begin{array}{c} \text{Observations} \\ \text{R}^2 \end{array}$	5,772 0.854	5,772 0.872	$5,772 \\ 0.847$	5,772 0.867	5,772 0.808	5,772 0.813	

Table OA5: Difference-in-Differences Estimates for Sales

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

NOTES: Results use a balanced panel of firms within the advertising dataset for which Shopify transaction data are observed. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification (3) with and without controls for log(total advertising spending + 1). Standard errors are clustered at the firm level.

Table OA6: Difference-in-Differences Estimates for New vs. Repeat Customers

	Dependent variable:				
	(1)	(2)	(3)	(4)	
	log(New cust	omer orders $+ 1$)	$\log(\operatorname{Repeat}$	customer orders $+ 1$)	
$After_t \times Treated$	-0.337^{***} (0.128)	-0.257^{**} (0.116)	-0.154 (0.141)	-0.097 (0.132)	
Ad spending controls Firm FE Month FE	No Yes Yes	Yes Yes Yes	No Yes Yes	Yes Yes Yes	
$\begin{array}{c} \text{Observations} \\ \text{R}^2 \end{array}$	$5,772 \\ 0.829$	$5,722 \\ 0.852$	$5,722 \\ 0.877$	$5,722 \\ 0.886$	

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

NOTES: Results use a balanced panel of firms within the advertising dataset for which Shopify transaction data are observed. The columns are $\log(1 + \text{ orders})$ coming from new and repeat customers. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification (3), where odd columns do not control for $\log(1 + \text{ total advertising spending})$ and even columns do control for them. Standard errors are clustered at the firm level.

		Dependent variable:					
	(1)	(2)	(3)	(4)	(5)	(6)	
	Total	orders	Orders from 1	new customers	Orders from repeat customers		
$After_t \times Treated$	-0.165^{***} (0.033)	-0.134^{***} (0.035)	-0.243^{***} (0.034)	-0.208^{***} (0.032)	-0.088^{**} (0.034)	-0.061^{*} (0.033)	
Ad spending controls	No	Yes	No	Yes	No	Yes	
Firm FE Month FE	Yes Yes	Yes Yes	Yes Yes	Yes Yes	Yes Yes	Yes Yes	
Observations θ	5,772 3.262^{***} (0.059)	5,772 3.597^{***} (0.066)	5,772 2.814^{***} (0.051)	5,772 3.142^{***} (0.057)	5,772 3.231^{***} (0.061)	5,772 3.449^{***} (0.066)	

Table OA7: Difference-in-Differences Estimates for New vs. Repeat Customers (Negative Binomial Specification)

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

NOTES: Results use a balanced panel of firms within the advertising dataset for which Shopify transaction data are observed. The first two columns use total orders as the dependent variable, the next two focus on new customer orders, and the last two on repeat customer orders. The odd columns do not control for $\log(\text{total ad spending} + 1)$, whereas the even columns do. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification (3) estimated using a negative binomial model. Standard errors are clustered at the firm level.



Figure OA6: Time-Varying Treatment Effects for Sales Outcomes

(c) DiD Estimates for New Orders

NOTES: Results use a balanced panel of firms within the advertising dataset for which Shopify transaction data are observed. The dependent variables on the top row from left to right: log(orders), Repeat order ratio. The dependent variable on the bottom row is log(new customer orders + 1). Plots are the time-varying estimates for the difference-in-differences specification (2) with controls for the log(total advertising spend + 1). Standard errors are clustered at the firm level.

Appendix E Data Representativeness

In this section, we discuss the representativeness of the two main datasets – the advertising and revenue datasets – that we employ in the main analyses. The primary threat to representativeness stems from firms self-selecting to share data with the respective providers. To address this, we provide transparency into the nature of this selection process and its implications for the empirical data. First, we outline the selection mechanisms to better understand which types of firms have the strongest incentives to opt in. Second, and most importantly, we compare the datasets to consumer-level benchmarks, including advertising spending and session count/revenue data, which are not subject to firm-side self-selection.

Summarizing the main results from the analyses below, we find that the advertising dataset aligns with the broader population of e-commerce firms in terms of total advertising spending but skews towards smaller-sized firms. The revenue dataset similarly aligns well with the broader population of e-commerce firms in terms of size, as measured by session counts, and trends in size over time, as measured by session counts and revenue, showing a slight skew towards relatively larger firms than those in the advertising dataset. As such, while our ability to assess representativeness is limited by the availability of external benchmark data, the external benchmark data that are available to be analyzed provides empirical support for the representativeness of the two datasets vis-à-vis a broad cross-section of e-commerce firms along key dimensions.

E.1 Advertising Data Representativeness

In this section, we assess the representativeness of the advertising data. The main identification threat is that we only observe data for firms that opt into their data being tracked by the analytics firm that gave us access to the advertising data. This may induce selection in the type of firms that we observe, threatening the external validity of the resulting estimates. As such, we discuss of the nature of selection into the advertising dataset, then empirically evaluate representativeness of the advertising dataset versus external benchmark data.

E.1.1 Selection into Sample

The analytics firm that provided us with the advertising data provides benchmarking of key business metrics across industries using aggregated and anonymized data sourced directly from participating advertisers. This analytics firm has a "give to get" business model in which firms that allow their data to be tracked are, in turn, given access to the anonymized and aggregated benchmarking data by the analytics firm. As such, firms opt into the advertising dataset to gain access to the performance of their marketing campaigns relative to other firms. It allows them to learn the acquisition costs of similar firms and for advertising campaigns that are targeting similar segments. Second, because the data are updated in near real-time, firms that opt in can better understand whether short-term changes in marketing performance are idiosyncratic to them or reflect broader market-level changes. The value proposition of the advertising data provider's analytics platform, which primarily revolves around competitive insights, is not inherently skewed towards a particular firm size. Ex ante, it is not obvious that participation in the advertising dataset would be systematically dominated by either smaller or larger firms.

E.1.2 Online Advertising Spending

To assess the representativeness of this dataset, we manually collected advertising spending across all media using Kantar's Vivvix Advertising Intelligence Product.⁶ For online advertising, Kantar generates its dataset using a combination of automated web crawlers that repeatedly scrape advertisements and a panel of 1.2 million consumers with technology installed on their devices to track exposure information. By pairing these exposure data with rate cards, Kantar estimates total advertising spending. Importantly, this methodology allows Kantar to provide comprehensive estimates of advertising spend across firms without

⁶https://www.vivvix.com/home

firm-side self-selection into the dataset, which was the main identification threat associated with our advertising dataset.⁷

Kantar-Vivvix is one of the largest advertising intelligence databases used in academic research and industry, covering approximately \$100 billion in annual advertising spending across 4 million brands and 3 million advertisers, and it is used by firms such as Procter and Gamble, Unilever, and Google. Its digital coverage is more comprehensive than that of its primary competitor, Nielsen Ad Intel. As such, while it may not have complete coverage of advertising spending, it offers coverage that is more complete than that of available alternatives, to the best of our knowledge.

While the Kantar-Vivvix dataset has wide coverage, it provides us with a cruder measure of firm behavior than our advertising dataset, which contains richer data about their online advertising campaigns, including the breakdown by platform and campaign type, as well as direct measures of conversion rates as reported by the advertising platforms. These measures are not tracked by Kantar-Vivvix, making the Kantar-Vivvix dataset not adequate for our main analysis, but valuable as an external benchmarking dataset to assess representativeness.

The Kantar-Vivvix dataset contains firms spanning many industries that fall outside of ecommerce, such as retail, consumer packaged goods, automotive, and financial services, which are not relevant to our analysis. Therefore, we collect data from BuiltWith⁸ to subset down to firms operating on Shopify, which serves as a proxy for the e-commerce retailer category. BuiltWith is a technology-profiling firm that identifies the underlying technologies used by websites, including e-commerce platforms such as Shopify. It does so without selection bias, as BuiltWith collects publicly available information, using web crawlers to examine HTML, JavaScript, CSS, and other code on those web pages, thus identifying embedded technologies without firm participation or opt-in. The resulting collection of 7,900 firms within this Kantar-Shopify dataset is broadly representative of a wide cross-section of the

⁷We use the version of the Kantar-Vivvix dataset that provides comprehensive spending coverage across mobile, desktop, and video, thus enabling credible measurement of online advertising spend.

⁸https://builtwith.com



Figure OA7: Advertising Data vs. Kantar-Vivvix (Shopify) Online Ad Spending Distribution

NOTES: This figure and table compare the kernel density estimate and summary statistics of online advertising spending between January 1, 2021 and March 31, 2021 for the focal advertising data (Advertising Data) and the Kantar-Vivvix Shopify subsample (Kantar-Vivvix). The figure presents a visual comparison using a kernel density estimate, while the table provides summary statistics of the logarithm of total online advertising spending. "SD" represents the empirical standard deviation. Both samples include firms with advertising spending in each week during this period.

e-commerce retailer category.⁹

We compare the total advertising spend distribution of the resulting Kantar-Shopify firms with the corresponding advertising spend distribution for the firms within the advertising dataset over the period of time between January 1, 2021 and March 31, 2021, since the Kantar-Vivvix data collects mobile ad spending starting at the beginning of 2021, and April 2021 is when ATT is rolled out. We include only firms with advertising spending in each week during this period. This allows us to assess the representativeness of the advertising data to the e-commerce retailer category in terms of total advertising spend.

⁹As the Kantar-Vivvix data do not provide a domain name, we match firms within Kantar-Vivvix and firms in the advertising dataset manually by firm name.

We compare the distribution and summary statistics of the log of total online advertising spend in Figure OA7. Overall, these results indicate that while both distributions span a similar range of firm sizes (i.e., have common support), the advertising dataset shows an overrepresentation of smaller firms compared to the broader e-commerce retailer population as captured by the Kantar-Vivvix benchmark, with a mean and median that are both approximately 7.2% smaller. These findings suggest that the advertising dataset reasonably approximates the e-commerce retailers in terms of total advertising spending, with a slight overrepresentation of smaller firms.

E.2 Revenue Data

As with the advertising dataset, we assess the representativeness of the revenue dataset by first discussing the nature of selection into it, and then empirically comparing it against two benchmark datasets. The first benchmark dataset consists of publicly disclosed merchant revenues from Shopify, a widely used e-commerce platform. The second benchmark dataset comes from SimilarWeb, a leading provider of web traffic and performance metrics. We use these datasets to empirically evaluate the representativeness of the revenue data in terms of firm size and changes in revenue over time.

E.2.1 Selection into Sample

To better understand the nature of selection into the revenue dataset, we describe the main incentives that firms have to use our data provider, Grips Intelligence. Grips acquires its data through services and analytics that firms can access in exchange for providing their data. These services provide insights into business performance, in absolute terms and relative to competing firms. Grips' platform does not exclusively cater to a particular size category of firms because, similar to the reasoning summarized above for the advertising dataset, the competitive benchmarking data that Grips provides can be valuable to firms regardless of their size. As such, it is not a priori obvious that the resulting sample would be dominated by one size category of firms rather than another.

For this research, Grips compiled the dataset by first identifying all firms with active API access as of December 2023 that had data points for each of the years 2019, 2020, 2021, and 2022, resulting in an initial pool of 1,807 candidate firms. We then filtered these candidate firms down to those with complete coverage over the full observation period, no missing revenue data at the monthly level (e.g., due to tracking errors), and no missing session count data at the daily level over the entirety of the observation period, resulting in a final collection of 773 firms which are used in our main analysis.

E.2.2 Representativeness Across Time

We evaluate whether changes in total revenue across time within the revenue data reflect changes in total revenue across time within a broader population of e-commerce firms. Establishing alignment in revenue over time mitigates concerns about dataset-specific biases/artifacts, which is important since our identification strategy leverages across-time comparisons. We assess this in two ways. First, we use publicly disclosed financial data from Shopify's 2023 Investor Day presentation.¹⁰ Second, we utilize data from SimilarWeb, a leading and widely-used provider of digital intelligence and analytics. We first present the analysis using Shopify data, then present the corresponding analysis using the SimilarWeb data.

The financial data disclosed by Shopify in the aforementioned Investor Day presentation provide a population-level view of revenue generated by Shopify merchants, making the resulting analysis a natural complement to the representativeness analysis performed above on the advertising dataset, as both analyses leverage Shopify firms as an empirical benchmark for assessing our datasets.

Importantly, the Shopify public disclosure data segments merchants into acquisition cohorts based on the year in which those merchants began selling on Shopify. This cohort-based

¹⁰Shopify Investor Day 2023, page 126. Available at: https://s203.q4cdn.com/784886181/files/doc_presentations/Shopify-Investor-Day-2023-Presentation.pdf.

segmentation is particularly valuable for constructing an "apples-to-apples" comparison with the revenue dataset, which includes only merchants that began operating before the beginning of the pre-treatment period, which is one year prior to the rollout of ATT (i.e., before April 2020). Correspondingly, because merchants are segmented into annual cohorts, we compute revenue figures using the Shopify benchmark data for firms that began selling on Shopify prior to 2020 as this allows for the closest match between the two respective datasets in terms of merchant acquisition dates. By aligning cohorts in both datasets, we isolate revenue trends for existing (i.e., pre-treatment) merchants while minimizing potential confounding effects from newly onboarded firms.



Figure OA8: Time Series of Overall Revenue: Revenue Dataset vs. Shopify

NOTES: This figure compares total revenue each quarter, as observed through the revenue data, to total revenue across Shopify merchants, as observed through cohort-level data publicly disclosed by Shopify. Firms within the revenue dataset began operating before April 2020; only merchants that began selling through Shopify prior to 2020 underlie the Shopify figures. Data series are mean-scaled relative to the Shopify sample to facilitate visual comparison across time.

Figure OA8 shows the resulting comparison of Shopify revenue to our revenue dataset over time, after scaling the latter so that it has the same empirical mean as the former so as to facilitate visual comparability. The two resulting time series show consistent trends over time, supporting the notion that spending trends over time in the revenue dataset are representative of broader trends within the e-commerce retailer category and are not artifacts of the revenue dataset. The empirical correlation of these two time series is 97.2%.

We provide a second across-time comparison using data from SimilarWeb. As noted above, SimilarWeb is a provider of digital intelligence and analytics that estimates domainspecific sessions for over 100 million websites using diverse clickstream data sources, including anonymous traffic from millions of devices and partnerships with DSPs, ISPs, and other measurement firms. The nature and breadth of SimilarWeb's data sources minimize firmside selection bias, making it an appropriate benchmark.

For the sake of consistency with our other representativeness analyses, we obtained data on the top 7,000 e-commerce firms tracked by SimilarWeb (ranked by session counts from 2019 to 2021). From within this set, we then filtered down to the firms that maintained positive sessions in 2019, mirroring the selection criteria used for the revenue data.



Figure OA9: Time Series of Overall Sessions: Revenue Dataset vs. SimilarWeb

NOTES: This figure compares session counts over time across three samples: (1) the revenue dataset (mean-scaled), (2) all firms tracked by SimilarWeb, and (3) the subset of SimilarWeb firms operating on Shopify (mean-scaled). Data series are mean-scaled relative to the full SimilarWeb sample to facilitate visual comparison across time.

We compare the revenue data to two samples from the SimilarWeb data: all firms and the subset of firms that operate on Shopify (as identified using BuiltWith). The former sample provides a comprehensive view of e-commerce retailers across multiple platforms, offering broader insights into industry-wide patterns. The latter Shopify-focused collection of firms is more complementary to the Shopify public disclosures and the Kantar-Vivvix-Shopify dataset in that all three datasets represent populations of e-commerce retailers using Shopify. Taken together, these analyses enable a more comprehensive assessment of the robustness/sensitivity of our dataset's representativeness over time as we vary firm sizes and size measures.

The results are shown in Figure OA9. We again see strong correspondence between trends in session counts over time between the revenue and both the SimilarWeb and SimilarWeb-Shopify datasets. The empirical correlations between the revenue and SimilarWeb/SimilarWeb-Shopify time series are 87.2% and 79.4%, respectively.

Taking these results together, the close alignment of revenue and session count trends across these two benchmarking datasets suggests that the revenue data capture broader ecommerce market temporal dynamics with reasonable fidelity, and that these findings are robust across different measures of firm activity and sample definitions.

E.2.3 Representativeness in Firm Size

Finally, we compare the distribution of firm size in the revenue data to benchmark data from SimilarWeb using both the full SimilarWeb sample and the subset of firms operating on Shopify. The results, shown in Figure OA10, demonstrate that the distribution of yearly session counts for calendar year 2020 in the revenue dataset aligns reasonably well with both SimilarWeb samples. The revenue dataset's mean (14.71) falls between that of the full SimilarWeb sample (14.99) and the Shopify-only subsample (14.42), with similar patterns for the median values. The revenue dataset exhibits slightly more dispersion (SD = 1.85) than both the full SimilarWeb sample (SD = 1.63) and the Shopify-only subsample (SD = 1.30).

Overall, these comparisons suggest that the revenue dataset spans a similar range of firm sizes as the benchmark samples, though with different frequency distributions. The revenue



Figure OA10: Revenue Dataset vs. SimilarWeb Session Distribution and Summary Statistics

NOTES: This figure displays the kernel density estimate of yearly session counts (log-transformed) across three samples: revenue dataset retailers, SimilarWeb's full e-commerce sample, and SimilarWeb's Shopify-only retailers. The accompanying table presents summary statistics for these distributions, also on a logarithmic scale, including mean, median, standard deviation (SD), 25th and 75th percentiles, and sample size (N). All data are based on firms with recorded session counts.

dataset exhibits wider variance in firm sizes compared to the SimilarWeb samples, with an empirical mean that falls between the full sample and Shopify-only sample means. The wider variance observed is likely explained by the data collection methodology for the SimilarWeb data, with both samples being derived from the largest 7,000 firms by size; restricting to the larger firms naturally compresses the range of firm sizes and thus reduces variance.

We further note that the revenue dataset spans the full spectrum of firm sizes, enabling credible analysis of heterogeneous treatment effects (HTE) by firm size (Tables ?? and ??). Through these HTE analyses, we find that smaller firms are more negatively impacted by ATT than larger firms. As a result, our main effect estimates could be considered representative of mid- to large-sized e-commerce retailers, and a conservative lower bound on the

effects for target populations with greater representation of smaller e-commerce retailers.

In summary, our analyses demonstrate that both datasets have full support across the spectrum of e-commerce retailers, but with different frequency distributions compared to benchmark populations. The advertising dataset oversamples smaller firms, while the revenue dataset includes relatively more larger firms. These differences in sample composition are important for interpretation, which is why we conduct heterogeneous treatment effect analyses in Section 4 of the main paper that explicitly examine how ATT's impact varies with firm size.

Appendix F Conceptual Framework

This section provides a conceptual model to characterize the equilibrium effects of ATT where firms allocate online advertising spending between two advertising platforms that differ in the degree to which they rely on behavioral targeting. We ultimately characterize how ATT should lead to overall spending allocations to differ between these two platforms, consistent with the empirical results we describe in the main text.

F.1 Firm Behavior

There is a unit mass of retail firms indexed by $(\pi, \theta) \in \mathbb{R}_+ \times [0, 1]$, where π is the profit per each acquired customer and θ is the retailer's type vis-à-vis its preference for ad network to acquire customers. (π, θ) is distributed according to a distribution Φ , which admits strictly positive density ϕ for all interior (π, θ) . They purchase ads from two outlets, F and G. If a firm of type θ purchases ads $a_F \geq 0$ and $a_G \geq 0$ from F and G, respectively, they acquire a mass of consumers,

$$q(a_F, a_G, \theta) := f(a_F, \theta) + g(a_G, \theta) - h(a_F, a_G, \theta).$$

The interpretation is that the firm acquires $f(a_F, \theta)$ and $g(a_G, \theta)$ from the two networks, but out of them $h(a_F, a_G, \theta)$ accounts for the multi-homing consumers who receive ads from both networks and would have been acquired only from an ad from either network. We thus subtract them to avoid double counting.

We make the following assumptions:

Assumption 1. (i) $f(0, \cdot) = g(0, \cdot) = 0$ and $h(a_F, a_G, \theta) = 0$ for all θ if $a_F a_G = 0$. (ii) f, g and g are twice differentiable, and $\frac{\partial f(0,\theta) - h(0, \cdot, \theta)}{\partial a_F} = \frac{\partial g(0,\theta) - h(\cdot, 0,\theta)}{\partial a_G} = \infty$ for all θ ; and (iii) $h(a_F, a_G)$ is strictly supermodular.

The first two are self-evident, clearly justified by the setup. The assumption (iii) means that as a_G increases, the marginal benefit of a_F fall since the multihoming consumers are more likely to be reached from both networks.

The next assumption captures how θ represents the retailer's relative preference for F and G.

Assumption 2. $f(a_F, \theta)$ is increasing in (a_F, θ) , and $g(a_G, \theta)$ is increasing in $(a_G, -\theta)$, and $q(a_F, a_G, \theta)$ is strictly concave in (a_F, a_G) . Further, $f(a_F, \theta) - h(a_F, a_G, \theta)$ is strictly supermodular in (a_F, θ) and $g(a_G, \theta) - h(a_F, a_G, \theta)$ is strictly supermodular in $(a_G, -\theta)$.

One interpretation is that F is like the Meta ad network, which specializes in behavioral targeting, which some firms prefer relative to G (e.g., Google), whose ads are less behaviorally targeted. The supermodularity assumption means that the marginal benefit of the ad at F increases in θ , and the marginal benefit of the ad at G decreases in θ .

We next add an assumption that implies that both a_F and a_G are "normal" goods for the firm.

Assumption 3. For any $(a'_F, a'_G) \neq (a_F, a_G)$ such that $a'_F + a'_G \geq a_F + a_G$, we have either

$$\frac{\partial (f(a'_F, \theta) - h(a'_F, a'_G, \theta))}{\partial a_F} < \frac{\partial (f(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_F}$$

or

$$\frac{\partial (g(a'_G, \theta) - h(a'_F, a'_G, \theta))}{\partial a_G} < \frac{\partial (g(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_G}.$$

Specifically, if the firm purchases higher total units of advertising from the two outlets, the marginal benefit for at least one advertising must be lower.

Example 1 (Microfoundation based on urn-ball model). Suppose there are two outlets Fand G. There are three types of consumers: those who single-home F, those who single-home G, and those who subscribe to both outlets. Let x_i be the measure of consumers of type $i \in \{F, G, FG\}$. Suppose a firm with type θ purchases a_j ads from outlet j. An amount aof ads placed at outlet j has the effective units $\lambda^j(\theta)a$ of ads directed at type θ . The scaling factor $\lambda^j(\theta)$ captures the targeting ability of j for type θ and the relevance of ad product to type θ . The precise interpretation is that instead of units a randomly landing the eyeball of a consumer at j, it is as if units $\lambda^j(\theta)a$ are randomly landing at the representative user. We can assume that $\lambda^F(\theta)$ is an increasing function. For example, we can take $\lambda^F(\theta) := \lambda^F + \theta$, and $\lambda^G(\theta) := \lambda^G + (1 - \theta)$, for some constants $\lambda^F, \lambda^G > 0$.

Let $X_j = x_j + x_{FG}$ is the total subscribers of j = F, G. Suppose a firm purchases ads (a_F, a_G) . Take a representative consumer single homing F. The probability of such a consumer "missing" the ads by that firm is

$$\left(1 - \frac{1}{X_F m}\right)^{\lambda^F(\theta)a_F m} \to e^{-\lambda^F(\theta)a_F/X_F} \text{ as } m \to \infty$$

So, the number of single-homing customers acquired by F is: $x_F \left(1 - e^{-\lambda^F(\theta)a_F/X_F}\right)$. Similarly, the number of single-homing customers acquired by G is $x_G \left(1 - e^{-\lambda^G(\theta)a_G/X_G}\right)$.

Now consider the dual-homing consumers. The probability of such a consumer "missing" the ads by that firm is

$$\left(1 - \frac{1}{X_F m}\right)^{\lambda^F(\theta)a_F m} \left(1 - \frac{1}{X_G m}\right)^{\lambda^G(\theta)a_G m} \to e^{-\lambda^F(\theta)\frac{a_F}{X_F} - \lambda^G(\theta)\frac{a_G}{X_G}}, \ as \ m \to \infty$$

So the number of dual-homing consumers acquired by the firm is:

$$x_{FG}\left(1-e^{-\lambda^{F}(\theta)\frac{a_{F}}{X_{F}}-\lambda^{G}(\theta)\frac{a_{G}}{X_{G}}}\right).$$

Hence, the total number of consumers acquired is:

$$\begin{aligned} x_F \Big(1 - e^{-\lambda^F(\theta)a_F/X_F} \Big) + x_G \Big(1 - e^{-\lambda^G(\theta)a_G/X_G} \Big) + x_{FG} \Big(1 - e^{-\lambda^F(\theta)\frac{a_F}{X_F} - \lambda^G(\theta)\frac{a_G}{X_G}} \Big) \\ = & X_F \Big(1 - e^{-\lambda^F(\theta)a_F/X_F} \Big) + X_G \Big(1 - e^{-\lambda^G(\theta)a_G/X_G} \Big) \\ & - x_{FG} \Big(1 - e^{-\lambda^F(\theta)a_F/X_F} - e^{-\lambda^G(\theta)a_G/X_G} + e^{-\lambda^F(\theta)\frac{a_F}{X_F} - \lambda^G(\theta)\frac{a_G}{X_G}} \Big) \\ = & X_F \Big(1 - e^{-\lambda^F(\theta)a_F/X_F} \Big) + X_G \Big(1 - e^{-\lambda^G(\theta)a_G/X_G} \Big) - x_{FG} \Big(1 - e^{-\lambda^F(\theta)a_F/X_F} \Big) \Big(1 - e^{-\lambda^G(\theta)a_G/X_G} \Big) \end{aligned}$$

The microfoundation for our model is therefore:

$$f(a_F, \theta) := X_F \left(1 - e^{-\lambda^F(\theta)a_F/X_F} \right),$$

$$g(a_G, \theta) := X_G \left(1 - e^{-\lambda^G(\theta)a_G/X_G} \right),$$

$$h(a_F, a_G) := x_{FG} \left(1 - e^{-\lambda^F(\theta)a_F/X_F} \right) \left(1 - e^{-\lambda^G(\theta)a_G/X_G} \right).$$

For a range of (a_F, a_G) , the above assumption is satisfied: for $a_j \lambda^j(\theta) < X_j$, $f(a_F, \theta)$ is supermodular in (a_F, θ) and $g(a_G, \theta)$ is supermodular in $(a_G, -\theta)$, and h is supermodular in (a_F, a_G) .

The current specification does not satisfy Assumption 1-(ii), but its sole purpose is to facilitate the analysis (based only on first-order conditions), so it is not essential. \Box

Now, we are in a position to characterize the firm's behavior with regard to the optimal purchase of advertising. The firm with type (π, θ) solves:

$$\max_{a_F, a_G} u(a_F, a_G; \theta, p_F, p_G) := \pi q(a_F, a_G, \theta) - p_F a_F - p_G a_G,$$

where p_F and p_G are the per unit price for ads placed at F and G.

Let $a_F(p_F, p_G, \pi, \theta)$ and $a_G(p_F, p_G, \pi, \theta)$ denote the optimal solution to the problem, and let $v(p_F, p_G, \pi, \theta)$ denote the maximized value.

Proposition 1. (i) $(a_F, -a_G)$ is increasing in $(\theta, -p_F, p_G)$. (ii) $v(p_F, p_G, \pi, \theta)$ is supermodular in $(\theta, -p_F, p_G)$; (iii) $v(p_F, p_G, \pi, \theta)$ is increasing in π ; (iv) if $(p'_F, p'_G) > (p_F, p_G)$, then $a_F(p'_F, p'_G, \pi, \theta) + a_G(p'_F, p'_G, \pi, \theta) < a_F(p_F, p_G, \pi, \theta) + a_G(p_F, p_G, \pi, \theta)$.

Proof. By Assumption 1-(iii) and by Assumption 2, the direct objective function is strictly supermodular in $(a_F, -a_G, \theta, -p_F, p_G)$. Further, the optimal solution is unique, given the strict concavity assumption. The first two results then follow from these two observations. The last result is also obvious, established easily by a revealed preference argument. For (iv), suppose to the contrary $a'_F + a'_G \ge a_F + a_G$, where $a'_F := a_F(p'_F, p'_G, \pi, \theta)$, $a'_G := a_G(p'_F, p'_G, \pi, \theta), a_F := a_F(p_F, p_G, \pi, \theta)$, and $a_G := a_G(p_F, p_G, \pi, \theta)$. By the first order condition:

$$\pi \frac{\partial (f(a'_F, \theta) - h(a'_F, a'_G, \theta))}{\partial a_F} = p'_F \ge p_F = \pi \frac{\partial (f(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_F}$$

and

$$\pi \frac{\partial (g(a'_G, \theta) - h(a'_F, a'_G, \theta))}{\partial a_G} = p'_G \ge p_G = \pi \frac{\partial (g(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_G}.$$

We thus have a contradiction to Assumption 3.

The characterization is clear. Firms with higher θ purchase relatively more add from F than from G. The relative add demand exhibits substitution effects; as $(p_F, -p_G)$ rises, firm reduces a_F and increases a_G .

We now consider an effect brought about by ATT:

Assumption 4. ATT decreases F's effectiveness by shifting (f,h) to a new function, (\hat{f},\hat{h}) , satisfying the above assumptions, such that (i) $\hat{f}(a_F,\theta) - \hat{h}(a_F,a_G,\theta) < f(a_F,\theta) - h(a_F,a_G,\theta)$ for all $a_F > 0$ and for all θ ; (ii) $\frac{\partial(\hat{f}(a_F,\theta) - \hat{h}(a_F,a_G,\theta))}{\partial a_F} < \frac{\partial(f(a_F,\theta) - h(a_F,a_G,\theta))}{\partial a_F}$, and

 $\frac{\partial (g(a_F,\theta) - \hat{h}(a_F,a_G,\theta))}{\partial a_G} > \frac{\partial (g(a_F,\theta) - h(a_F,a_G,\theta))}{\partial a_F} \text{ for all } a_F \ge 0 \text{ and for all } \theta; \text{ and (iii) } f(a_F,\theta) - h(a_F,a_G,\theta) - \hat{h}(a_F,a_G,\theta) - \hat{h}(a_F,a_G,\theta) - \hat{h}(a_F,a_G,\theta) = \hat{h}(a_F,$

With the urn-ball model example, ATT can be modeled as the reduction of the adefficiency parameter for F: that is, $\lambda^F(\cdot)$ is replaced by $\hat{\lambda}(\cdot) < \lambda(\cdot)$, which then affects f and h in the way consistent with Assumption 4.

Let $\hat{a}_F(p_F, p_G, \pi, \theta)$ and $\hat{a}_G(p_F, p_G, \pi, \theta)$ denote the optimal solution to the problem under \hat{f} , and let $\hat{u}(p_F, p_G, \pi, \theta)$ and $\hat{v}(p_F, p_G, \pi, \theta)$ denote the direct and indirect objective functions, respectively.

Proposition 2. $\hat{a}_F(p_F, p_G, \pi, \theta) < a_F(p_F, p_G, \pi, \theta)$ and $\hat{a}_G(p_F, p_G, \pi, \theta) > a_G(p_F, p_G, \pi, \theta)$, and $\hat{a}_F(p_F, p_G, \pi, \theta) + \hat{a}_G(p_F, p_G, \pi, \theta) < a_F(p_F, p_G, \pi, \theta) + a_G(p_F, p_G, \pi, \theta)$. Further, the loss from the shift $v(p_F, p_G, \pi, \theta) - \hat{v}(p_F, p_G, \pi, \theta)$ is nonnegative for all firms, and strictly positive for firms with $a_F(p_F, p_G, \pi, \theta) > 0$, and is increasing in θ .

Proof. The first statement follows from the fact that the direct objective before and after the shift is supermodular in $(a_F, -a_G)$ and from Assumption 4-(iii). The second statement follows from Assumption 3. Namely, suppose to the contrary $\hat{a}_F + \hat{a}_G \ge a_F + a_G$. Then, by Assumption 4-(iii),

$$\pi \frac{\partial (\hat{f}(\hat{a}_F, \theta) - h(\hat{a}_F, \hat{a}_G, \theta))}{\partial a_F} \le \pi \frac{\partial (f(\hat{a}_F, \theta) - h(\hat{a}_F, \hat{a}_G, \theta))}{\partial a_F} \le \pi \frac{\partial (f(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_F} = p_F$$

of

$$\pi \frac{\partial (g(\hat{a}_G, \theta) - h(\hat{a}_F, \hat{a}_G, \theta))}{\partial a_G} \le \pi \frac{\partial (g(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_G} = p_G$$

where the last inequalities follow from the fact that (a_F, a_G) is an optimal decision given (p_F, p_G) . Both inequalities become strict unless $(\hat{a}_F, \hat{a}_G) = (a_F, a_G)$, in which case we have a contradiction. If $(\hat{a}_F, \hat{a}_G) = (a_F, a_G)$, then the first inequality becomes strict, which violate the first order condition under ATT.

The last statement follows from fact that

$$tu(a_F, a_G; \theta, p_F, p_G) + (1-t)\hat{u}(a_F, a_G; \theta, p_F, p_G)$$

is supermodular in (θ, t) , which implies that the corresponding indirect objective function indexed by parameters (θ, t) is also supermodular.

The implication is clear. After the shift, the firms substitute their ad purchase away from F toward G, and all firms are worse off, and the firms with higher θ suffer higher losses.

F.2 Equilibrium of Ads Markets

We now consider the market equilibrium. First, we assume that each firm incurs fixed cost $\kappa > 0$, so the firm that never earns enough to cover the cost will exit the market.

Let $\pi(\theta)$ denote marginal active type (π, θ) such that $v(\pi(\theta), \theta, p_F, p_G) = \kappa$. Then, the demand for platform i = F, G is

$$D_i(p_F, p_G) := \int_0^1 \int_{\pi(\theta)}^\infty a_i(\pi, \theta, p_F, p_G) \phi(\pi|\theta) d\pi \phi(\theta) d\theta.$$

By Proposition 1, the ad demand for F, $D_F(p_F, p_G)$, is decreasing in $(p_F, -p_G)$, and the ad demand for G, $D_G(p_F, p_G)$, is increasing in $(p_F, -p_G)$.

We assume platforms i = F, G, incurs costs $c_i(A_i)$ for delivering total mass of ads A_i , where c_i is increasing and strictly convex.

We consider that markets are competitive so that ad prices are determined at levels that clear the markets: (p_F, p_G) are **market-clearing** or **equilibrium** if

$$p_F = C'_F(D_F(p_F, p_G))$$
 and $p_G = C'_G(D_G(p_F, p_G)).$

Proposition 3. Suppose the advertising technology shifts from (f, g) to (\hat{f}, g) as assumed in Assumption 4. The equilibrium exists both before and after the change. The equilibrium
prices change from (p_F, p_G) to (\hat{p}_F, \hat{p}_G) , where $\hat{p}_F < p_F$. In equilibrium, $\hat{D}_F(\hat{p}_F, \hat{p}_G) < D_F(p_F, p_G)$.

Proof. Suppose to the contrary $\hat{p}_F \ge p_F$. There are two possibilities. Suppose first $\hat{p}_G \le p_G$. In this case, note first that, for each type (π, θ) ,

$$\hat{a}_F(\hat{p}_F, \hat{p}_G, \pi, \theta) < a_F(\hat{p}_F, \hat{p}_G, \pi, \theta) \le a_F(p_F, p_G, \pi, \theta),$$

where the first inequality follows from Proposition 2, and the second follows from Proposition 1-(i). Consequently, we have

$$\hat{D}_F(\hat{p}_F, \hat{p}_G) = \int_0^1 \int_0^\infty \hat{a}_F(\pi, \theta, \hat{p}_F, \hat{p}_G) \phi(\pi|\theta) d\pi \phi(\theta) d\theta$$
$$< \int_0^1 \int_0^\infty a_F(\pi, \theta, p_F, p_G) \phi(\pi|\theta) d\pi \phi(\theta) d\theta$$
$$= D_F(p_F, p_G).$$

Then, we have

$$\hat{p}_F = C'_F(\hat{D}_F(\hat{p}_F, \hat{p}_G)) < C'_F(D_F(p_F, p_G)) = p_F,$$

where we use the market clearing condition and the convexity of C_F . We thus have a contradiction.

Next, $\hat{p}_F \ge p_F$ and $\hat{p}_G > p_G$. Then, by Proposition 1-(iv), we have, for all (π, θ) ,

$$\sum_{i=F,G} a_i(p_F, p_G, \pi, \theta) > \sum_{i=F,G} a_i(\hat{p}_F, \hat{p}_G, \pi, \theta) \ge \sum_{i=F,G} \hat{a}_i(\hat{p}_F, \hat{p}_G, \pi, \theta).$$

This means that

$$\sum_{i=F,G} D_i(p_F, p_G) > \sum_{i=F,G} D_i(\hat{p}_F, \hat{p}_G) \ge \sum_{i=F,G} \hat{D}_i(\hat{p}_F, \hat{p}_G).$$

Hence, either $D_F(p_F, p_G) > \hat{D}_F(\hat{p}_F, \hat{p}_G)$, or $D_G(p_F, p_G) < \hat{D}_G(\hat{p}_F, \hat{p}_G)$. The former will again

contradict $\hat{p}_F \ge p_F$, whereas the latter will contradict $\hat{p}_G \ge p_G$.

The richness of the heterogeneity and the general equilibrium limits the extent to which the effect of ATT on the equilibrium outcome is characterized analytically. Nevertheless, we can summarize the results and their implications as follows:

- 1. Proposition 2 shows that at the individual firm level, ATT causes firms (advertisers) to substitute away from the Facebook network to the Google network, all else including ad prices equal.
- 2. Proposition 2 also shows that all else equal, including ad prices, ATT causes revenue loss for all firms but more so with higher Facebook dependency (i.e., higher θ).
- 3. Proposition 3 analyzes the general equilibrium effect: with ATT, the price of Facebook ads and their overall demand/quantity fall.
- 4. While the richness of the model limits the analytical results to those stated in Proposition 3, we can draw further implications.
 - (a) The second statement of Proposition 3 means that a significant proportion of, or possibly all, firms reduce their purchase of Facebook ads. This will likely imply that the equilibrium price of Google ads goes up after the ATT shock. To see this, suppose otherwise. Those firms that reduce their purchase of Facebook ads must increase their demands of Google ads, which follows from the submodularity of payoff function in (a_F, a_G) together with Assumption 4-(ii). As long as this effect is significant, the equilibrium price of Google ad will be higher.
 - (b) This last point also makes it plausible that the relative expenditure for Facebook ads to that for Google ads falls with the ATT shock. The substitution effect derived in Proposition 2 together with the market-wide effect obtained in Proposition 3 mean that unless the Google ad price goes up too high, the relative

proportion of the spending for Facebook ads relative to Google ads likely falls after the ATT shock.