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 ‖

December 2024

Abstract

In the backdrop of heightened data privacy concerns, governments and firms have introduced various privacy regulations. 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 estimate ATT's impact on advertising that relies on off-platform data for targeting and measurement through a within-firm analysis, comparing conversion-optimized Meta campaigns, which were affected by ATT, to click-optimized Meta campaigns, which were not. We find that conversion-optimized campaigns saw a 36% reduction in click-through rates after ATT. While firms responded by shifting spend from Meta to the Google ecosystem, firms with higher baseline Meta dependence nevertheless experienced a 37% decline in firm-wide revenue relative to firms with lower baseline Meta dependence. 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 Digital Economics Paris Seminar, Econometric Society Interdisciplinary Frontiers Economics and AI/ML Meeting, GMU Program on Economics & Privacy Empirical Research Workshop, IIOC, Marketing Strategy Meets Wall Street Conference, Toulouse Economics of Platforms Seminar 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

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

‖University 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, reaching $506 billion in total spending worldwide in 2021.[1] Its growth is driven by two factors: precise targeting using granular 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, regulations 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 increase advertising costs and harm firms by reducing their ability to target consumers effectively, potentially leading to welfare losses for both firms and customers.

In this paper, we quantify the economic costs of one such privacy policy – 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 (Wagner, 2021, 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, TikTok, and Pinterest are also impacted by the ATT policy change, but as Meta and Google are the dominant players in this industry, we focus on them. Second, how

[1] https://www.statista.com/statistics/237974/online-advertising-spending-worldwide/

[2] Unlike general brand advertising, which builds equity over time without direct response metrics (Borkovsky et al., 2017).

[3] Here 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.

did firms reallocate 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 spend and performance across Meta, Google, and TikTok for 1,221 firms. The other comes from Grips Intelligence,[4] a leading data analytics and market intelligence firm, providing transaction and revenue data for 773 firms.

To characterize the first-stage 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 which 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% relative reduction in click-through rates for conversion-optimized campaigns. Additionally, Meta's online advertising spend 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 characterize the downstream implications of this on firm revenues using the revenue dataset that encompasses a broader range of e-commerce firms. Given the degradation of Meta advertising performance, we measure the revenue effect by stratifying firms in terms of their pre-ATT reliance on Meta and conduct a difference-in-differences analysis.[5] We find that after ATT went into effect, more Meta-dependent firms saw a decrease in total orders

[4] https://gripsintelligence.com/

[5] We also consider a specification where firms are stratified according to the share of their revenue that comes from iOS users and find similar results. Both iOS and Meta reliance expose firms more to the ATT policy change, but, as these measures are not highly correlated, capture ATT exposure in different ways.

by 22.3% and overall revenue by 39.4% relative to less Meta-dependent firms. Furthermore, we find that this effect is disproportionately borne by the smaller firms within our sample.

Our results have several important policy and managerial implications. The large and negative impact on revenues indicates that opt-in privacy regulation has a significant economic cost for firms that rely on targeted advertising for revenue generation, and especially for smaller firms. In particular, the magnitude of the revenue reductions suggests that privacy regulation can threaten the viability of business models, such as those of smaller 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 of composition of firms that can succeed in the product market. Finally, while we do not directly observe Apple's advertising platform in our sample, our results also speak to ongoing antitrust concerns around the potentially anti-competitive impacts of ATT (CMA, 2022; Sokol & Zhu, 2021) by showing the impact on competing advertising platforms, especially Meta and Google. 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.

Related Literature Our research joins a growing literature studying the economics of online privacy (Acquisti et al., 2016; Goldfarb & Que, 2023; Miklós-Thal et al., 2024), and exploring the impact of privacy regulation specifically, such as the EU's Cookie Law in 2009 (Goldfarb & Tucker, 2011), the 2010 AdChoices program (Johnson et al., 2020), the iOS privacy nutrition labels (Bian et al., 2021), the EU's General Data Protection Regulation (GDPR) in 2018 (Aridor et al., 2023; Goldberg et al., 2024; Johnson, 2023), and the proposed EU ePrivacy Regulation (Miller & Skiera, 2023).

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 conversion optimized" 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. Our work extends and complements 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.

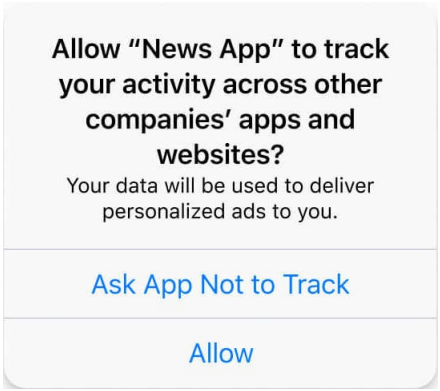
In contemporaneous work, Cecere and Lemaire (2023) also study the effect of ATT on predicted, aggregated ad outcomes and find that ATT reduced targeting efficiency on Meta. We complement this work by using platform-observed advertising data to similarly find a reduction in targeting efficiency and use our revenue data to quantify the downstream economic costs of reduced targeting efficiency. 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 Apple's App Tracking Transparency (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 as well as 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

Figure 1: ATT Data Sharing Prompt



by each app. This feature officially launched on April 25, 2021.[6] 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 to observe what those users did after leaving the app.[7]

The IDFA had two primary uses for mobile display advertising via platforms such as Meta. First, it provided a view of the consumer activity across applications, which could serve as an input for targeting. Second, it enabled Meta to link conversions to advertisements more easily.[8] If a consumer opts out through ATT, however, then 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.[9] Thus, both the loss in off-platform data and

[6]https://techcrunch.com/2020/06/22/apple-ios-14-ad-tracking/

[7]Unlike other privacy regulations 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 to remaining on the App Store was to include the prompt provided by Apple.

[8]Effective 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).

[9]Campaigns targeting non-impacted operating systems remained unchanged, but, if a campaign targeted

conversion measurement issues contribute to an overall degradation in targeting by reducing the data observed by firms (Johnson et al., 2022; Runge & Seufert, 2021).

2.2 Data Overview

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

The first 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. The second data source we denote as the *advertising* dataset, which comes from an anonymous advertising analytics provider, and provides granular data on Meta, Google, and TikTok advertising expenditures and performance for 1,221 firms. 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 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 only relies on first-party data to track relevant metrics. As a result, this dataset remains unaffected by ATT as long as firms do not alter the data they report through the platform. From the available domains, 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.[10]

We present firm-month level summary statistics for the revenue dataset during the pre-

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.

[10]Last-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.

Table 1: Dataset Summary Statistics

		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

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 "share" variables for the advertising data refer to the share of monthly average online advertising spend on Meta from the set of Meta, Google, and TikTok advertising.

ATT period (April 2020 to April 2021) in Table 1. The distribution of monthly revenue exhibits significant right skewness, with a median of $359,000 and a mean of $4.9 million. For the median domain, 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.[11]

2.2.2 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) as well as the total number of conversions associated with the advertising campaign (conversions). The measurement of

[11] 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.

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 this data at different levels of granularity. For Meta, we observe performance broken down based on 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, Display). We present a set of summary statistics for the advertising dataset in Table 1, indicating the mean online advertising spending is $115,390 per month and that the online advertising share across different platforms heavily skews towards Meta.

2.2.3 Data Representativeness

As both datasets contain firms that opt into data sharing, a natural question arises regarding which firms choose to participate and the broader population they represent. In Online Appendix D, 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 "population-level" 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 conclude that both datasets reasonably capture a broad class of e-commerce retailers, with the revenue dataset skewing toward relatively larger firms than the advertising dataset.

Given these datasets and their respective compositions, our analyses proceed as follows.

[12] See 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.

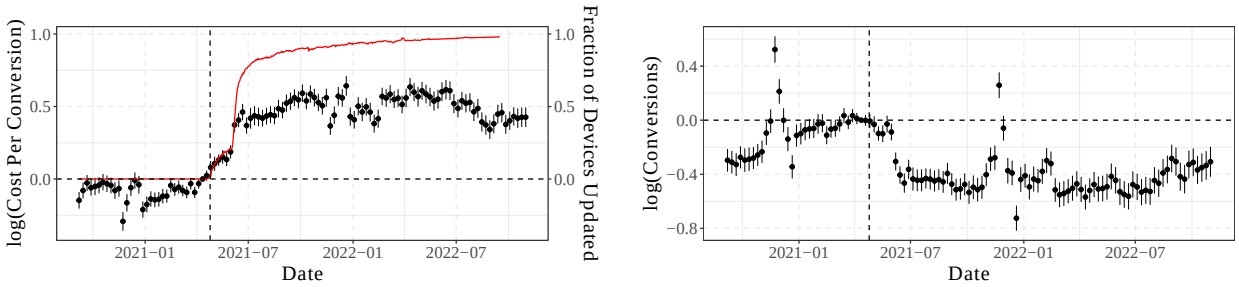
First, we use the advertising dataset to assess the impact of ATT on the efficacy of ad campaigns reliant on off-platform data through a within-firm comparison of off- versus on-platform ad campaign performance while controlling for unrelated factors (e.g., firm size and type). Next, we use the advertising 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-stage 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 First-Stage: Impact on Advertising Effectiveness

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

Figure 2: Event Study on Conversions and Cost per Conversion



(a) Event Study for Log(Conversion Cost)

(b) Event Study for Log(Conversions)

NOTES: Results use a balanced panel of firms using Meta advertising from the advertising dataset. The figures plot the event study coefficients for log(Cost Per Conversion) on the left and log(Conversions) on the right. We note that both of these variables have measurement issues after ATT. Table A.1 presents the 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

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.[13] We restrict attention to a balanced panel of firms with Meta advertising spending from September 2020 until October 2022 and estimate the following specification:

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

where α_i denotes the firm fixed effects.

Figure 2 plots the estimated β_t for each week and, beyond the spikes around the holiday season, it is clear that after ATT the number of conversions drops dramatically and the cost per conversion increases.[14] 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. The gradual adoption of iOS 14.5 coincides with a gradual increase in the cost per conversion that then nearly discontinuously increases by nearly 30% as Apple nudged a large portion of users to adopt iOS 14.5 in early June, resulting in an overall 73.2% increase 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. 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 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 since it enables 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

[13] 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.

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

to determine whether there were real degradations in targeting caused by the introduction of ATT, we next exploit the fact that ATT only impacts the ability to measure conversions and not events on Meta-owned platforms such as advertising clicks.

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 accurately measure after ATT, (2) they are the most popular campaign objective whose measurement is not impacted by ATT,[15] and (3) clicks are positively correlated with conversions.[16]

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.[17] We consider the following specification for firm i, advertising campaign objective j, and month t:

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

where T_j is an indicator for whether the campaign j is a conversion-optimized campaign, α_{ij} denotes firm-campaign fixed effects, and κ_t denotes month fixed effects.

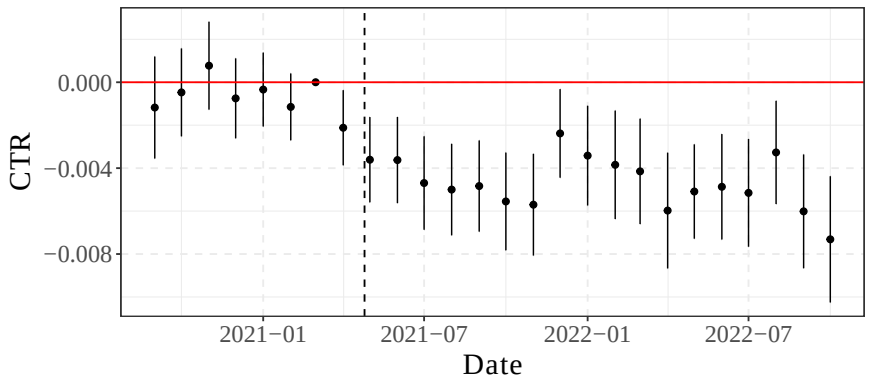
We estimate specification (2) using the click-through rate for firm i and campaign objective j for each month t as Y_{ijt} and on a balanced panel of firms that utilized both click-

[15] Table OA1 provides a breakdown of the market shares of different campaign objectives before ATT.

[16] Table 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.

[17] For instance, Meta adjusted from deterministic to probabilistic attribution models in an attempt to recover some of the lost targeting efficiency as a result of ATT.

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



NOTES: Results use a balanced panel of firms from the advertising dataset who use both click and conversion optimized campaigns before ATT. Standard errors are clustered at the firm level. The plots represent the time-varying treatment effects for the the click-through rate between conversion and click-optimized ads. Table A.3 presents the aggregate post-ATT estimates.

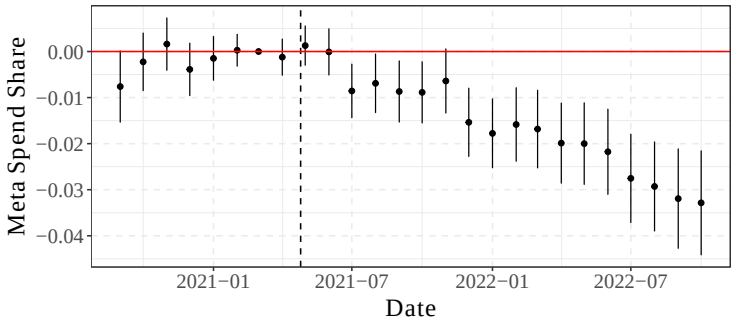
optimized and conversion-optimized campaigns pre-ATT.[18] Figure 3 shows identical performance between campaign types before ATT, followed by a sharp decline in click-through rates for conversion-optimized campaigns after ATT's introduction. The relative reduction in click-through rates is 0.004 for conversion-optimized campaigns, representing a 36.6% decrease from the baseline rate of 0.011. While we cannot directly measure the impact of ATT on conversions due to measurement limitations, Table A.2 shows that the pre-ATT correlation between clicks and conversions was strong for conversion-optimized campaigns, suggesting a proportional 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.

[18]This sample selection criterion results in a subset comprising 44.9% of the firms included in Figure 2. Table A.4 shows that the event studies report a similar impact on conversion-optimized advertisements for this subset of firms.

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 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 the Online Appendix B.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 spend to Google.

Figure 4: Event Study of Online Ad Spend Share (Meta)



NOTES: Results use a balanced panel of firms with non-zero online advertising spend in the advertising dataset. The plot represent the event study estimates for Meta spend share. Table A.3 presents the aggregate post-ATT estimates. 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 to 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 E 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 spend share on Meta, calculated as their spend on Meta divided by their total spend on Meta, Google, and TikTok during the same period. We then estimate specification (1). Figure 4 presents the estimates which show little change in market share before the onset of ATT and a gradual decrease in the share of Meta after ATT.[19] The mean market share for Meta ads was 0.75 in the baseline period but fell by 4.4% by the end of our sample period. In Online Appendix B.2, we use an across-firm difference-in-differences analysis to show that this reallocation was more pronounced for firms with higher pre-ATT Meta dependence.

These results indicate a meaningful reallocation of ad spending, suggesting that to measure 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. In an ideal world, there would be randomized variation in advertising firms' exposure to ATT, which did not occur. Our approach to understanding its effects, therefore, relies on variation across firms in the extent to which they were vulnerable to being impacted by ATT based on their pre-ATT characteristics.

Empirical Strategy: Our empirical strategy employs an across-firm difference-in-differences approach to compare pre- and post-ATT revenue for firms differing in their vulnerability to ATT. Our treatment indicator, capturing vulnerability to ATT, is derived from each firm's pre-ATT reliance on Meta as a source of revenue. Specifically, we first calculate for each firm the average share of revenue coming from Facebook/Instagram sessions over the one-year

[19] In Online Appendix B.2 we show the time series of log(Spend) across these two advertising platforms and show that a similar pattern holds for two relevant quantity variables: clicks and impressions.

period before ATT's introduction (April 2020 to April 2021). We then use a median split of this Meta revenue share to classify firms into 'high exposure' (treatment) and 'low exposure' (control) groups. On average, firms in the treatment group attribute 8.17% of their pre-ATT revenue to Meta traffic, compared to 0.42% for the control group.

In addition, as ATT only applies to iOS users and not Android users, we also consider a specification defining ATT vulnerability based on a median split of firm revenue attributable to iOS users, following the same procedure.[20] This captures a slightly different, yet complementary, dimension of dependence on pre-ATT targeted advertising relative to our baseline strategy.[21] While neither dependence on Meta or iOS users is an exogenous firm characteristic, we again rely on the parallel trends assumption and test for the presence of different revenue trends between treated and control firms in the pre-period. Formally, we estimate:

$$Y_{it} = \sum_t \beta_t \Big(\text{Month}_t \times T_i \Big) + \alpha_i + \kappa_t + \epsilon_{it}, \tag{3}$$

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.

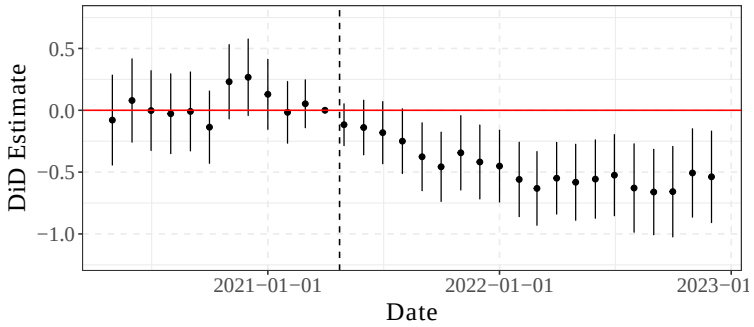
Overall Revenue Changes: Results for the main specification are shown in Table 2, with the corresponding time-varying treatment effects in Figure 5. We find that in the pre-period, there are no significant differences between the treated and untreated firms in terms of log monthly revenue. In the post-period, beginning in month 1 we see a clear downward trend in revenue for the firms most exposed to ATT, with significant differences beginning within 4 months. This is again consistent with the gradual timing of adoption of iOS 14.5

[20] Using this definition of exposure, firms in the treatment group derive 36.51% of their pre-ATT revenue from iOS traffic, versus 11.84% for firms in the control group.

[21] While 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 treated under the two definitions. 29% of firms are considered treated under both definitions, 42% are considered treated under only one, and 29% are considered treated under neither.

amongst 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.

Figure 5: Time-Varying Treatment Effects for Revenue (Meta Share Treatment)



Notes: Results use the revenue dataset. The estimates present the time-varying treatment effects for log(Total Revenue). The treatment indicator is based on a dummy indicating the firm-level pre-ATT share of revenue from Meta traffic is above the median. Standard errors are clustered at the firm-level.

The coefficient estimate in column (1) of Table 2 suggests a decrease in revenue of 37.1% for more Meta-dependent firms relative to less Meta-dependent firms. Column (2) considers an alternative specification that allows for separate time trends for each e-commerce category, suggesting a relative revenue decline of 32.7%. 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. Columns (4) and (5) show that this effect is driven by small firms, which are defined as those with below-median pre-ATT average monthly revenue.[22] Finally, in column (6) we study the impact on the log number of transactions, finding that they decline by 21.4%. In Table A.5 and Figure A.1 we present results using the alternative treatment indicator based on the pre-ATT share of revenue from iOS users and find consistent results, suggesting that the estimated revenue effects are robust to how we categorize firms as vulnerable to ATT.[23]

[22]We plot time-varying treatment effects separately for small and large firms in Figures A.3a and A.3b, respectively, and for transactions in Figure A.2.

[23]In Online Appendix C.1, we also present data plots and robustness checks using tercile and quartile splits instead of median splits. We find consistent results across these specifications, with a greater reliance on iOS users and Meta traffic being associated with larger post-ATT decreases in revenue.

Table 2: Primary Revenue Estimates (Meta Treatment)

	Dependent Variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
	log(Revenue)					log(Transactions)
	All Firms			Small Firms	Large Firms	All Firms
$\text{After}_t \times \text{Treated}$	-0.463***	-0.396**		-1.132***	0.158	-0.241**
	(0.165)	(0.180)		(0.302)	(0.121)	(0.100)
$0-3$ months $\times$ Treated			-0.138			
			(0.151)			
4+ months $\times$ Treated			-0.499***			
			(0.171)			
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Category $\times$ Month FE		Yes				
N	24868	24868	24868	12468	12400	24868
R^2	0.73	0.74	0.73	0.58	0.67	0.83
Marginal Effects (%)	-37.06%	-32.70%	–	-67.76%	17.12%	-21.42%
Share Treated (%)	8.17%	8.17%	8.17%	10.22%	5.96%	8.17%
Share Not Treated (%)	0.42%	0.42%	0.42%	0.34%	0.49%	0.42%

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

NOTES: Results use the revenue dataset and the treatment indicator that relies on the share of revenue from Meta traffic. The first 5 columns use log(revenue) as the dependent variable and column 6 uses log(# transactions). The rows present the estimated average treatment effect coefficient using the difference-in-differences specification. Standard errors are clustered at the firm level.

These findings indicate that privacy policies have substantially harmed some e-commerce firms, with effect estimates larger than what one might expect from the direct revenue share attributed to Meta advertising. Several factors likely contribute to these large effects. First, 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. Second, these losses represent foregone growth and not absolute revenue declines: revenue growth at more Meta-reliant firms lags behind that of less Meta-reliant firms, but those less reliant firms continue to grow. Third, 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.[24] Fourth, because firms previously relied on off-platform data to refine ad targeting across all platforms, the loss of this data under ATT decreases targeting effectiveness for every user segment, not just iOS. Finally, we readily acknowledge that many factors beyond ATT – some unknown and unknowable – also influence revenue outcomes, contributing to uncertainty in our estimates.

5 Conclusion

As policymakers consider extending or implementing new privacy regulations that limit the ability of firms 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 more Meta-dependent firms experienced a 37.1% relative reduction in revenue that was primarily borne by small firms.

There are several policy and managerial takeaways from this paper. Our estimates suggest large economic costs of opt-in privacy regulation. 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 deter entry of these firms into product markets. These findings highlight the importance of developing balanced approaches to privacy regulation that protect consumer rights while also supporting the competitiveness of small businesses in the digital economy.

[24] We provide suggestive evidence that the negative revenue impacts are largely from reduced new customer acquisition, but also find negative point estimates on repeat customers, in analysis of an auxiliary dataset in Online Appendix C.2.

References

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.

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*.

CMA. (2022). *Mobile ecosystems market study*. https://www.gov.uk/cma-cases/mobile-ecosystems-market-study

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*.

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.

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/

Li, D., & Tsai, H.-T. (2022). Mobile apps and targeted advertising: Competitive effects of data exchange. *Available at SSRN 4088166*.

Miklós-Thal, J., Goldfarb, A., Haviv, A., & Tucker, C. (2024). Frontiers: Digital hermits. *Marketing Science (Articles in Advance)*, *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.

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.

Wagner, K. (2021). Facebook users said no to tracking. Now advertisers are panicking. *Bloomberg*. https://www.bloomberg.com/news/articles/2021-07-14/facebook-fb-advertisers-impacted-by-apple-aapl-privacy-ios-14-changes

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

Appendix A Omitted Tables and Figures

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

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

*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 implementation. The dependent variables are the click-through rate, and the treatment is whether the campaign is a conversion-optimized campaign. Standard errors are clustered at the firm level.

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

	Dependent variable:	
	(1)	(2)
	$\log(1 + \text{Conversions}_{ijt})$	
$\log(1 + \text{Clicks}_{ijt})$	0.252^{***}	0.195^{***}
	(0.017)	(0.013)
$\log(1 + \text{Clicks}_{ijt}) \times \mathbb{1}(\text{j is Conversion-Optimized Campaign})$	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<0.01

NOTES: All results use the advertising dataset and we use a balanced panel of Meta firms who spend on click-optimized and conversion-optimized campaigns. We estimate the following specification on data from the pre-ATT period: $\log(1 + \text{conversions}_{ijt}) = \beta\Big(\log(1 + \text{clicks}_{ijt}) \times \mathbb{1}(\text{j is Conversion-Optimized Campaign})\Big) + \alpha_{ij} + \kappa_t + \epsilon_{ijt}$. Standard errors are clustered at the firm level.

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

	Dependent variable:		
	(1)	(2)	(3)
	log(Cost Per Conversion)	log(Conversions)	Meta Online Spend Share
After_t	0.549^{***}	-0.302^{***}	-0.014^{***}
	(0.019)	(0.032)	(0.003)
Firm FE	Yes	Yes	Yes
Observations	61,091	61,091	31,746
R^2	0.837	0.855	0.931
Marginal Effects	73.15%	-26.07%	-

*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 implementation. The dependent variables are the log(Cost Per Conversion), the log(Conversions), and the spend share of Meta ads. Columns (1) and (2) are estimated over the sample of advertisers with a balanced panel of Meta advertising spend, while column (3) is estimated over the sample of advertisers with a balance panel of any online advertising spend. We note that the dependent variables in columns (1) and (2) have measurement issues after ATT. Standard errors are clustered at the firm level.

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

	Dependent variable:			
	(1)	(2)	(3)	(4)
	log(Conversions)		log(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
R^2	0.839	0.861	0.819	0.848
Marginal Effects	−30.30%	−22.43%	74.89%	71.42%

*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. Standard errors are clustered at the firm level.

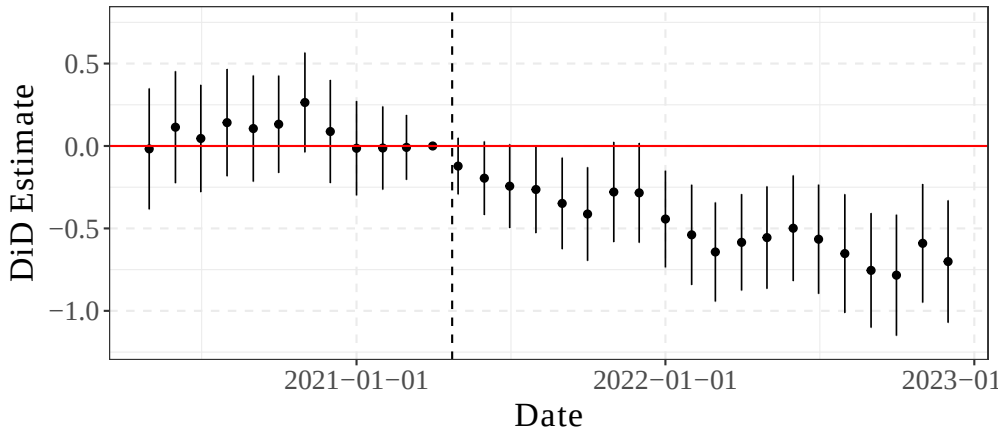
Table A.5: Alternative Revenue Estimates (iOS Treatment)

	Dependent Variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
	log(Revenue)					log(Transactions)
	All Firms			Small Firms	Large Firms	All Firms
$\text{After}_t \times$ Treated	-0.522***	-0.352**		-0.999***	0.079	-0.237**
	(0.165)	(0.167)		(0.290)	(0.116)	(0.100)
$0-3$ months $\times$ Treated			0.040			
			(0.103)			
4+ months $\times$ Treated			-0.483***			
			(0.159)			
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Category $\times$ Month FE		Yes				
N	24868	24868	24868	12468	12400	24868
R^2	0.73	0.74	0.73	0.58	0.67	0.83
Marginal Effects (%)	-40.67%	-29.67%	–	-63.18%	8.22%	-21.10%
Share Treated (%)	36.51%	36.51%	36.51%	35.10%	37.73%	36.51%
Share Not Treated (%)	11.84%	11.84%	11.84%	11.35%	12.42%	11.84%

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

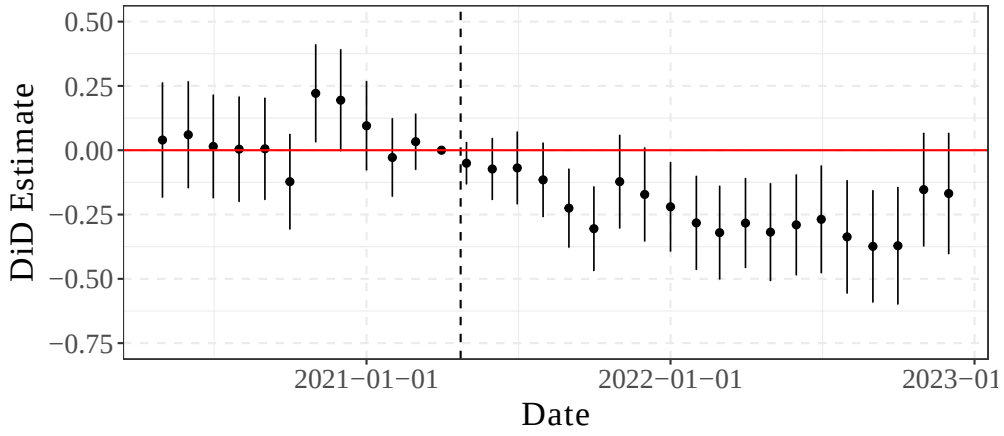
NOTES: Results use the revenue dataset and the treatment indicator that relies on the share of revenue from iOS traffic. The first 5 columns use log(revenue) as the dependent variable and column 6 uses log(# transactions). The rows present the estimated average treatment effect coefficient using the difference-in-differences specification. Standard errors are clustered at the firm level.

Figure A.1: Time-Varying Estimates for Revenue (iOS Share Treatment)



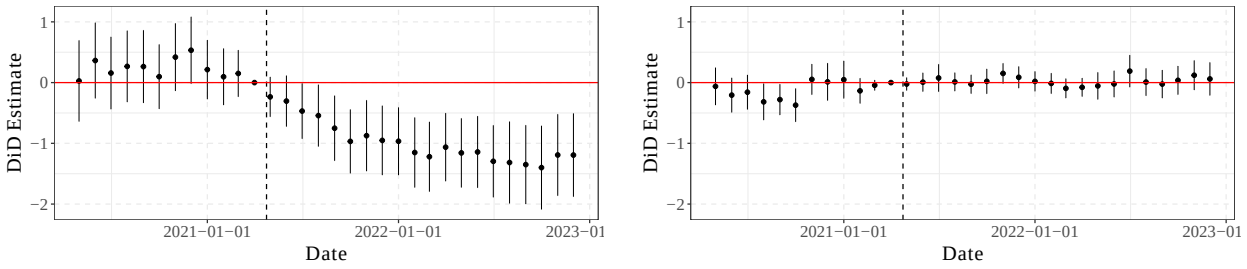
NOTES: Results use the revenue dataset. The estimates present the time-varying treatment effects for log(Total Revenue). The treatment indicator is based on a dummy indicating the firm-level pre-ATT share of revenue from iOS traffic is above the median.

Figure A.2: Time-Varying Estimates for Transactions (Meta Share Treatment)



NOTES: Results use the revenue dataset. The estimates present the time-varying treatment effects for log(Transactions). The treatment indicator is based on a dummy indicating the firm-level pre-ATT share of revenue from Meta traffic is above the median.

Figure A.3: Time-Varying Estimates for Small and Large Firms (Meta Share Treatment)



(a) Small Firms

(b) Large Firms

NOTES: Results use the revenue dataset. The estimates present the time-varying treatment effects for log(Revenue) for firms whose pre-ATT revenue was below the median (Panel a and above the median (Panel b). The treatment indicator is based on a dummy indicating 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.

Online Appendix: For Online Publication Only

Appendix B Additional Analyses for Advertising

B.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 firm's 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, 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

[1] In 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.

[2] This includes all campaign objectives except for messages and event responses since it is ambiguous how to categorize them.

Table OA1: Spend Share of Meta Advertising Campaign Objectives

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%
Video Views	On-platform Reach	0.55%
Post Engagement	On-platform Actions	0.15%
Store Visits	On-platform Actions	0.02%
Page Likes	On-platform Actions	0.02%

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.

the vast majority of total spending. Second, link clicks are the most popular 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.

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 off-platform 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 from off-platform conversion-optimized campaigns to on-platform action-optimized campaigns.

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

We consider three outcome measures: the campaign objectives' share of spending, share of impressions, and an indicator for whether spend was non-zero. The former two 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.

Table OA3: Meta Campaign Objective Substitution

	Dependent variable:		
	(1)	(2)	(3)
	Spend Share	Impression Share	$\mathbb{1}(\text{Spend}_t > 0)$
$\text{After}_t \times$ On-Platform Actions	−0.007	0.004	0.015
	(0.007)	(0.009)	(0.011)
$\text{After}_t \times$ 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
R^2	0.839	0.750	0.535

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

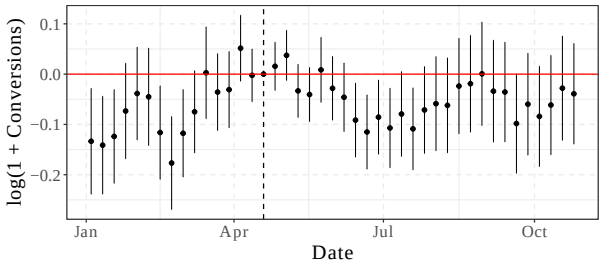
NOTES: All results use the advertising dataset. We estimate this on a balanced panel of firms who have non-zero spend on Meta. The dependent variables are the share of spend (column 1), share of impressions (column 2), and whether there is non-zero spend (column 3). The left out category is off-platform 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 use 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 which, consistent with the aggregate spending in Table OA2, show a precise null effect on substitution to on-platform objectives on the intensive-margin. 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 calculated using a balanced panel of firms who use both click- and conversion-optimized objectives before ATT.

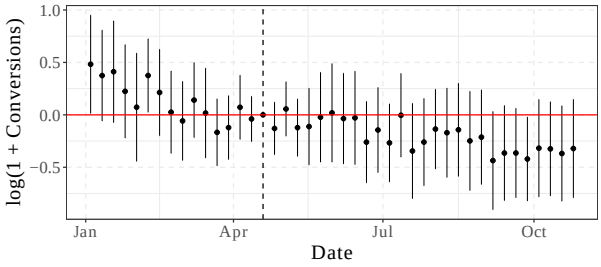
B.2 Additional Analyses for Advertising Platform Substitution

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 due to ATT. Then, we explore absolute trends in spending patterns between Meta and Google to show that they follow similar patterns after ATT, but noticeably diverge after ATT. Finally, we explore whether reallocation was more likely by more Meta-dependent firms.

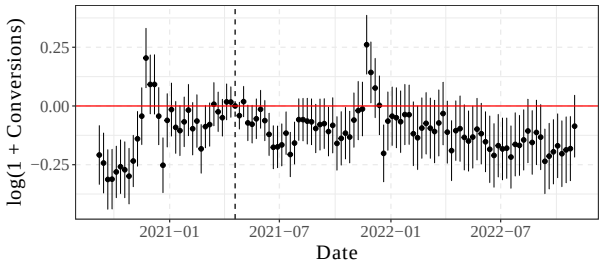
Figure OA1: Event Study for Google Conversions



(a) Event Study for Google Search Conversions



(b) Event Study for Google Display Conversions



(c) Event Study for Overall Google Conversions

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.

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 estimates for the relative changes in performance of Google advertising. To do so, we esti-

mate 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 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.

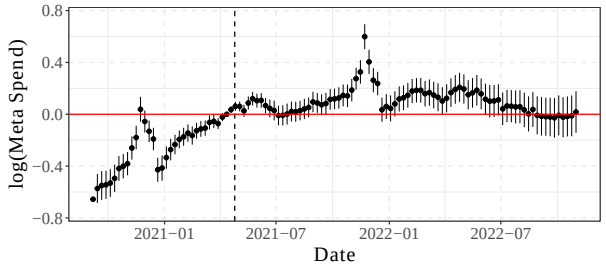
Changes in Total Advertising Spending: Given that we have some evidence that Google advertising is less impacted than Meta advertising, we explore whether and to what extent firms changed their online advertising shares towards Google as opposed to Meta. To do so, we focus mainly on changes in total spending. We consider 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 only allows us to observe end outcomes, so that we can only observe the average price of the ads that firms actually purchase, and not the prices themselves.

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 plot the estimates in Figure OA2. Figure OA2 shows that spending is increasing both on 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 OA1. We cannot disentangle whether this comes from supply-side or demand-side effects, but additionally show in Figure OA3 that we get similar relative reductions in

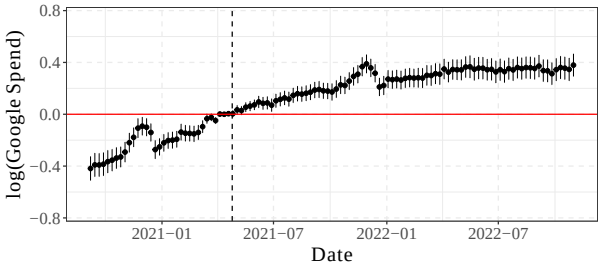
[4]For the individual services, we only report results from January 1st until October 31st 2021 since Google launched their Performance Max product in November 2021, which led to substitution within the Google Ecosystem that is orthogonal to ATT.

quantity of Meta ads using either clicks or impressions as our quantity measure.

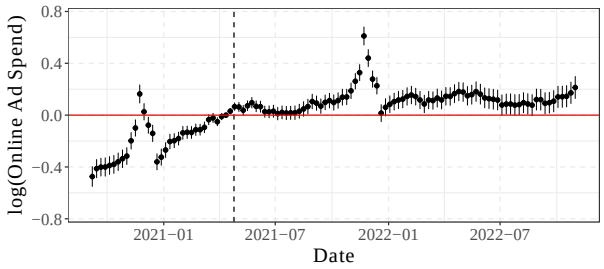
Figure OA2: Event Study for Online Advertising Spending



(a) Event Study for Meta Spend



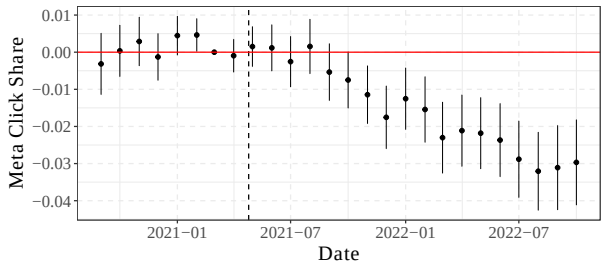
(b) Event Study for Google Spend



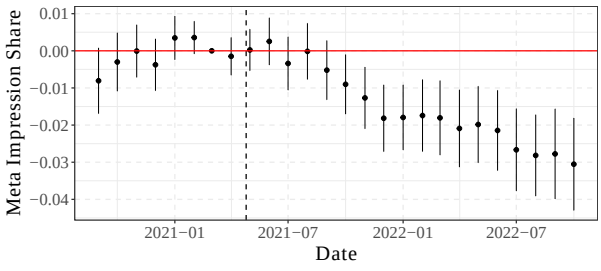
(c) Event Study for Total Advertising Spend

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) considers 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.

Figure OA3: Event Study on Meta Online Ad Share of Clicks and Impressions



(a) Event Study for Online Ad Click Share



(b) Event Study for Online Ad Impression Share

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 monthly level. Panels (a) and (b) consider the dependent variable of the share of observed clicks and impressions, respectively, attributed to Meta.

Table OA4: Advertising Platform Share Estimates

	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
	(0.008)	(0.008)	(0.008)
Google Display	0.005***	0.028***	0.013***
	(0.002)	(0.006)	(0.004)

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

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

Which firms are more likely to reallocate? We now explore whether more Meta-dependent firms 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 above-average Meta advertising spend as a proportion of total advertising spend in the advertising dataset.[5] We consider a balanced panel of firms who spend non-zero dollars on any advertising platform throughout the same sample period as before and our outcome variables are the online advertising market share of impressions, clicks, and spending across 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.

[5]The mean is reported in Table 1 as 0.75.

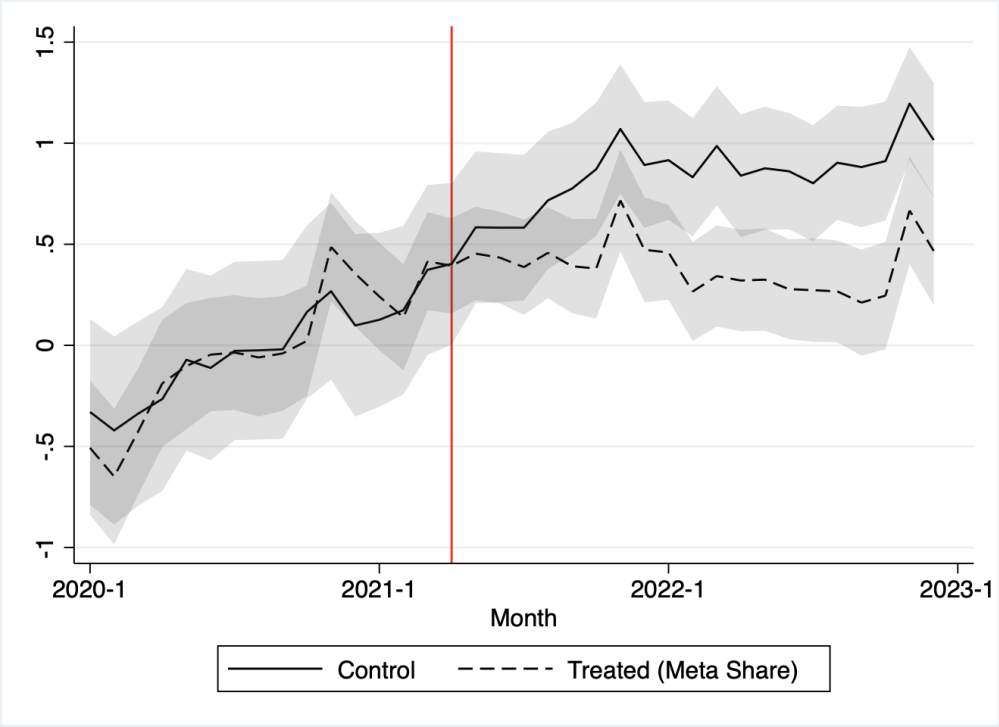
Furthermore, rows (4) and (5) of Table OA4 show the change in market share across different Google products and there is a greater increase in the share of Google Display relative to Google Search.[6]

Appendix C Additional Analyses for Revenue

C.1 Additional Analyses of Revenue Dataset

In this section we consider several robustness exercises to complement the main analyses done in Section 4.

Figure OA1: Log Revenue for Low and High Meta Shares

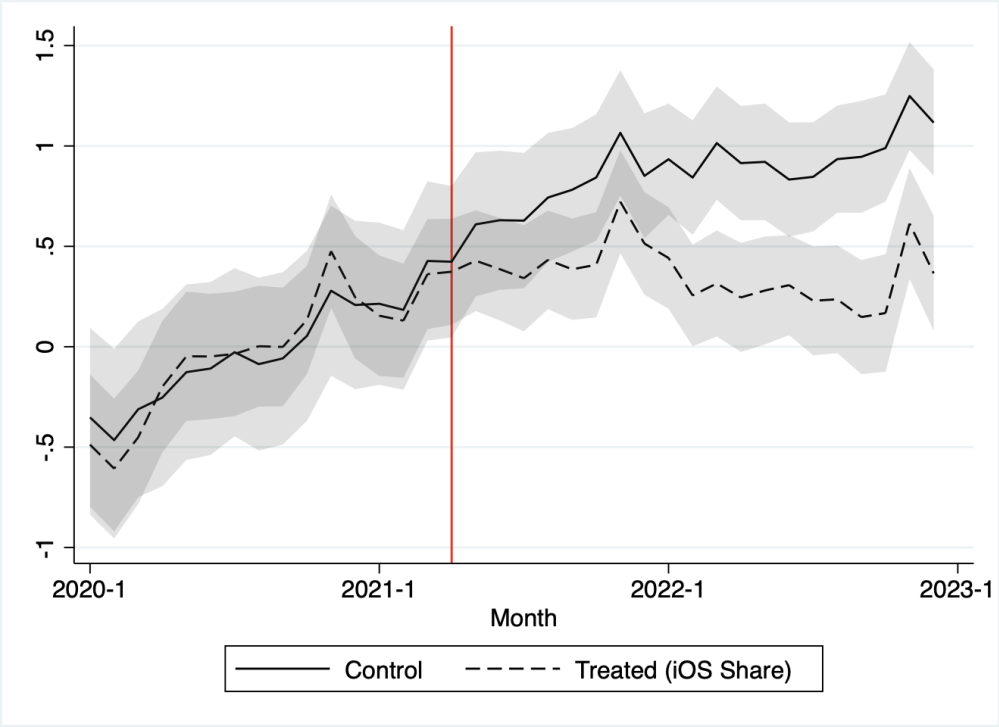


NOTES: Results using the revenue dataset. High and low Meta shares are calculated using a median split of pre-ATT revenue from Meta traffic. Plot shows log(revenue) demeaned using the pre-ATT mean along with 95% confidence intervals.

In Figures OA1 and OA2 we plot the demeaned log monthly revenue for the treated and

[6]The results for rows (4) and (5) are estimated over 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.

Figure OA2: 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.

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.

As an additional robustness check, we consider more conservative versions of the median split. We present the results in Table OA1, where we show the median split results alongside results using the top/bottom tercile and top/bottom quartile for both treatment definitions. We find consistent results across these specifications, with a greater reliance on iOS users and Meta traffic being associated with larger post-ATT falls in revenue.

Table OA1: Revenue Estimates Robustness: Alternative Splits

	Dependent variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
	Meta Treatment			iOS Treatment		
	Median Split	Terciles	Quartiles	Median Split	Terciles	Quartiles
After×Treated	-0.463***	-0.612***	-0.894***	-0.522***	-0.617***	-0.804***
	(0.165)	(0.233)	(0.301)	(0.165)	(0.228)	(0.291)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
N	24868	16570	12434	24868	16580	12446
R^2	0.73	0.71	0.69	0.73	0.71	0.71
Marginal Effects (%)	-37.06%	-45.77%	-59.10%	-40.67%	-46.04%	-55.25%

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

NOTES: Results use the revenue dataset and in all columns the dependent variable is log of total revenue. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification for both the Meta treatment definition and the iOS treatment definition. Median split results correspond to the main results in the paper. Tercile/quartile split results stratify firms based on whether they are in the top/bottom quartile or tercile, with the middle tercile/quartiles dropped. Standard errors are clustered at the firm level.

C.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, this data provides 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 has no measurement issues as a result of ATT. Notably, the measurement of repeat customers relies on data unaffected by the changes from ATT since they are typically user-provided email addresses or phone numbers. 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 analyses. 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 primarily contains smaller firms. As such, the estimates in this section will be relatively imprecisely estimated and underestimates of 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 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 OA2 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 OA2. While this reports similar effect sizes for the ratio of repeat customers, we no longer find statistically significant reductions in revenues or orders though we still find sizeable and negative point estimates. For the repeated order fraction, 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 this measure 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 using the log of new and repeat orders are presented in columns

(2) and (3) of Table OA3. Column (2) shows a statistically significant 28.5% decrease in orders coming from new customers, but column (3) shows a negative, but statistically insignificant, effect on repeat customer orders. The second row of Table OA3 shows that this result is robust to controlling for total online advertising spend. 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 OA3 considers the time-varying difference-in-differences 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 OA4 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 OA4 indicate that there is moderate overdispersion in the data, supporting the suitability of the negative binomial regression model.

Table OA2: Estimates on Sales

	Dependent variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
	log(Orders)		log(Revenue)		Repeated Order Ratio	
$\text{After}_t \times \text{Treated}$	-0.242^{*}	-0.171	-0.230^{*}	-0.156	3.994^{***}	3.563^{**}
	(0.129)	(0.118)	(0.131)	(0.119)	(1.415)	(1.419)
Ad Spending Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,772	5,772	5,772	5,772	5,772	5,772
R^2	0.854	0.872	0.847	0.867	0.808	0.813

*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 is observed. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification with and without controls for log(total advertising spending + 1). Standard errors are clustered at the firm level.

Table OA3: Estimates on New vs. Repeat Customers

	Dependent variable:			
	(1)	(2)	(3)	(4)
	log(New Customer Orders + 1)		log(Repeat Customer Orders + 1)	
$\text{After}_t \times \text{Treated}$	-0.337^{***}	-0.257^{**}	-0.154	-0.097
	(0.128)	(0.116)	(0.141)	(0.132)
Ad Spending Controls	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	5,772	5,722	5,722	5,722
R^2	0.829	0.852	0.877	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 is observed. The columns are log(1 + orders) coming from new and repeat customers. The rows present the estimated average treatment effect coefficient using the difference-in-differences specification, where odd columns do not control for log(1 + total advertising spending) and even columns do control for them. Standard errors are clustered at the firm level.

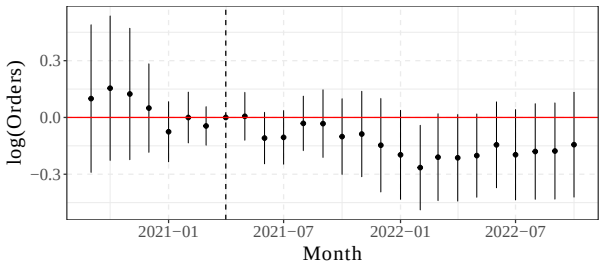
Table OA4: Estimates on New vs. Repeat Customers (Negative Binomial Specification)

	Dependent variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
	Total Orders		Orders from New Customers		Orders from Repeat Customers	
$\text{After}_t \times \text{Treated}$	-0.165^{***}	-0.134^{***}	-0.243^{***}	-0.208^{***}	-0.088^{**}	-0.061^{*}
	(0.033)	(0.035)	(0.034)	(0.032)	(0.034)	(0.033)
Ad Spending Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,772	5,772	5,772	5,772	5,772	5,772
θ	3.262^{***} (0.059)	3.597^{***} (0.066)	2.814^{***} (0.051)	3.142^{***} (0.057)	3.231^{***} (0.061)	3.449^{***} (0.066)

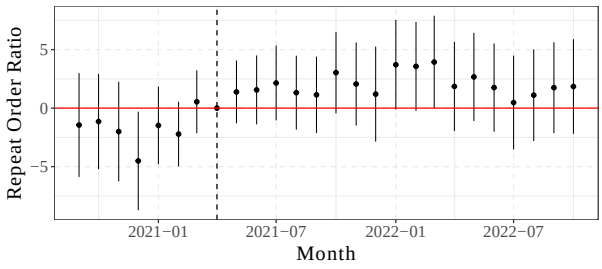
*p<0.1; **p<0.05; ***p<0.01

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

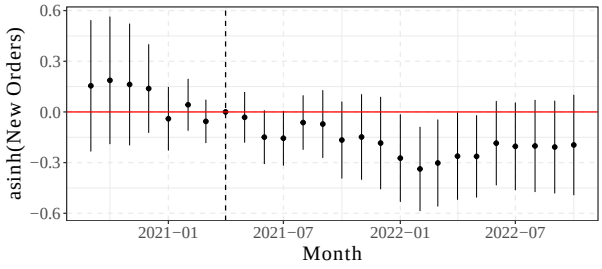
Figure OA3: Difference-in-Difference Estimates for Sales Outcomes



(a) DiD Estimates for Orders



(b) DiD Estimates for Repeated Order Ratio



(c) DiD Estimates for New Orders

NOTES: Results use a balanced panel of firms within the advertising dataset for which Shopify transaction data is observed. The dependent variables on the top row from left to right: log(Orders), Repeated Customer Order Ratio. The outcome variable on the bottom row is log(New Customer Orders + 1). Plots are the time-varying estimates for the specification including controls for the log(Total Advertising Spend + 1). Standard errors are clustered at the firm level.

Appendix D 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 is available to be analyzed provides empirical support for the representativeness of the two datasets vis a vis a broad cross-section of e-commerce firms along key dimensions.

D.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 who 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.

D.1.1 Selection into Sample

The analytics firm that provided us with the advertising data specializes in benchmarking key business metrics across industries using aggregated and anonymized data sourced directly from participating e-commerce retailers. This analytics firm has a "give to get" business model in which firms who 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 is 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 benchmarking 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.

D.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.[7] 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 this 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

[7] https://www.vivvix.com/home

firm-side self-selection into the dataset, which was the main identification threat associated with our advertising dataset.[8]

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, used by firms such as Procter and Gamble, Unilever, and Google. It has more comprehensive digital coverage than its primary competitor, Nielsen Ad Intel. As such, while it may not have complete coverage of advertising spending, it offers more complete coverage than available alternatives, to the best of our knowledge.

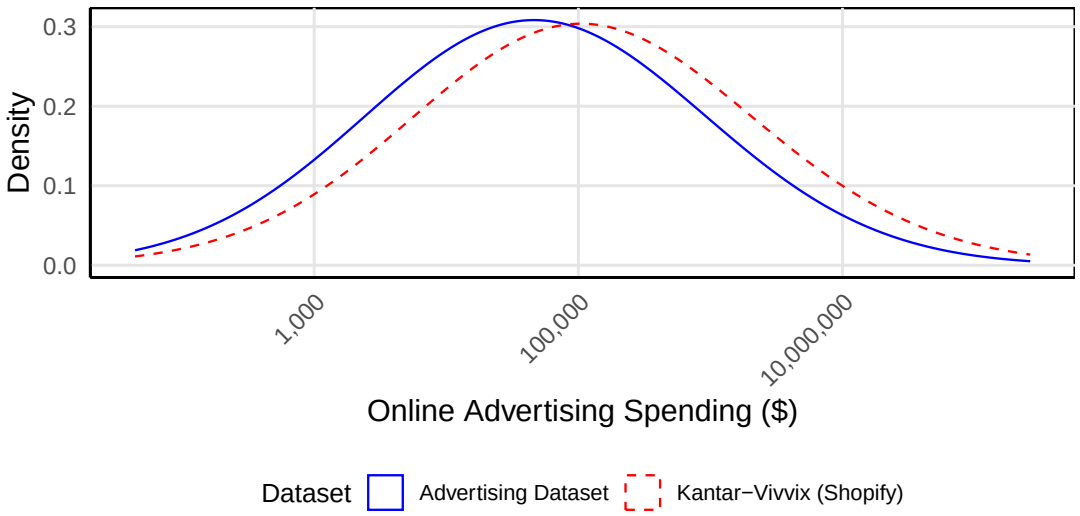
While the Kantar-Vivvix dataset has wide coverage, it provides us with a cruder measure of firm behavior than our advertising dataset, which contains richer data about their online advertising campaigns, including the breakdown by platform, campaign type, and direct measures of conversion rates as reported 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 e-commerce, such as retail, consumer packaged goods, automotive, and financial services, which are not relevant to our analysis. Therefore, we collect data from BuiltWith[9] 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, which identifies 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 e-commerce

[8]We 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.

[9]https://builtwith.com

Figure OA1: Advertising Data vs. Kantar-Vivvix (Shopify) Online Ad Spending Distribution



Dataset	Mean	Median	SD	25th	75th	N
Advertising Data	10.79	10.71	1.89	9.52	12.10	1475
Kantar-Vivvix (Shopify)	11.63	11.55	2.03	10.39	12.89	3248

NOTES: This figure and table compare the kernel density estimate and summary statistics of online advertising spending between 01/01/2021 and 03/31/2021 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.

retailer category.[10]

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 only collects mobile ad spending starting at the beginning of 2021, and April 2021 is when ATT occurs. We only include firms with advertising spending in each week during this period. This allows us to assess representativeness of the advertising data to the e-commerce retailer category in terms of total advertising spend.

[10] As the Kantar-Vivvix data does not provide a domain name, we match firms within Kantar-Vivvix and firms in the advertising dataset manually by company name.

We compare the distribution and summary statistics of the log of total online advertising spending in Figure OA1. Overall, these results indicate that the two distributions are broadly similar, although the advertising dataset skews slightly toward relatively smaller firms, 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.

D.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 comes from Shopify, a widely used e-commerce platform, as reported by Shopify in its public disclosures. 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.

D.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 get access to 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 over 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.

D.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 given 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.[11] 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.

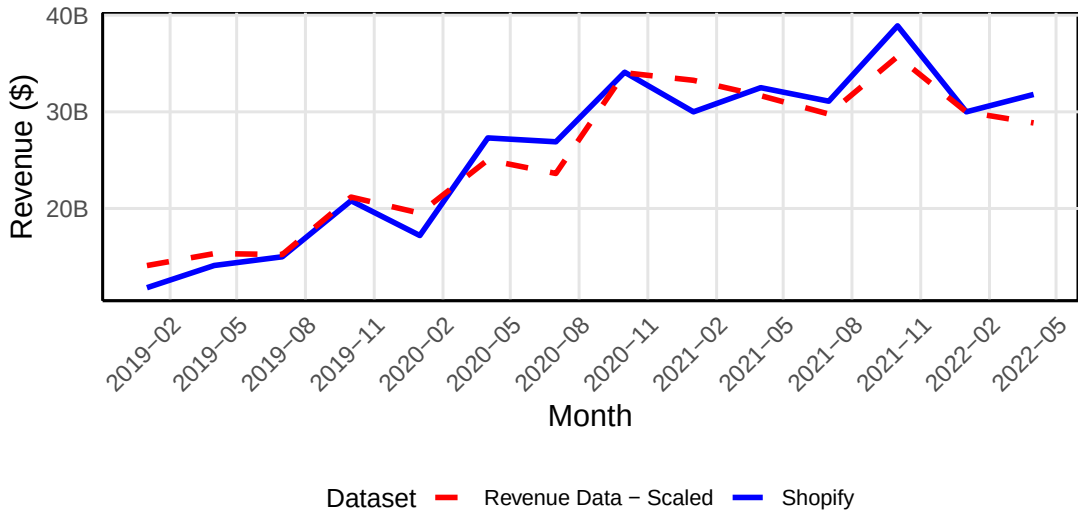
The financial data disclosed by Shopify in the aforementioned Investor Day presentation provides a population-level view of revenue generated by Shopify merchants, making the resulting analysis a natural complement to the representativeness analysis performed above on the advertising dataset, as both 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 those merchants began selling on Shopify. This cohort-based segmentation is particularly valuable for constructing an "apples-to-apples" comparison with the revenue dataset, which includes only merchants who began operating before the beginning

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

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 OA2: 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.

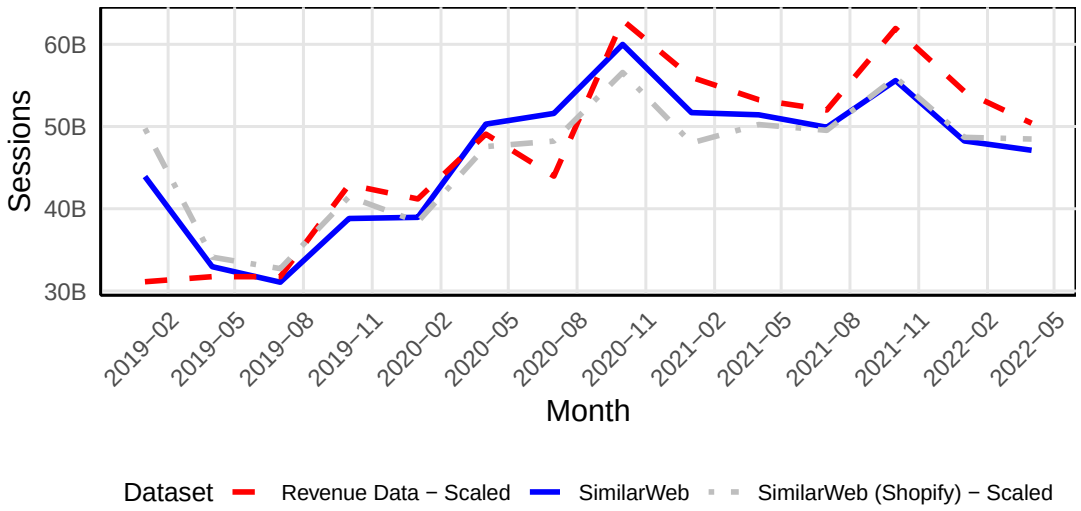
Figure OA2 shows the resulting comparison of Shopify revenue to our revenue dataset over time, after scaling the latter so that it has the same empirical mean as the former to facilitate visual comparability. The two resulting time series show consistent trends over time, supporting the notion that spending 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 domain-specific 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 firm-side selection bias, making it an appropriate benchmark.

For 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 OA3: 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.

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.

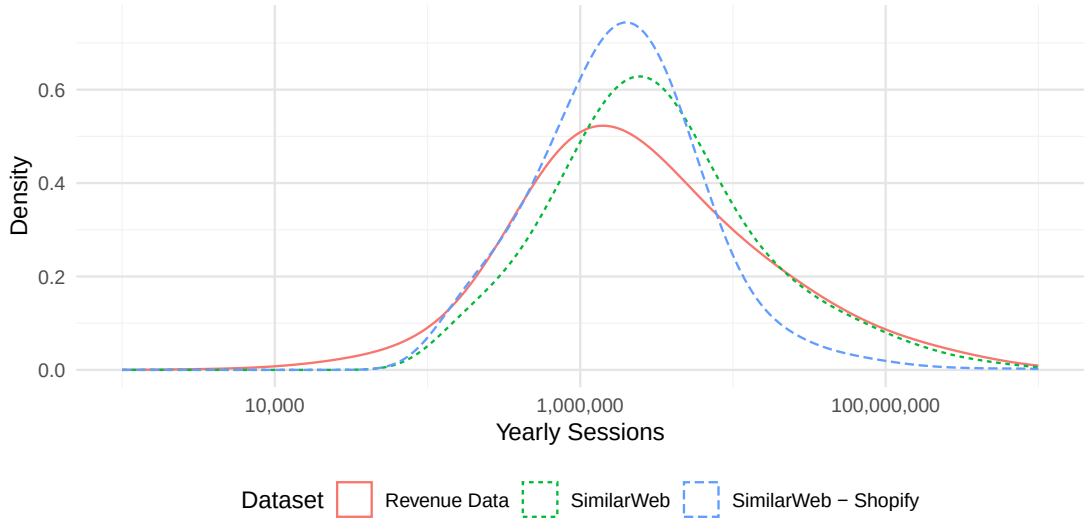
The results are shown in Figure OA3. We again see strong correspondence between trends in session counts over time between the revenue and both the SimilarWeb and SimilarWeb-Shopify datasets. The empirical correlations between the revenue and SimilarWeb/SimilarWeb-Shopify time series are 87.2% and 79.4%, respectively.

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

D.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 OA4, 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) the the Shopify-only subsample (SD = 1.30).

Figure OA4: Revenue Dataset vs. SimilarWeb Traffic Distribution and Summary Statistics



Dataset	Mean	Median	SD	25th	75th	N
Revenue Data	14.71	14.49	1.85	13.54	15.75	773
SimilarWeb	14.99	14.83	1.63	13.92	15.87	7000
SimilarWeb - Shopify	14.42	14.40	1.30	13.59	15.17	2868

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.

Overall, these comparisons suggest the revenue dataset captures a representative cross-section of e-commerce firms in terms of size, with an empirical mean that falls between the SimilarWeb full sample and the SimilarWeb Shopify-only sample, but with wider cross-sectional variation in firm sizes. 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 2 and A.5). 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 more representation of smaller e-commerce retailers.

Appendix E 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.

E.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-a-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 F and 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 a of 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

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

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:

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

The microfoundation for our model is therefore:

$$
\begin{aligned}
f(a_F,\theta)&:=X_F\Big(1-e^{-\lambda^F(\theta)a_F/X_F}\Big),\\
g(a_G,\theta)&:=X_G\Big(1-e^{-\lambda^G(\theta)a_G/X_G}\Big),\\
h(a_F,a_G)&:=x_{FG}\Big(1-e^{-\lambda^F(\theta)a_F/X_F}\Big)\Big(1-e^{-\lambda^G(\theta)a_G/X_G}\Big).
\end{aligned}
$$

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. □

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:

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

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 $a'_F + a'_G \geq a_F + a_G$, 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:

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

and

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

We thus have a contradiction to Assumption 3. □

The characterization is clear. Firms with higher θ purchase relatively more ads from F than from G. The relative ads 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 $\frac{\partial(g(a_F, \theta) - \hat{h}(a_F, a_G, \theta))}{\partial a_G} > \frac{\partial(g(a_F, \theta) - h(a_F, a_G, \theta))}{\partial a_F}$ for all $a_F \geq 0$ and for all θ; and (iii) $f(a_F, \theta) - h(a_F, a_G, \theta) - [\hat{f}(a_F, \theta) - \hat{h}(a_F, a_G, \theta)]$ is increasing in θ.

With the urn-ball model example, ATT can be modeled as the reduction of the ad-

efficiency 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 $\hat{a}_F + \hat{a}_G \geq a_F + a_G$. Then, by Assumption 4-(iii),

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

of

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

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.

E.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 which 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

$$p_F = C_F'(D_F(p_F, p_G)) \text{ and } p_G = C_G'(D_G(p_F, p_G)).$$

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 $\hat{p}_F \geq p_F$. There are two possibilities. Suppose first $\hat{p}_G \leq p_G$.

In this case, note first that, for each type (π, θ),

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

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

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

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

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

This means that

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

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

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, the ATT causes firms (advertisers) to substitute away from Facebook network to Google network, all else including ad prices equal.

2. Proposition 2 also shows that all else equal, including ad prices, the 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 the ATT, the price of Facebook ad and its 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.