Online Appendix: For Online Publication Only

Appendix A Additional Analyses for Advertising

A.1 Meta Campaign Objective Substitution

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

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

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

[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 because ambiguity arises when attempting to categorize them.

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

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.

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

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

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

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 specification (**??**) on a balanced panel of firms that have non-zero spend on Meta. The dependent variables are the share of spend (column 1), share of impressions (column 2), and whether there is non-zero spend (column 3). The left out category is off-platform conversions. Standard errors clustered at the firm level.

We consider a balanced panel of Meta firms that have positive spend on any campaign objective throughout the sample period and estimate the within-firm difference-in-differences specification (**??**). As with the main analyses, this allows us to control for differences across firms. Table OA3 displays the results that, consistent with the aggregate spending in Table OA2, show a precise null effect on substitution to on-platform objectives on the 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 **??** estimated using a balanced panel of firms that use both click- and conversion-optimized objectives before ATT.

A.2 Additional Analyses for Advertising Platform Substitution

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

Figure OA1: Event Study for Google Conversions

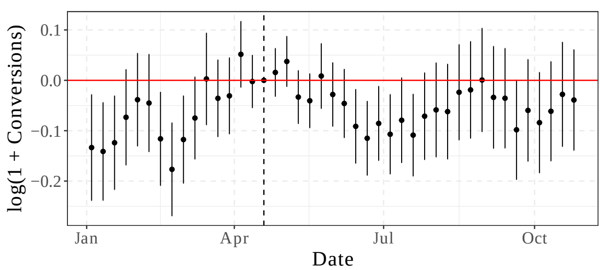

(a) Event Study for Google Search Conversions

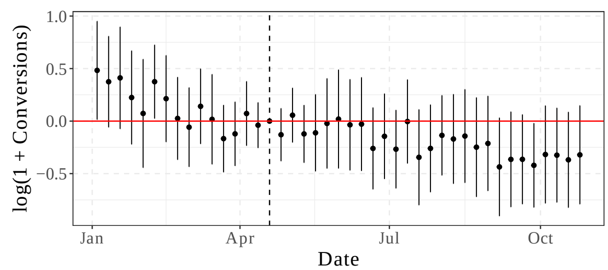

(b) Event Study for Google Display Conversions

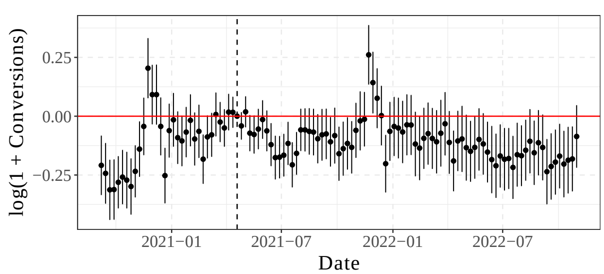

(c) Event Study for Overall Google Conversions

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

Figure OA2: Event Study for Online Advertising Spending

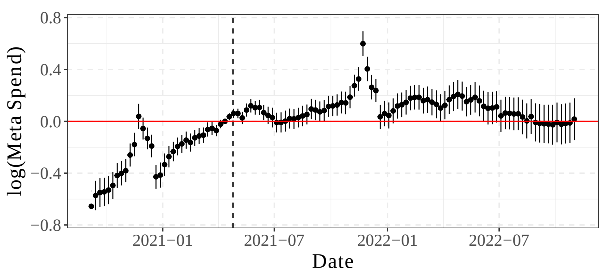

(a) Event Study for Meta Spending

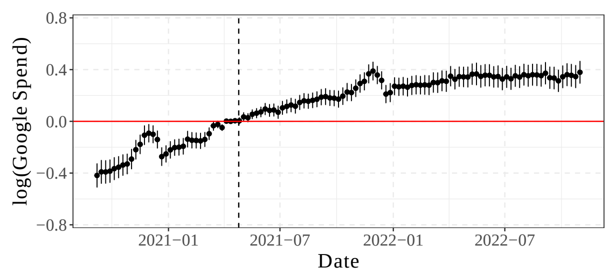

(b) Event Study for Google Spending

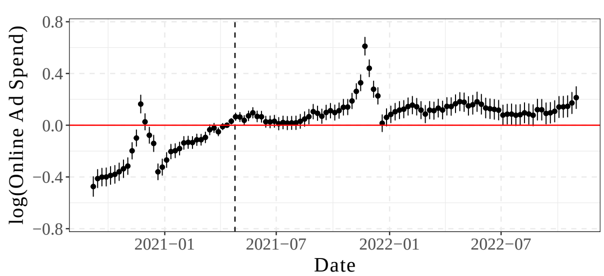

(c) Event Study for Total Advertising Spending

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

Changes in Total Advertising Spending: Given that we have some evidence that Google advertising is less impacted than Meta advertising, we explore whether and to what extent

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

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

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

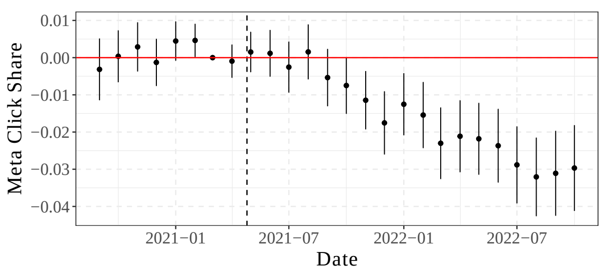

(a) Event Study for Online Ad Click Share

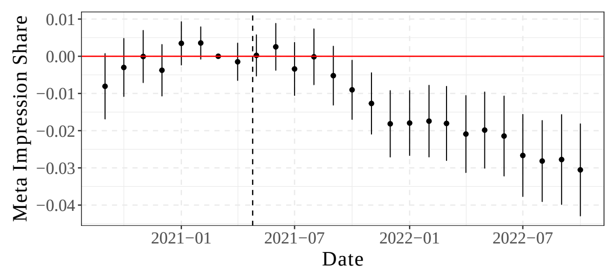

(b) Event Study for Online Ad Impression Share

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

In the main text, we show that the share of online advertising spend on Meta advertising declines. In this section, we estimate the event study specification (**??**) for total online advertising spending as well Google and Meta advertising spending individually, and we plot the estimates in Figure OA2. Figure OA2 shows that spending is increasing on both Google and Meta advertising before ATT at a similar rate, and that, after ATT, the spending on Google continues on a similar trend, whereas on Meta it slowly declines over time. The increasing time trend for online advertising spending is consistent with the revenue time trend shown in Figure OA4. We cannot disentangle whether this comes from supply-side or demand-side effects, but we additionally show in Figure OA3 that we get similar relative reductions in quantity of Meta ads using either clicks or impressions as our quantity measure.

Table OA4: Advertising Platform Substitution 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 using specification (**??**) with the different specified dependent variables (columns) and subsetted to the relevant platforms (rows). The first three rows present the results of the difference-in-differences specifications for share of spend, impressions, and clicks on overall spending on Meta, TikTok, and Google. The final two rows present the same dependent variables for Google Search and Google Display products. The results for Google Search and Google Display are estimated over the period January to October 2021 as Google launched its popular Performance Max product in November 2021, which led to substitution within the products in the Google ecosystem. Standard errors are clustered at the firm level.

Which firms are more likely to reallocate? We now explore whether firms more dependent on Meta were more likely to reallocate their spending by estimating the across-firm difference-in-differences specification (**??**). 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 that spend non-zero dollars on any advertising platform throughout the same sample period as before and the considered dependent variables are the online advertising market share of impressions, clicks, and spending across the different platforms. The first three rows of Table OA4 show the results for the online

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

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

Appendix B Additional Analyses for Revenue

B.1 Robustness Checks for Revenue Effects

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

B.1.1 Raw Data Plots

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

Figure OA4: Log Revenue for Low and High Meta Shares

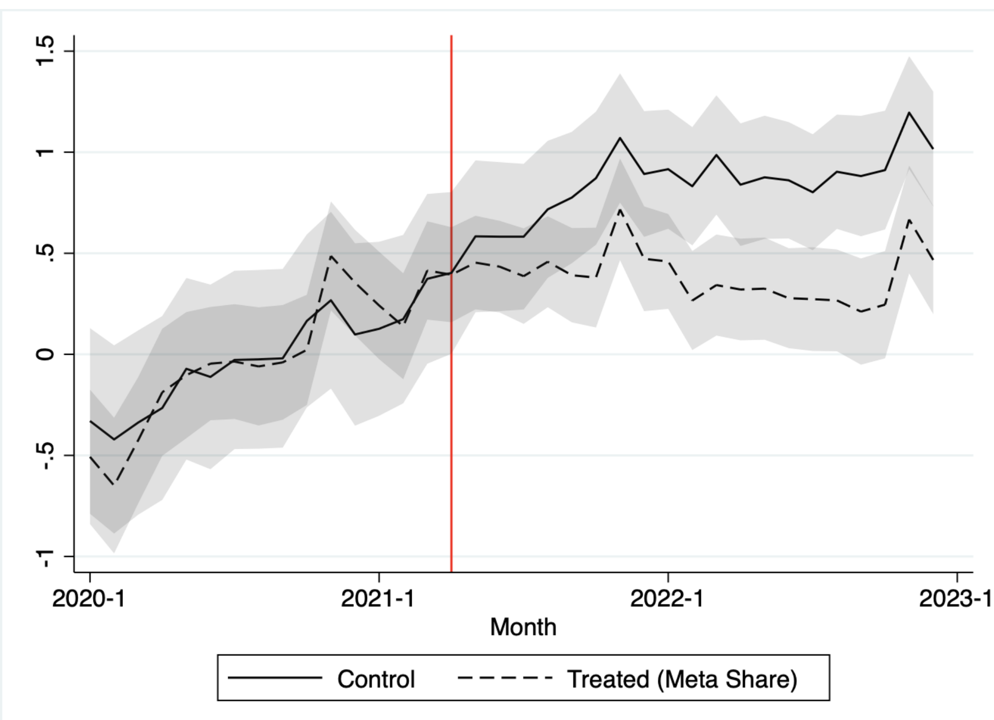

Notes: Results using the revenue dataset. High and low Meta shares are calculated using a median split of pre-ATT revenue from Meta traffic. Plot shows log(revenue) demeaned using the pre-ATT mean along with 95% confidence intervals.

Figure OA5: Log Revenue for Low and High iOS Shares

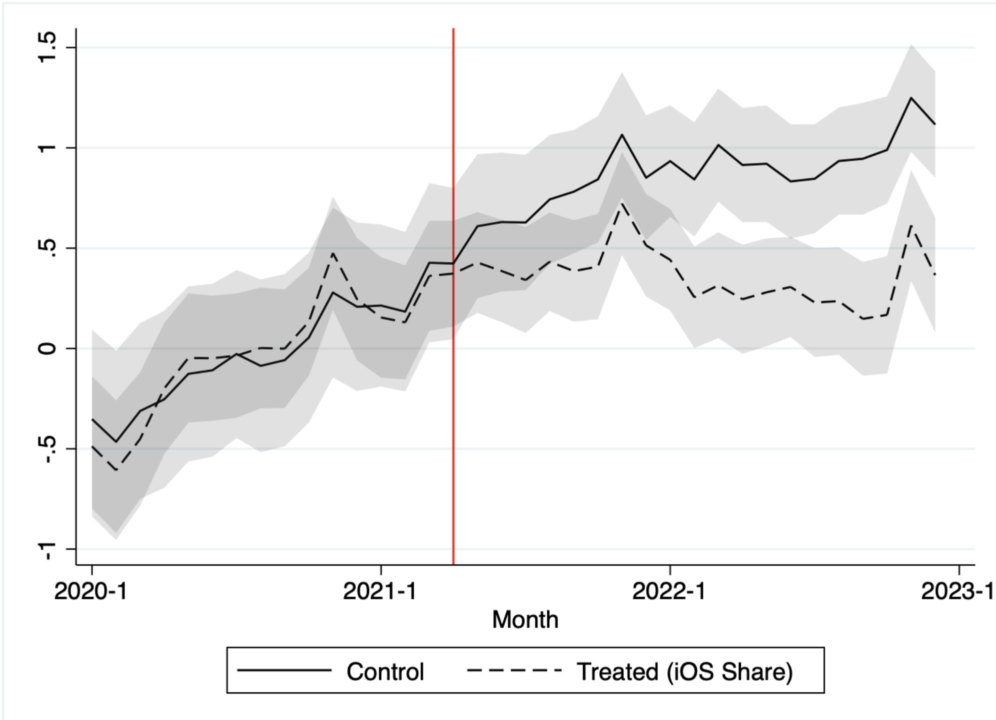

NOTES: Results using the revenue dataset. High and low iOS shares are calculated using a median split of pre-ATT revenue from iOS traffic. Plot shows log(revenue) demeaned using the pre-ATT mean along with 95% confidence intervals.

B.2 Reduced New Customer Acquisition as a Mechanism

In this section we characterize the extent to which revenue reductions were due to a decline in new customer acquisition. We do so by analyzing a secondary dataset associated with the advertising dataset. In this secondary dataset, we observe aggregated revenue data from the set of firms in the advertising dataset that provide access to their Shopify account and can directly link this to their advertising spending. Importantly for our purposes, these data provide us with a complete view of revenue for the firms. In particular, we observe the total revenue, the number of orders, and the fraction of orders that come from repeat customers. The Shopify data have no measurement issues as a result of ATT. Notably, the measurement of repeat customers relies on data unaffected by the changes from ATT since they are typically user-provided email addresses or phone numbers. Thus, the ability to separately measure new and repeat customers allows us to characterize the effects on new customer acquisition.

We note two significant limitations to this analysis. First, the number of firms that provide access to revenue data and are present through the full sample is relatively small. Second, the firms in the advertising dataset are strongly Meta-dependent and our representativeness exercises highlight that the advertising dataset skews towards smaller firms. As such, the estimates in this section will be relatively imprecisely estimated and underestimate the overall effect on orders and revenues, relative to our analyses in Section **??**. 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 (**??**) defining treatment as whether Meta advertising spend was above the mean within the set of firms, and we use the measure of the fraction of orders processed by a merchant that come from repeated customers. We consider monthly sales measures and, after joining with the advertising data, compute each firm's pre-ATT spend share for Meta advertising.

Figure OA5 presents the results for the dependent variables that we observe from Shopify: log(revenue), log(order count), and repeat order ratio. The takeaway across each of these is consistent: there is a reduction in orders and revenue of 20-22% and the fraction of total orders coming from repeat customers has increased. To understand whether this is simply a result of shifting advertising spend, we estimate the difference-in-differences specification controlling for the log of advertising spend. These results are presented in the second row of Table OA5. While this reports similar effect sizes for the ratio of repeat customers, we no longer find statistically significant reductions in revenues or orders though we still find comparable and economically large negative point estimates. For the repeat order ratio, the pre-ATT baseline for the share of orders coming from repeat customers was 33.93%, implying a 10.5% increase in the share of orders coming from repeat customers.

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

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

We conduct several robustness checks to validate the result that the primary reduction in orders comes from new customers. Figure OA6 considers the time-varying difference-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 OA7 presents the results for total, new customer, and repeat customer orders respectively. The results and effect sizes are largely consistent with our earlier analyses showing that the reduction in orders from new customers seems to be the driving force for the overall reduction in orders. The estimates for θ in Table OA7 indicate that there is moderate overdispersion in the data, supporting the suitability of the negative binomial regression model.

Table OA5: Difference-in-Differences Estimates for Sales

	Dependent variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
	log(Orders)		log(Revenue)		Repeat 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 are 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 OA6: Difference-in-Differences Estimates for 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 are 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 OA7: Difference-in-Differences Estimates for 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	
$\text{After}_t \times$ 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 are observed. The first two columns use total orders as the dependent variable, the next two focus on new customer orders, and the last two on repeat customer orders. The odd columns do not control for log(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 OA6: Time-Varying Treatment Effects for Sales Outcomes

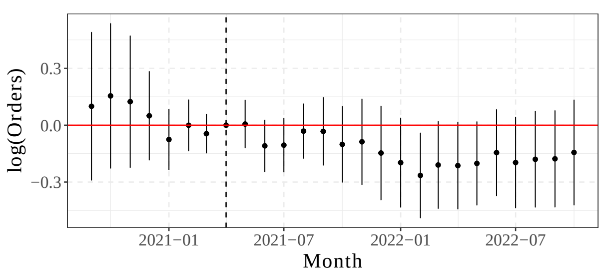

(a) DiD Estimates for Orders

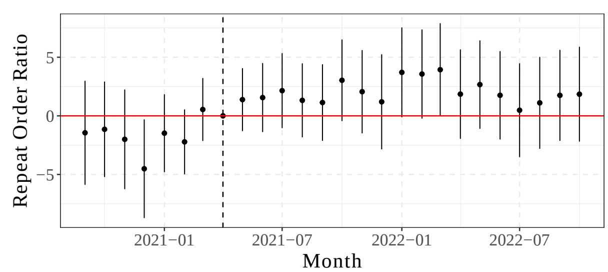

(b) DiD Estimates for Repeat Order Ratio

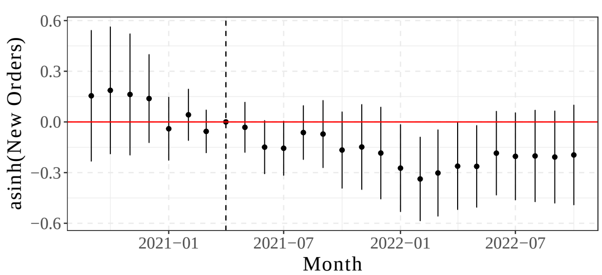

(c) DiD Estimates for New Orders

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

Appendix C Data Representativeness

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

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

C.1 Advertising Data Representativeness

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

C.1.1 Selection into Sample

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

C.1.2 Online Advertising Spending

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

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

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

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

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

The Kantar-Vivvix dataset contains firms spanning many industries that fall outside of 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[8] to subset down to firms operating on Shopify, which serves as a proxy for the e-commerce retailer category. BuiltWith is a technology-profiling firm that identifies the underlying technologies used by websites, including e-commerce platforms such as Shopify. It does so without selection bias, as BuiltWith collects publicly available information, using web crawlers to examine HTML, JavaScript, CSS, and other code on those web pages, thus identifying embedded technologies without firm participation or opt-in. The resulting collection of 7,900 firms within this Kantar-Shopify dataset is broadly representative of a wide cross-section of the

[7]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.

[8]https://builtwith.com

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

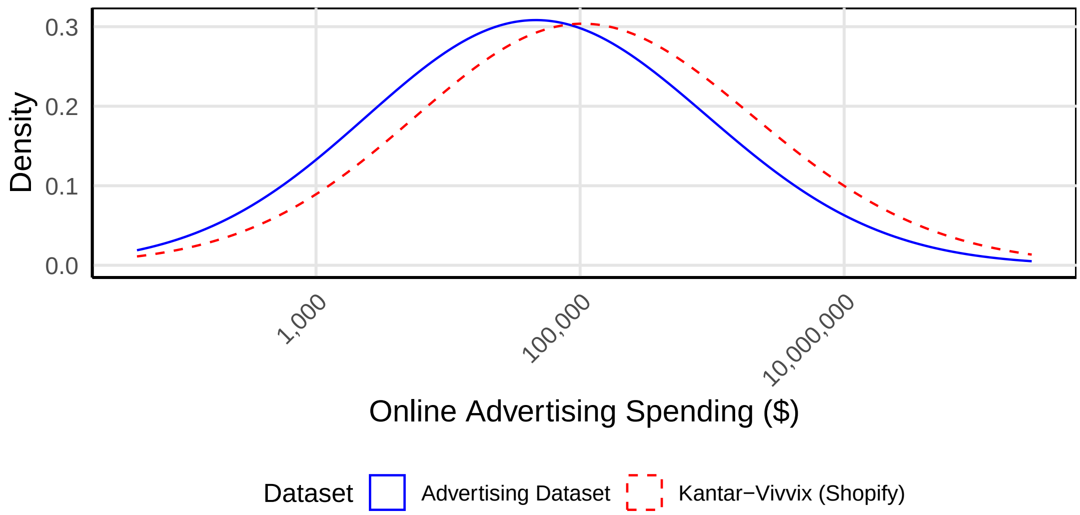

Dataset	Mean	Median	SD	25th	75th	N
Advertising Data	10.79	10.71	1.89	9.52	12.10	1475
Kantar-Vivvix (Shopify)	11.63	11.55	2.03	10.39	12.89	3248

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

e-commerce retailer category.[9]

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

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

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

C.2 Revenue Data

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

C.2.1 Selection into Sample

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

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

C.2.2 Representativeness Across Time

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

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

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

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

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

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

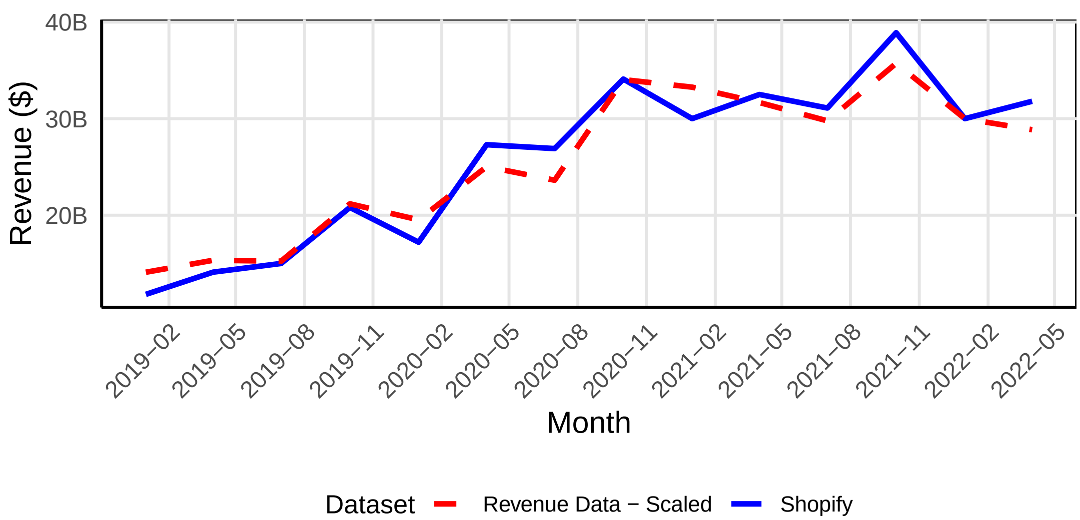

NOTES: This figure compares total revenue each quarter, as observed through the revenue data, to total revenue across Shopify merchants, as observed through cohort-level data publicly disclosed by Shopify. Firms within the revenue dataset began operating before April 2020; only merchants that began selling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

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

We provide a second across-time comparison using data from SimilarWeb. As noted above, SimilarWeb is a provider of digital intelligence and analytics that estimates 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 the sake of consistency with our other representativeness analyses, we obtained data on the top 7,000 e-commerce firms tracked by SimilarWeb (ranked by session counts from 2019 to 2021). From within this set, we then filtered down to the firms that maintained positive sessions in 2019, mirroring the selection criteria used for the revenue data.

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

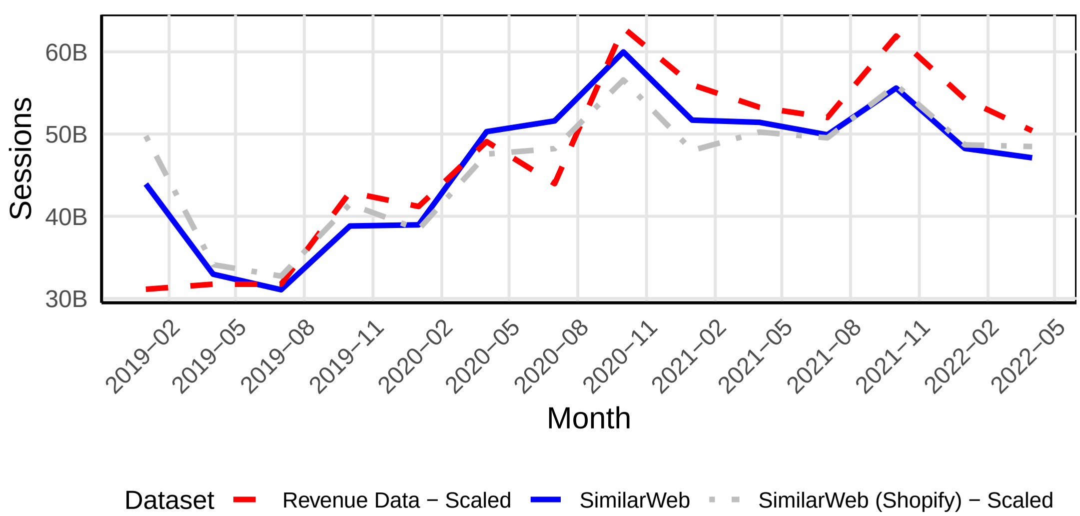

NOTES: This figure compares session counts over time across three samples: (1) the revenue dataset (mean-scaled), (2) all firms tracked by SimilarWeb, and (3) the subset of SimilarWeb firms operating on Shopify (mean-scaled). Data series are mean-scaled relative to the full SimilarWeb sample to facilitate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

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

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

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

C.2.3 Representativeness in Firm Size

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

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

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

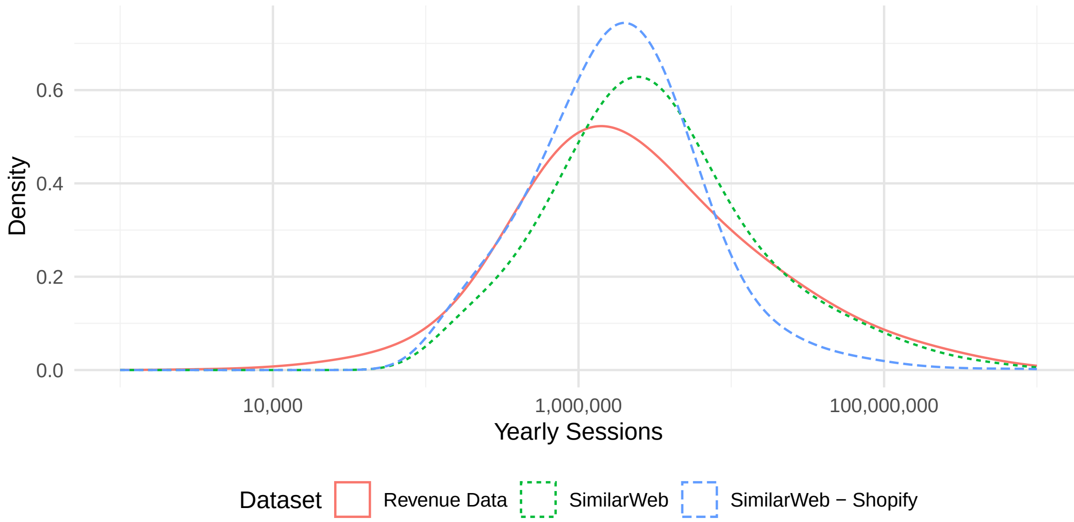

Dataset	Mean	Median	SD	25th	75th	N
Revenue data	14.71	14.49	1.85	13.54	15.75	773
SimilarWeb	14.99	14.83	1.63	13.92	15.87	7000
SimilarWeb - Shopify	14.42	14.40	1.30	13.59	15.17	2868

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.

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

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

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

Appendix D 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.

D.1 Firm Behavior

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

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

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

We make the following assumptions:

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

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

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

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

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

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

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

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

or

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

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

Example 1 (Microfoundation based on urn-ball model)**.** *Suppose there are two outlets 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: $x_F\Big(1-e^{-\lambda^F(\theta)a_F/X_F}\Big)$.

Similarly, the number of single-homing customers acquired by G is $x_G\Big(1-e^{-\lambda^G(\theta)a_G/X_G}\Big)$.

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:

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

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

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

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

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

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. $\square$

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.

D.2 Equilibrium of Ads Markets

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

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

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

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

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

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

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

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

2. Proposition 2 also shows that all else equal, including ad prices, ATT causes revenue loss for all firms but more so with higher Meta dependency (i.e., higher θ).

3. Proposition 3 analyzes the general equilibrium effect: with ATT, the price of Meta ads and their overall demand/quantity fall.

4. While the richness of the model limits the analytical results to those stated in Proposition 3, we can draw further implications.

   (a) The second statement of Proposition 3 means that a significant proportion of, or possibly all, firms reduce their purchase of Meta 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 Meta 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 Meta 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 Meta ads relative to Google ads likely falls after the ATT shock.