

NBER WORKING PAPER SERIES

DIGITAL ADVERTISING AND MARKET STRUCTURE:
IMPLICATIONS FOR PRIVACY REGULATION

Daniel Deisenroth
Utsav Manjeer
Zarak Sohail
Steven Tadelis
Nils Wernerfelt

Working Paper 32726
<http://www.nber.org/papers/w32726>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024, Revised February 2026

Authors are listed alphabetically. We are grateful to JP Dubé, Sam Goldberg, Sandy Jap, Ginger Zhe Jin, Thierry Magnac, Dina Mayzlin, Nadav Tadelis, and Joel Waldfogel for constructive discussions. We also thank many seminar participants for valuable comments on our research. Deisenroth, Manjeer and Sohail are employees of Meta Inc., a firm that sells targeted advertising on its digital platform. Tadelis was a paid consultant for Meta Inc from June 2021 through December 2022 and was not paid for any part of this research. Wernerfelt was an employee at Meta Inc. from June 2016 until June 2023 and currently has a consulting relationship with Meta Inc. but was not paid for any part of this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Daniel Deisenroth, Utsav Manjeer, Zarak Sohail, Steven Tadelis, and Nils Wernerfelt. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Digital Advertising and Market Structure: Implications for Privacy Regulation
Daniel Deisenroth, Utsav Manjeer, Zarak Sohail, Steven Tadelis, and Nils Wernerfelt
NBER Working Paper No. 32726
July 2024, Revised February 2026
JEL No. D22, D40, L10, L59, M38

ABSTRACT

We study how changes in the effectiveness of targeted digital advertising shape firm behavior and product-market outcomes. We leverage the introduction of Apple's App Tracking Transparency (ATT) framework, which limited online advertisers' access to user data and thus substantially reduced their advertising effectiveness. Using proprietary advertiser-level data from Meta and industry-level outcomes from the Bureau of Labor Statistics, we show that greater exposure to the ATT shock led to lower ad spend, fewer establishments, and higher product prices. The findings suggest that digital advertising primarily expands markets and that privacy policies can have unintended consequences on product market competition.

Daniel Deisenroth
Meta Platforms, Inc
dbdeisenroth@meta.com

Utsav Manjeer
Meta Platforms, Inc
utsavmanjeer@meta.com

Zarak Sohail
University of California, Irvine
zaraks@uci.edu

Steven Tadelis
University of California, Berkeley
Haas School of Business
and NBER
stadelis@haas.berkeley.edu

Nils Wernerfelt
Northwestern University
Kellogg School of Management
Department of Economics
nils.wernerfelt@kellogg.northwestern.edu

1 Introduction

Today more than 75% of all advertising spend in the U.S. is digital (Feger, 2024). A key advantage of digital advertising over traditional channels is the ability to target ads to consumers using detailed data on their online activities. This harnessing of user data has, however, brought about major privacy and data protection laws that continue to emerge and shape the industry, yet we still have limited evidence on their inherent tradeoffs.

Privacy regulation can offer meaningful benefits to consumers—most notably by giving them greater control over their data, reducing unwanted tracking, and limiting the potential misuse of personal information. At the same time, these protections can make it harder for firms to target ads effectively, raising the possibility of meaningful downstream economic effects. This tension gives rise to a central policy question: if privacy regulation reduces digital advertising effectiveness, what are the consequences for firms and the product markets in which they compete? Prior theoretical and empirical work shows that reductions in advertising effectiveness can affect advertising behavior, competition, and prices in different ways depending on the context. Advertising can either raise or lower entry barriers, and which effect dominates in digital targeted advertising is an empirical question.

We address this question using a natural experiment created by Apple’s rollout of its App Tracking Transparency (ATT) framework, which sharply reduced the effectiveness of a substantial share of digital advertising. The following anecdote is illustrative of ATT’s impact on one business. In March 2022, NPR’s *Planet Money*¹ profiled a firm selling a highly niche product—ramps for small dogs—that had previously relied on targeted digital advertising to reach a narrow set of likely buyers. Prior to ATT, precise targeting lowered the cost of reaching these consumers, allowing the firm to operate profitably despite limited baseline demand. After ATT, customer acquisition costs rose sharply, reducing the set of consumers the firm could profitably reach, and thereby limiting its scale. More broadly, the firm’s reliance on targeted digital advertising reflects that of many small and medium-sized

¹“Tech giants and tiny dogs,” available at <https://www.npr.org/transcripts/1087425495>.

businesses in the United States, for whom digital advertising plays a particularly central role in reaching potential customers (Raggio et al., 2025).

Not all advertisers, however, were equally exposed to the negative impact of ATT, which is at the core of our identification strategy. Specifically, using detailed internal data from Meta, we estimate how much individual advertisers—and industries—were affected by this shock. We then compare outcomes for those more versus less exposed to ATT using a differences-in-differences design. For outcomes, we combine data on U.S. Meta advertisers with BLS data to examine how the shock influenced not only online advertising behavior, but also the number of establishments and product prices across U.S. industries.

We find that ATT had significant and economically meaningful effects on our outcomes. Advertisers more exposed to the shock reduced their ad spend by roughly 7-10% over the following year and shifted the campaigns they did run toward formats that were less impacted by ATT. Industries with higher exposure saw about a 1.4% decrease in the number of establishments and roughly a 3-4% increase in product prices. The Online Appendix provides extensive robustness checks, including alternative exposure measures, specifications, and advertiser samples.

Our paper makes two contributions. First, we provide large-scale empirical evidence on how digital advertising impacts market structure. Economists have long been interested in the relationship between advertising and market structure outcomes, going back to at least Marshall (1890). Theoretical work has modeled advertising as influencing demand in different ways. Following Bagwell (2007), the literature typically distinguishes among advertising that conveys product information (“informative”), alters consumer preferences (“persuasive”), or enhances the utility of consuming the advertised good (“complementary”).²

The upshot is that different views of digital advertising imply different market consequences. If, for example, advertising mainly informs consumers and expands the market, then making it less effective should contract demand, leading to fewer firms and potentially to higher

²Well known examples of each are Dixit and Norman (1978), Ozga (1960), and Stigler and Becker (1977). We note, however, that not all theories of advertising fall cleanly within these commonly used categories (e.g., the signaling based models following Milgrom and Roberts (1986)).

prices. But if instead advertising mostly creates brand loyalty and fuels business stealing, then reducing its effectiveness can lower wasteful rent-seeking, raise profits, attract entry, and put downward pressure on prices.

Empirical evidence supports both views of advertising in different settings. For example, evidence consistent with market-expansionary effects of advertising have been documented in direct-to-consumer pharmaceuticals (Shapiro, 2018), liquor (Milyo and Waldfogel, 1999), and yogurt (Akerberg, 2001). In contrast, evidence consistent with business stealing has been found in smoking (Eckard Jr, 1991), analgesics (Anderson et al., 2016), and junk food (Dubois et al., 2018).³ To our knowledge, no prior work has identified which of these dynamics best characterizes digital advertising, making it an open empirical question. Our results fit the market-expansion view. To make this concrete, the Appendix presents a simple model—following Sutton (1991)—that generates comparative statics matching our empirical patterns.

Our second contribution is to provide evidence on the tradeoffs embedded in today’s privacy policy debates.⁴ Over the past decade, privacy regulation has expanded dramatically: nearly 100 countries have adopted data-privacy laws, and an estimated 82% of the world’s population now lives under some form of national privacy legislation (Apacible-Bernardo and Bushey, 2025; Greenleaf, 2023). Compliance with these policies can be costly for firms. For instance, the International Association of Privacy Professionals’ annual survey found that many firms report expenditures in the millions of dollars to meet GDPR requirements (Hughes and Saverice-Rohan, 2018). Another report estimates that more than \$300 billion in market value was lost from major ad platforms following Apple’s launch of ATT (Fox, 2022). Despite these stakes, we still have limited empirical evidence to guide the relevant

³These recent empirical papers build on a large literature going back to Bain (1956), Telser (1964), Benham (1972), and others that were early attempts to measure the relationships between market structure outcomes and advertising. Early empirical work often ran cross-industry regressions of advertising investments on different market structure variables; even today, given the nature of the problem, all empirical research we are aware of is non-experimental and leverages natural experiments for identification.

⁴To be clear, ATT was a company-initiated product change, not government regulation. Still, the types of data affected by ATT have been—and continue to be—a central focus of regulatory efforts worldwide.

tradeoffs—both the privacy benefits to consumers and the potential downstream effects on firms and markets.⁵

In this contribution, we add to a growing empirical literature on the downstream effects of privacy regulation. This literature also relies heavily on natural experiments, often exploiting geographic and temporal variation in the rollout of privacy laws. For example, Goldfarb and Tucker (2011) use survey data to show that the EU’s Privacy and Electronic Communications Directive reduced ad effectiveness; Demirer et al. (2024) document how GDPR increased cloud-storage data costs for firms; and Goldberg et al. (2024) study GDPR’s effects on website traffic and revenue. Dubé et al. (2025) provides a recent and comprehensive overview of unintended effects of privacy regulation. A related set of papers examines how privacy regulations affect market concentration in specific product markets; consistent with our findings, this work generally reports increased concentration following regulatory enactment (Johnson et al., 2023; Janssen et al., 2022; Ghanavi et al., 2025). We extend this literature in a meaningful way by combining online and offline data and by examining a broad range of product markets. The closest paper to ours is Aridor et al. (2024), who use a similar identification strategy to study the effects of ATT on advertising effectiveness, budget allocation, and firm revenue. Consistent with our results, they find sizable negative effects of ATT on these outcomes. Their analysis, however, focuses on the advertising behavior of e-commerce firms only and does not extend to market structure impacts, which is the main focus of our study.

The rest of our paper is structured as follows. Section 2 provides institutional details on digital advertising and the nature of our natural experiment; Section 3 describes our data; Section 4 outlines our empirical approach; Section 5 contains our results; and Section 6 concludes. Our Appendix contains our model, and the Online Appendix describes further aspects of our data and details additional robustness checks

⁵In our paper we do not estimate privacy benefits to consumers. Prior work (e.g., Athey et al. (2017); Lin (2022)) generally finds low average consumer valuations, but these results do not imply that consumers place low value on all forms or dimensions of privacy. Understanding which aspects of privacy consumers value, and by how much, remains an important area for future research.

2 Institutional Details

Millions of businesses in the U.S. use digital advertising platforms to reach and engage with customers. In this section, we provide relevant, high level institutional background on how digital advertising works today on platforms like Meta and the role of user data in that process. For a more detailed account see Wernerfelt et al. (2025), from which this section borrows heavily.

First, we note that different advertisers may have many different objectives when running digital advertising campaigns. For example, an advertiser may want to increase brand awareness, grow an email mailing list, drive purchases, etc. Whenever an advertiser sets up a campaign, Meta asks the advertiser for their objective to align the platform’s targeting capabilities with the objective. The reason behind this question is as follows. Suppose that an advertiser wants to drive “likes” on their Facebook page. Given a budget and target audience, Meta could distribute the ad spend uniformly across the target audience. It is very likely, however, that different users will respond differently to the campaign. Suppose that as the campaign runs, the data show that only men click on the ad and “like” the advertiser’s Facebook page. Meta could continue to spend the advertiser’s budget on men and women, but that would presumably neither be good for the advertiser (as they would be spending ad budget on users who are not liking their Facebook page) nor for women (as they would see ads they are not interested in).

One way this scenario could be improved (in some, but not all circumstances) for both the advertisers and users would be for Meta to train a model in real time to predict the probability a user would like the advertiser’s Facebook page, and use that to dynamically adjust ad spend to focus on users who are likely to deliver the advertiser-specified objective. In this case, Meta would train a model, recognize that only men are liking the page, and shift the ad spend to focus on men, potentially improving outcomes for both the advertiser

and users. This process is commonly referred to as “delivery optimization” by ad platforms and underlies ad delivery at many major ad platforms today.⁶

Delivery optimization is arguably the major upside of digital advertising. With delivery optimization, Meta can leverage its large amounts of data and advanced algorithms to find the users to show ads to who are most likely engage as the advertiser wants. Such a process is not feasible over TV, print, radio, etc., and may be particularly valuable for small businesses who face niche demand or do not know who their customer base is. Indeed, most advertisers on Meta are small and medium sized businesses (e.g., Chouaki et al. (2022)), and a large share of their advertising is optimized for sales.

Note, however, that sales are a different outcome than likes on an important dimension. “Likes” are directly observed by Meta on its platforms. This is referred to as “first party data” or “on platform data” for Meta. A complication here is that many outcomes the advertisers care about—such as sales—happen off platform from Meta. For example, upon clicking on a Facebook ad, a user is often taken to Safari, Chrome, or an advertiser’s mobile app. Left to their own data, Meta would have no way of knowing what actions the user took on those other platforms, and thus it would have no way of training prediction models for important outcomes like purchases. This is where third party data and ATT come in.

2.1 Digital Advertising and Third Party Data

Third party data are data generated outside of the focal ad platform’s own domains. For example, if a user adds an item to their shopping cart on an Amazon webpage or in the Instacart app, that would be third party data to Meta but first party to Instacart and Amazon. In isolation, Meta would have no way of knowing what that user did on Instacart or Amazon, while both Instacart and Amazon can observe that directly.

⁶Note in our description, delivery is not based on *incremental* actions. A common criticism of delivery optimization is that the conversions may not be incremental. For example, the algorithms may just target users who were already going to purchase the product. We take no stance on this; whatever share of conversions are incremental, it remains an empirical question how much ATT affected firm outcomes. We note, however, that both Wernerfelt et al. (2025) and Tadelis et al. (2023) report in holdout experiments on Meta where some users are randomly withheld from seeing ads, a large share of advertisers see no purchases, suggesting that many reported conversions are in fact incremental.

In order to obtain data on users’ activity off of Meta, Meta uses online tracking technologies such as the ‘Meta pixel.’⁷ To roughly understand how this technology works, suppose an advertiser is selling shoes. They can install a Meta pixel on their website and configure it to track purchase events. Once installed, whenever a purchase event occurs on the website, the pixel ‘fires’ and transmits information to Meta about the event and user. Hence, if an advertiser has a Meta pixel installed, Meta can now get off platform data that can then be used to optimize ad delivery for purchases.⁸

The Meta pixel only sends information to Meta, but given how important these data are for digital advertising, nearly every major ad platform today has analogous technologies.⁹ Indeed, much of the digital advertising ecosystem today was built around use and access to such online tracking data. With that industry backdrop, Apple’s ATT policy was arguably the largest shock to this third party data ecosystem to date.

2.2 Apple’s App Tracking Transparency

After announcing the policy change in 2020, Apple launched its ATT feature as part of the iOS 14.5 software update for mobile devices, such as iPhones and iPads, on April 26, 2021. Apple’s stated aim was to “improve transparency and empower users” with greater control over how their data was collected and used by third party platforms and developers.¹⁰ Prior to the implementation of ATT, users could choose to restrict third party data sharing across


⁷The ‘Meta Software Development Kit,’ or Meta SDK, is another commonly used technology; it provides similar functionality to the Meta pixel but for mobile apps instead of websites.

⁸Note that the advertisers must install the pixel themselves, but that they have an incentive to as it will help optimize their ad performance. Without the third party data signals from pixels and SDK’s, ad delivery can only be optimized for on platform signals. To the extent that the on platform signals (e.g., clicks) are a good proxy for the off platform signals, there may be little loss of ad effectiveness. However, Wernerfelt et al. (2025) found that optimizing for the on platform outcomes greatly diminishes ad effectiveness when the advertiser’s objective is an off platform outcome.

⁹For example, Google’s technology is referred to as the Google tag and LinkedIn’s as the LinkedIn Insight Tag.

¹⁰Apple had many press releases with these stated objectives. See e.g., <https://www.apple.com/newsroom/2021/01/data-privacy-day-at-apple-improving-transparency-and-empowering-users/>. Importantly, this was the single major change in iOS 14.5; other updates included two new English Siri voice options, the ability to unlock an iPhone with the Apple Watch, and support for Xbox controllers.

Figure 1: Apple App Tracking Transparency User Interface



apps by using the “Limit Ad Tracking” setting. The introduction of ATT, however, required apps to obtain *explicit* permission from the user to enable third party data sharing.¹¹


Specifically, when a user opens an app for the first time after installing iOS 14.5, they are shown a pop up that asks if they would like to allow tracking from the focal app. As shown in Figure 1, users can either “Ask App Not to Track,” which will disable third party data tracking, or “Allow” the app to track third party data. All apps using the iOS platform must offer this prompt to users, the text was chosen and set by Apple, and if an app were caught failing to comply (e.g., by still trying to track users), Apple threatened to remove it from the App Store. Only if a user decides to allow the collection and use of their third party data (hereinafter “opt-in users”), can applications observe the user’s activity across any third party applications and websites.

As a result of this simple prompt, the share of iOS users who could no longer be tracked off platform increased dramatically, making it much harder for Meta and other advertising

¹¹In this sense, our natural experiment is a case study for the impact of switching to an ‘opt in’ consent model for privacy. Thus, our results may be relevant to future privacy policy discussions that are considering different models of consent, though how much we can extrapolate from the specific wording, interface, etc. depends on the specific setting. We thank JP Dubé for making this point.

platforms to optimize delivery for off platform outcomes. As an indicator of how iOS 14.5 spread throughout the iOS population, Figure 2 shows around 80% of iOS users had updated to iOS 14.5 (or higher) within three months of ATT’s launch (Gogi Gupta, 2022).¹² Conditional on having iOS 14.5 or higher, opt out rates were quite high, as we see in our data.

Figure 2: Adoption of iOS 14.5 and higher across time



Notes: This figure shows the adoption of iOS 14.5 and subsequent versions over time, with an accelerated adoption in June 2021 (Gogi Gupta, 2022).

In short, ATT was a substantial, focused shock to the digital advertising industry. Experimental evidence suggests that losing the data targeted by ATT could increase advertisers’ cost per incremental customer by more than 30% (Wernerfelt et al., 2025), which is a large enough shock to have meaningful downstream economic effects. Core to our identification strategy is the fact that this shock was not uniform across advertisers. Leveraging our access to detailed internal data from Meta, we are able to estimate this variation in exposure, which we describe in detail next.

¹²In June 2021, Apple started pushing iPhone users more aggressively to update their iOS – as they often do with new operating systems – hence the sharp increase then.

3 Data

We combine novel on platform data representing U.S. advertisers on Facebook and Instagram, together with publicly available data across U.S. industries from the Bureau of Labor Statistics (BLS) and map the two for our comprehensive analyses.

3.1 Meta’s On Platform Data

Meta’s internal data serves two key purposes for our analysis: servicing our identification strategy and providing key outcome variables. Below we outline each, as well as how we selected our sample of advertisers.

3.1.1 Identification Strategy

For our identification strategy, we derive an advertiser-level measure of exposure to ATT using historical ad impression data from Meta. Specifically, for each advertiser in our sample, we fix a date pre-ATT and look at every single Instagram and Facebook user who received an ad impression from that advertiser. Meta logs the device type of users, allowing us to see what share of these impressions went to iOS users. Thus for each advertiser we can calculate what share of their ad impressions went to iOS users, which we treat as our measure of exposure. For example, imagine that Advertiser-*A* delivered a total of 5,000 impressions of which 3,300 were delivered to iOS users, while Advertiser-*B* delivered a total of 4,000 impressions of which 1,600 were delivered to iOS users. Then Advertiser-*A*’s exposure is 66% ($3,300/5,000$) while Advertiser-*B*’s exposure is only 40% ($1,600/4,000$).

We note that for robustness we also created several other measures of exposure, including the share of the target audience that were iOS users, the share of ad spend that went to users who were iOS, and the share of ad spend that went to users who were iOS who eventually opted out of tracking after ATT. All these measures are correlated; in the Online Appendix we rerun our analyses with each and find similar results. These exposure measures highlight the upside of our detailed data from Meta.

Separately, for our results on establishments and prices, we must aggregate these measures of exposure up to the industry level to match to the granularity of our outcome data. For this, we took the total number of impressions from advertisers in a given industry and calculated what share of those went to iOS devices. We explain in more detail later how we matched these on platform measures of exposure to the off platform outcome variables.

3.1.2 Outcomes

We study four different outcomes for advertisers in our sample: (i) advertising spend, (ii) activity levels, (iii) advertising spend per active advertiser, and (iv) the share of campaigns that use the data affected by ATT.

Regarding (i) and (ii), for every advertiser account in our sample, we are able to observe when they advertise (activity levels), and if so, how much they spend (ad spend). Activity levels are recorded as number of months per quarter the account ran a campaign; spend is log transformed and winsorized at the 95th percentile given its high skew.^{13,14} For (iii), we take the total amount of ad spend and divide by the number of active accounts; this gives us a sense of conditional on advertising, how much advertisers are spending. For (iv), we focus on what share of campaigns an account created that were offsite conversion optimized. Recall that offsite conversion optimized means that targeting is optimized on offsite data, such as purchases, which is the data at risk from ATT. Instead, onsite conversion optimized (the complement) means that targeting is optimized for onsite data such as likes or clicks on ads. Following the marketing literature (e.g., Gordon et al. (2023)), we refer to these as “lower-” and “upper-” funnel campaigns, respectively, and we study how the relative share of such campaigns change across time after ATT.

¹³Rerunning our ad spend results under similar transformations (e.g., inverse hyperbolic sine) and with similar levels of trimming or winsorizing does not change the results meaningfully.

¹⁴November 2020 had a large, erratic amount of ad spend due to the first Black Friday/Cyber Monday during the COVID pandemic. To smooth the effects of that shock, in our aggregation we replaced that month for each account with the average of October and December spend; rerunning dropping November or smoothing using a within-account rolling average yields similar results.

These detailed data on advertiser accounts highlight the upside of our Meta data for our analysis. The sample is further able to be minimally selected, aiding representativeness, as we describe next.

3.1.3 Sample Selection

We fix a large sample of advertiser accounts and track their on platform behavior over time. For tractability, we focus our primary analyses on all U.S. advertisers active on Facebook or Instagram on July 31, 2020.¹⁵ Our choice of the sample period is motivated by two factors. First is to have a sufficiently long pre-ATT time period in our individual-level data. However, we are limited in how far back in the pre-ATT period we can go due to natural churn in advertisers over time. Second, we want to ensure that the advertisers in our sample are not idiosyncratically affected by the onset of COVID-19 during the first half of 2020. In the early period of COVID-19, advertising patterns on Meta were erratic and going back too far risks exposure to this thrash.

Finally, in our analysis of advertiser behavior, we restrict our sample to advertisers whom Meta estimates to be small advertisers. This metric is primarily based on level of historical ad spend, but also includes additional information such as whether a new ad account may belong to a larger parent company. We do this because small advertisers are most reliant on the data impacted by ATT and tend to advertise exclusively on lower funnel outcomes; larger advertisers rely more heavily on upper funnel advertising and so are impacted differently by the ATT shock. We present summary statistics for our sample, broken down by treatment and control groups in Table 1. In total, our data span 869,191 advertisers over 12 quarters from Q1 2020 to Q4 2022.

¹⁵Computational costs to derive our advertiser-level exposure metrics were quite substantial. For robustness we also collected the data on June 30, 2020 and reran analyses with that panel of advertisers. As shown in the Online Appendix, our results did not change meaningfully.

3.2 U.S. Industry Data

We leverage publicly available data from the Bureau of Labor Statistics (BLS) on industries at the 6-digit NAICS (North American Industry Classification System) code level for our outcome variables. We use two different sources for our outcome data on prices and establishments:

Prices. We use data from the BLS’ Producer Price Index (PPI) aggregated to the quarter-6-digit NAICS code level. The PPI measures the average change over time in the selling prices for all U.S. domestic producers of goods and services. It covers all manufacturing industries, as well as about 70% of service industries. We use the PPI instead of the Consumer Price Index (CPI) as the former contains prices for a much broader range of industries, allowing a much better match to the industry range of advertisers on Meta.

Establishments. We leverage quarterly data on the number of existing business establishments from the BLS’ Quarterly Census of Employment and Wages (QCEW). The QCEW uses microdata from the federal-state Unemployment Insurance programs to get counts of the number of operating establishments at the level of our data. This rich dataset covers all private sector establishments, as well as state and local governments covered under the UI program. We use the 6-digit NAICS industry codes to utilize this data at the state-quarter-industry level.

3.2.1 Matching Meta’s Industry Verticals to BLS industries

Meta internally classifies each advertiser account according to one of 176 industries. To map our advertiser-level measures of exposure to ATT to the BLS data, we thus aggregate our exposure measures to the industry level on Meta, and then match those industries to 6-digit NAICS codes. We did this matching manually, going over each 6-digit NAICS code description as well as that of Meta’s industry classification. This mapping was not one-to-one; for example, sometimes multiple Meta industries corresponded to one NAICS code and vice

versa. As a result, the merged dataset contained a smaller set of 98 new industry definitions that at times included multiple Meta industries or NAICS codes. These form the basis for our analyses on the number of establishments and prices; in the Online Appendix we describe the matching process in more detail and provide several examples.

4 Empirical Strategy

We employ a difference-in-differences framework to estimate the impact of ATT’s data sharing restrictions on our outcome variables. Given that ATT was rolled out globally, our empirical strategy exploits variation in the extent to which different advertisers and industries were more versus less impacted by the ATT shock.

4.1 Share of iOS Impressions and Treatment Assignment


As described in Section 3.1.1, for every advertiser and every industry we calculate our measure of exposure as the pre-ATT share of Facebook and Instagram ad impressions that were delivered to iOS users, henceforth referred to as “ATT impact.” For all our analyses, we compute the median ATT impact across units in our sample and simply define ‘treatment’ as being above median and ‘control’ as being below median. Note this means our counterfactual is not the effect of going from no ATT to ATT but rather from being below to above median exposed, similar to Aridor et al. (2024).^{16,17} As context, across advertiser accounts, our median measure of ATT exposure was 0.557, with 25th and 75th percentiles 0.381 and 0.695.

¹⁶We note that this captures a measure of exposure to the ATT shock. It does not necessarily capture how reliant an advertiser is on digital advertising. For example, if a firm is highly exposed but can easily substitute to TV or another unaffected channel, the consequences may not be large for them. In such cases, we may see drops in advertising behavior on Meta but no impact in our offline datasets. The fact that we still find evidence of an effect of ATT in our off platform results suggests such substitution was limited for many firms.

¹⁷There is an instrumental variables interpretation to our analysis. Intuitively, we are interested in the regression of market structure outcomes on digital advertising effectiveness, where our measurement of ATT impact is an instrument for the latter. Our analysis produces reduced-form estimates from regressing the market structure outcomes directly on our ATT impact measure. We thank Joel Waldfogel for making this observation.

Table 1 shows pre-ATT summary statistics for our main on platform variables (Panel A) and off platform variables (Panel B) for each of the above and below median exposure units. The summary statistics for the on and off platform variables do not seem to differ substantially across treatment and control, evidence of balance across our treatment assignment. Given the sample sizes, a large share of the differences are statistically significant, but across Panels A and B the average absolute difference as a percent of the below median exposure mean is only 3.4%. Similarly, in Figure 3 we provide the distribution of advertiser accounts over the main Meta industries across treatment and control. The mean absolute difference in percentage points is 0.39, suggesting the groups are relatively comparable along these observables.¹⁸

Figure 3: Distribution over listed industries for advertiser accounts above and below median ATT exposure.



¹⁸Professional services has disproportionately more high exposure accounts, but the difference there is only 2.6 percentage points and as we show in the Online Appendix, rerunning our analyses dropping that industry does not materially change our results.

Table 1: Summary statistics for control and treatment groups.

Variable	Below Median Exposure to ATT						
	Mean	Std. Dev.	10th Percentile	25th Percentile	Median	75th Percentile	90th Percentile
<i>Panel A. On platform variables</i>							
Ad Spend (log)	2.505	2.451	0.000	0.000	2.443	4.512	5.867
Active Advertisers	1.223	1.211	0.000	0.000	1.000	2.000	3.000
Share Upper Funnel	0.916	0.253	0.750	1.000	1.000	1.000	1.000
Ad Spend per Active Advertiser	137.118	923.271	5.000	12.500	32.907	92.759	245.210
<i>Panel B. Off platform variables</i>							
Number of Establishments (log)	5.616	1.726	3.367	4.394	5.591	6.794	7.841
Prices (log)	4.985	0.242	4.698	4.768	4.983	5.171	5.343

Variable	Above Median Exposure to ATT						
	Mean	Std. Dev.	10th Percentile	25th Percentile	Median	75th Percentile	90th Percentile
<i>Panel A. On platform variables</i>							
Ad Spend (log)	2.659	2.440	0.000	0.000	2.914	4.622	5.870
Active Advertisers	1.295	1.211	0.000	0.000	1.000	3.000	3.000
Share Upper Funnel	0.919	0.251	0.800	1.000	1.000	1.000	1.000
Ad Spend per Active Advertiser	140.776	1,798.870	5.980	14.460	34.484	89.180	227.630
<i>Panel B. Off platform variables</i>							
Number of Establishments (log)	5.726	1.865	3.332	4.357	5.638	6.886	8.252
Prices (log)	5.146	0.387	4.799	4.888	5.169	5.285	5.330

Notes: Panels A and B present pre-ATT summary statistics for units in our treatment and control conditions. Summary statistics for all on- and off-platform variables are calculated over the course of our data and vary across units and time (quarter) pre-ATT. “Active Advertisers” is measured as the number of months each quarter an advertiser account ran a campaign. Spend is measured in USD.

4.2 Empirical Specification: The Differences-in-Differences Model

Our main specification takes the following standard form:

$$Y_{it} = \beta_0 + \beta_1(Treat_i \times Post_t) + \theta_t + \lambda_i + \epsilon_{it}, \quad (1)$$

where i represents an advertiser and t a time period (quarter). Y_{it} represents the outcome of interest (e.g., ad spend for advertiser i in period t). $Treat_i$ is an indicator variable which equals 1 for treated advertisers (above median ATT impact) and equals 0 for control advertisers (below median). Similarly, $Post_t$ equals 1 for post-ATT time periods and 0 for pre-ATT time periods. Finally, θ_t and λ_i are time and advertiser fixed effects, respectively, and ϵ_{it} is the error term. Our coefficient of interest is β_1 . For our analyses of prices and establishments, the units of analyses are industries and industry-states, respectively.¹⁹

Our identification relies on a parallel-trends assumption: that absent ATT, outcomes for the two exposure groups would have evolved similarly over time. Note that within each exposure group, advertisers may respond heterogeneously to ATT; the estimated interaction coefficient captures the difference in group-level average responses to ATT between more- and less-exposed advertisers.

This assumption could break if, for example, iOS usage is correlated with being rich or young and these demographics disproportionately cut back on spending post-ATT launch. In the data, the firms that were targeting rich or young people would thus appear to be more exposed to the ATT shock, but their drop in business would be due not to ATT. In the particular cases of income and age, these hypotheses do not appear consistent with broader evidence on spending trends during this era. For example, data from Kochhar and Sechopoulos (2022) and Lee et al. (2021) suggest lower income individuals experienced disproportionately worse economic outcomes during COVID and the recovery; similarly, data

¹⁹In January 2021, Meta emailed all of its advertisers and explicitly warned them that ATT would lower ad effectiveness. Even months prior to that, Apple had publicly stated ATT was coming. This creates a few different treatment regimes – the Apple warning, the Meta warning, and then the actual ATT launch. Advertiser behavior around the first two dates did not change meaningfully but was much more pronounced around the launch, suggesting minimal anticipatory effects. For this reason we treat the analysis as a binary launch in Q1 2021.

from Chetty et al. (2024) show that the sharpest pandemic-era spending contractions were driven by older households. Thus, the demographic groups most likely to own iOS devices were not the ones that disproportionately reduced spending in 2021, making it unlikely that income- or age-related macro trends explain our ATT results.

More broadly, we perform several analyses detailed later in the main text and in the Online Appendix to provide evidence on the plausibility of the parallel trends assumption in our context. In particular, we provide evidence on pre-trends, rerun our analyses dropping out each vertical, analyze effects across different decile splits of ATT exposure, and implement a continuous differences-in-differences estimator. We discuss these analyses in turn, but our results appear robust across these checks.

5 Results


In the next subsections, we report our main results on how the negative shock to digital advertising effectiveness impacted advertising behavior, the number of firms in different industries, and product prices.

5.1 Effects of ATT on Advertiser Behavior

We begin by examining whether ATT led to detectable shifts in advertiser behavior. We study changes in overall ad spend and activity levels and, conditional on being active, the types of campaigns advertisers run. We find that overall spend falls because fewer advertisers remain active, even though surviving advertisers increase their per-advertiser spending slightly. While this pattern may seem surprising, it is rationalized by our theoretical model in the Appendix. We also find that advertisers shift away from lower funnel campaigns that rely most on the data affected by ATT, consistent with Aridor et al. (2024).

Advertiser Spend. We first estimate our differences-in-differences specification with lagging and leading terms on advertising spend at the advertiser-account-quarter level, using an

Figure 4: Effects of ATT on (log) advertiser spend.




indicator for whether an account had above-median ATT exposure. As shown in Figure 4, we find evidence that total ad spend drops, declining for about a year until the coefficients level off at around a 10 percent relative reduction in ad spend. (The overall average post-ATT effect is about a 7 percent reduction, as shown in Table 2.) This result suggests a persistent, sizable impact of ATT on investments in digital advertising.²⁰

Active Advertisers. The reduction in advertiser spend could reflect changes along either the extensive or intensive margin—the former due to fewer active advertisers, and the latter due to lower spending per advertiser. To study the extensive margin, we rerun our differences-in-differences specification using a count of how many months in each quarter an advertiser

²⁰A common finding in several papers on the effects of privacy regulation on firms is that smaller businesses are hurt more than larger businesses (using varying definitions of ‘small’ and ‘large’; Korganbekova and Zuber (2023), Campbell et al. (2015), Wernerfelt et al. (2025)). We also run the analysis for the ‘large’ advertisers using the same internal classification that Meta uses. We find a nonsignificant effect on ad spend (+2.5%, $p = 0.816$), though the standard errors are wide. There is no significant difference versus the main result in the text, and at conventional levels we are only powered to rule out effects with a magnitude of 21%. Hence, large advertisers may be affected the same as small advertisers, but regardless, we provide evidence that small advertisers are substantially hurt.

Figure 5: Effects of ATT on advertiser activity.




was active on Meta (i.e., ran at least one campaign). Figure 5 similarly shows that activity levels drop, declining for about a year until settling. (Table 2 shows an average estimate of -0.0313 , which corresponds to a reduction of 3.7% of the reference point mean.)

Ad Spend per Active Advertiser. Despite decreases in the total ad spend and the activity levels of advertisers, we find that the average amount each active advertiser spent each quarter increased post-ATT (Figure 6). The average effect reported in Table 2 is an increase of 0.817 dollars per quarter, which is close to a 1 percent increase at the reference point mean.

Our market-expansion model in the Appendix can rationalize these results. The model assumes that when an advertiser invests more in advertising, they generate interest in the broader product category, creating opportunities for competing sellers to gain sales. This “public good” aspect of advertising implies that a decline in ad effectiveness leads more firms


Figure 6: Effects of ATT on ad spend per active advertiser.



to exit and stop advertising, while the remaining advertisers spend more because there is less free riding.

Share of Upper Funnel Campaigns. Turning to campaign type, recall that the data affected by ATT is most valuable for lower funnel campaigns that rely on offsite signals such as purchases or app downloads, whereas upper funnel campaigns use on platform signals (e.g., clicks or likes) and are not directly affected. Because not all advertisers use upper funnel campaigns, we restrict the sample to those that had run at least one such campaign in the two months before our sample begins. As shown in Table 2, the share of upper funnel-optimized campaigns increases by 0.44 percentage points ($p = 0.0019$) post-ATT. Figure 7 shows this effect over time, with coefficients trending upward, suggesting continued substitution beyond our observation window.

Figure 7: Effects of ATT on the share of upper funnel campaigns.



The average effects from our above results are presented in Table 2:

Table 2: Difference-in-differences estimates of ATT effects on advertiser behavior

	Ad Spend (log)	Activity Level	Spend per Active Advertiser	Share Upper Funnel Campaigns
DiD Estimate	-0.0741*** (0.0028)	-0.0313*** (0.0015)	0.8169** (0.3252)	0.0044*** (0.0013)
Fixed Effects	Account + Time	Account + Time	Account + Time	Account + Time
Adjusted R^2	0.582	0.5938	0.6703	0.713
N	8,268,051	9,560,606	868,912	621,310
Mean Dep. Variable	1.460	0.838	101	0.862

Notes: The table reports difference-in-differences estimates of the effects of ATT on advertiser behavior. All specifications include account and time fixed effects. Standard errors are clustered at the account level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Effects of ATT on the Number of Establishments

We now turn to our off platform results. Unlike the on platform data, the number of business establishments is available at the industry-state-quarter level. To get industry-state level data in iOS shares, we leverage the state-level locations of advertisers in our sample combined with our measurement of ATT impact. Specifically, we assign each advertiser account to a state based on the location of the account’s registered administrator; we then aggregate impact across all advertisers in a given industry-state to derive our measure of exposure.²¹ We can then compare industry-states that were above versus below median exposure in the following regression:

$$Y_{ist} = \beta_0 + \beta_1(Treat_{is}xPost_t) + \theta_t + \lambda_{is} + \epsilon_{ist},$$


where Y_{ist} is the log-transformed count of the number of establishments in industry i , state s , and quarter t . We also include fixed effects for both time and industry-state pairs as well as lags and leads of the $Post_t$ term for our main analysis.

Figure 8 presents the event study estimates for this analysis and shows a steady and continuing decline in the number of establishments in the most impacted state-industry pairs. As reported in Table 3, relative to state-industry pairs below median of ATT impact, those above the median experienced a 1.4% decline in the total number of active business establishments ($p < 0.001$).

To put these results in context, with about 6.8 million establishments in the U.S. across the industries in Q1 2021, a back of the envelope calculation suggests that the 1.4% decrease in the number of establishments equates to roughly 95,000 establishments. This is a sizable impact. We can use the on platform results from Section 5.1 together with public data released by Meta as a plausibility test of sorts. Recall that Table 2 shows a 3.7 percent reduction in the number of active advertisers in the most impacted group. This is, of course, a lower bound on the impact of moving from no ATT-impact to full ATT-impact. Next, in

²¹Some advertisers use agencies or other firms to run ads on their behalf. Hence, an advertiser in state-A may use a marketing firm in state-B who is identified by Meta as the advertiser. This introduces classical measurement error, which in turn creates an attenuation bias against finding significant results.

Figure 8: Effect of ATT on the number of establishments in the US.



their Q3 2020 earnings call, Meta (then Facebook) reported having more than 10 million active advertiser accounts worldwide.²² Assuming accounts are geographically distributed proportionate to advertising revenue, the concurrent regional revenue breakdown for the company would estimate around 5 million active advertisers in North America, which we can assume is mostly in the U.S. Our 3.7% effect would translate into roughly 185,000 accounts shutting down compared to the offline impact of 95,000 fewer establishments. Naturally, not all businesses that stopped advertising on Meta have shut down, making these two estimates consistent in order of magnitude.

Finally, we note that several economic models can rationalize this comparative static, including the simple one in our Appendix. Straightforwardly, if advertising helps increase demand for product markets, then if costs of advertising increase, market sizes decrease and firms exit. How large such effects may be—or if other advertising mechanisms (e.g., persuasion) dominate—is an empirical question; our results suggest that in our context, digital advertising plays a nontrivial market expansion role.


²²See <https://investor.atmeta.com/investor-events/event-details/2020/Facebook-Q3-2020-Earnings/default.aspx>

5.3 Effects of ATT on Prices

Our analysis of prices is at the industry-quarter level, meaning we return to our appropriately aggregated main specification (1). Our point estimate from Table 3 finds industries that were more affected by ATT saw their prices rise 3.5% relative to the less affected industries. The confidence interval is wide, however, and the point estimate is marginally significant at conventional levels ($p = 0.0519$).

Moving away from the average effect, Figure 9 presents the event study estimates and shows an upward trend in prices post-ATT for industries more affected by ATT relative to those less affected, which seems to stabilize after about one year. The slow increase in prices is consistent with the facts documented by Nakamura and Steinsson (2008). In their study of tariff- and cost-driven shocks in the coffee industry, retail prices adjust only about once every nine to ten months, implying that firms incorporate higher costs into prices gradually over several quarters. A similar dynamic would be expected when the relevant cost shock comes from higher advertising costs rather than input prices. If we estimate the impact on prices in 2022 and beyond (so incorporating interaction terms for being above median ATT exposure and the remaining months of 2021 vs. beyond), we find a increase of 4.3% ($p = 0.0402$), a substantial and economically meaningful magnitude.

Figure 9: Effect of ATT on prices in the U.S. (PPI).



Again, this increase can be rationalized via our model in the Appendix. Intuitively, if advertising effectiveness and investment fall, there can be weaker market-wide demand as fewer individuals learn about the products. However, this effect can be overpowered by the decrease in competition from the reduction of market size and firm exit, leading to net higher prices. This comparative static thus is consistent with our result on the number of establishments. More broadly, similar comparative static predictions could come from models of imperfect pass through of marketing costs or standard ‘informative’ (Bagwell, 2007) models of advertising.

Our results for the number of establishments and prices are shown in Table 3:

Table 3: Effects of ATT on number of establishments and prices.

	(log) Num. Establishments		(log) Prices (PPI)	
	Overall average	2022-on	Overall average	2022-on
DiD Estimate	-0.01427*** (0.0042)	-0.0205*** (0.0058)	0.0352* (0.0177)	0.04347** (0.0207)
Fixed Effects	State*Industry + Time		Industry + Time	
Adjusted R^2	0.9977	0.9977	0.9842	0.9844
N	50,280	50,280	742	742
Mean Dep. Variable at reference date	5.66	5.66	5.08	5.08

Notes: Standard errors clustered at the state–industry level for the establishments regressions and at the industry level for the price regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4 Robustness Checks

In the Online Appendix, we detail robustness checks on the above results. First, instead of comparing above versus below median ATT exposure, we rerun our analyses comparing the top versus bottom decile, quintile, etc. Second, to assuage concerns about differential time trends by industries that could be driving our results, we rerun all our analyses dropping out each vertical at a time. Third, we rerun our analyses with several different measures of ‘exposure’ to ATT: (i) the share of iOS users in advertisers’ pre-ATT target audiences, (ii) the share of ad spend pre-ATT that went to iOS users, (iii) the share of ad spend pre-ATT that

went to iOS users who opted out by August 2021, and (iv) the share of ad spend pre-ATT that went to iOS users who opted out by March 2023. Fourth, we test our comparative statics with a continuous differences-in-differences estimator from de Chaisemartin et al. (2025). Finally, we rerun our analyses choosing a different pre-ATT date to fix our panel of advertisers (June 2020 instead of July 2020). Our results are quite consistent across these checks.

6 Discussion

Digital advertising has rapidly grown to become the dominant marketing channel for a large share of businesses. By exploiting what is arguably the largest shock to digital advertising effectiveness to date—Apple’s rollout of App Tracking Transparency (ATT)—and combining detailed data from Meta with industry outcomes from the BLS, we shed light on how this channel shapes market structure. Our findings indicate that ATT led not only to material shifts in advertiser behavior but also to increased market concentration and higher consumer prices in the United States.

The anecdote with which we opened the paper underscores these dynamics. Alpha Paw, a company selling a niche CPG product—ramps for small dogs—relied heavily on pre-ATT targeting to reach a small, well-defined customer base. Because niche products benefit disproportionately from precise targeting, ATT’s reduction in signal quality sharply raised customer acquisition costs and eroded the firm’s profitability. This experience was widely echoed in the business press, where numerous small businesses reported meaningful declines in advertising efficiency, higher customer acquisition costs, and lower profits following ATT.²³

The harms to such firms matter because small businesses are widely viewed as a central engine of innovation, job creation, and economic dynamism. A large empirical literature documents the outsized role that young and small firms play in net job creation and productivity growth (e.g., Decker et al. (2014)). If privacy policies substantially reduce advertising

²³See, e.g., Mims (2021) and Haggin and Vranica (2021).

effectiveness—particularly for niche products or firms without established brands—these regulations may inadvertently dampen an important source of economic growth.

These concerns have not gone unnoticed by policymakers. Apple was fined by the French competition authority for aspects of ATT that were found to disadvantage advertisers relative to Apple’s own targeting capabilities, and other European regulators have initiated investigations into whether ATT undermines competition in digital advertising markets.²⁴ These developments reflect an emerging recognition that privacy-preserving design choices can have competitive implications, especially when implemented unilaterally by a dominant platform.

None of this implies that privacy protections lack value. Consumers may benefit from reduced tracking, greater control, and lower risks of data misuse. Yet the empirical literature has struggled to quantify these benefits in a comparable manner, and existing estimates of privacy valuations remain limited and context-dependent. Moreover, a large empirical literature has documented numerous unintended consequences of the GDPR—including reduced firm revenue, higher compliance costs, and increased market concentration—highlighted comprehensively in Dubé et al. (2025). Our findings therefore point to a broader research agenda: developing rigorous empirical methods for measuring both the costs and benefits of privacy regulation, and understanding how different forms of online activity—and the firms that rely on them—are differentially affected.

Overall, our results highlight a fundamental tradeoff at the core of modern privacy policy. Protecting consumer data may serve important social objectives, but these protections can also generate meaningful economic frictions. Quantifying this balance remains a central challenge for researchers and policymakers alike.

²⁴See, e.g., Coyer (2025).

References

- Akerberg, D. A. (2001). Empirically distinguishing informative and prestige effects of advertising. *RAND Journal of Economics*, pages 316–333.
- Anderson, S. P., Ciliberto, F., Liaukonyte, J., and Renault, R. (2016). Push-me pull-you: comparative advertising in the otc analgesics industry. *The RAND Journal of Economics*, 47(4):1029–1056.
- Apacible-Bernardo, A. and Bushey, K. (2025). Data protection and privacy laws now in effect in 144 countries. *International Association of Privacy Professionals*. <https://iapp.org/news/a/data-protection-and-privacy-laws-now-in-effect-in-144-countries/>.
- Aridor, G., Che, Y.-K., Hollenbeck, B., Kaiser, M., and McCarthy, D. (2024). Evaluating the impact of privacy regulation on e-commerce firms: Evidence from apple’s app tracking transparency. *SSRN Working Paper No. 4698374*.
- Athey, S., Catalini, C., and Tucker, C. (2017). The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research.
- Bagwell, K. (2007). The economic analysis of advertising. *Handbook of industrial organization*, 3:1701–1844.
- Bain, J. S. (1956). *Barriers to new competition: their character and consequences in manufacturing industries*. Harvard University Press.
- Benham, L. (1972). The effect of advertising on the price of eyeglasses. *The Journal of Law and Economics*, 15(2):337–352.
- Campbell, J., Goldfarb, A., and Tucker, C. (2015). Privacy regulation and market structure. *Journal of Economics & Management Strategy*, 24(1):47–73.
- Chetty, R., Friedman, J. N., and Stepner, M. (2024). The economic impacts of covid-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics*, 139(2):829–889.
- Chouaki, S., Bouzenia, I., Goga, O., and Roussillon, B. (2022). Exploring the online micro-targeting practices of small, medium, and large businesses. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–23.
- Coyer, C. (2025). France’s \$160 million apple fine puts app tracking on alert. Bloomberg Law. April 2, 2025.
- de Chaisemartin, C., Ciccia, D., d’Haultfoeulle, X., and Knau, F. (2025). Two-way fixed effects and differences-in-differences in heterogeneous adoption designs without stayers. *SciencesPo Institutional Repository*.
- Decker, R., Haltiwanger, J., Jarmin, R., and Miranda, J. (2014). The role of entrepreneurship in us job creation and economic dynamism. *Journal of Economic Perspectives*, 28(3):3–24.
- Demirer, M., Hernandez, D. J. J., Li, D., and Peng, S. (2024). Data, privacy laws and firm production: Evidence from the gdpr. *NBER Working Paper No. w32146*.
- Dixit, A. and Norman, V. (1978). Advertising and welfare. *The Bell Journal of Economics*, pages 1–17.
- Dubé, J.-P., Lynch, J. G., Bergemann, D., Demirer, M., Goldfarb, A., Johnson, G., Lambrecht, A., Lin, T., Tuchman, A., and Tucker, C. (2025). Frontiers: The intended and unintended consequences of privacy regulation for consumer marketing. *Marketing Science*.

- Dubois, P., Griffith, R., and O’Connell, M. (2018). The effects of banning advertising in junk food markets. *The Review of Economic Studies*, 85(1):396–436.
- Eckard Jr, E. W. (1991). Competition and the cigarette tv advertising ban. *Economic Inquiry*, 29(1):119–133.
- Feger, A. (2024). Digital makes up over three-quarters of total ad spend in the us. *eMarketer*.
- Fox, M. (2022). \$315 billion in market value has been erased from these 4 companies since apple’s ios privacy changes went into effect last year. *Business Insider*. <https://markets.businessinsider.com/news/stocks/facebook-meta-stock-apple-idfa-ios-privacy-change-social-media-2022-2>.
- Ghanavi, R., Hosseini, S., and Tucker, C. E. (2025). Privacy regulation and ad-tech consolidation. *Available at SSRN 5374198*.
- Gogi Gupta, Jordan Maddocks, A. P. . Z. K. (2022). ios update adoption history. *Gupta Media*.
- Goldberg, S. G., Johnson, G. A., and Shriver, S. K. (2024). Regulating privacy online: An economic evaluation of the gdpr. *American Economic Journal: Economic Policy*, 16(1):325–358.
- Goldfarb, A. and Tucker, C. E. (2011). Privacy regulation and online advertising. *Management science*, 57(1):57–71.
- Gordon, B. R., Moakler, R., and Zettelmeyer, F. (2023). Close enough? a large-scale exploration of non-experimental approaches to advertising measurement. *Marketing Science*, 42(4):768–793.
- Greenleaf, G. (2023). Global data privacy laws 2023: 162 national laws and 20 bills. <http://dx.doi.org/10.2139/ssrn.4426146>.
- Haggin, P. and Vranica, S. (2021). Apple’s privacy change is hitting tech and e-commerce companies. here’s why. *Wall Street Journal*. October 22, 2022.
- Hughes, J. T. and Saverice-Rohan, A. (2018). Iapp-ey annual privacy governance report. *International Association of Privacy Professionals*. Accessed: January 27, 2026.
- Janssen, R., Kesler, R., Kummer, M. E., and Waldfogel, J. (2022). Gdpr and the lost generation of innovative apps. *NBER Working Paper No. w32146*.
- Johnson, G. A., Shriver, S. K., and Goldberg, S. G. (2023). Privacy and market concentration: Intended and unintended consequences of the gdpr. *Management Science*, 69(10):5695–5721.
- Kochhar, R. and Sechopoulos, S. (2022). Covid-19 pandemic pinches finances of america’s lower- and middle-income families. Technical report, Pew Research Center.
- Korganbekova, M. and Zuber, C. (2023). Balancing user privacy and personalization. *Working paper*.
- Lee, S. Y. T., Park, M., and Shin, Y. (2021). Hit harder, recover slower? unequal employment effects of the covid-19 shock. Technical report, National Bureau of Economic Research.
- Lin, T. (2022). Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*, 41(4):663–681.
- Marshall, A. (1890). *Principles of Economics*. MacMillan and Co.
- Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.

- Milyo, J. and Waldfogel, J. (1999). The effect of price advertising on prices: Evidence in the wake of 44 liquormart. *American Economic Review*, 89(5):1081–1096.
- Mims, C. (2021). As apple and facebook clash over ads, mom-and-pop shops fear they’ll be the victims. Wall Street Journal. April 10, 2021.
- Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464.
- Ozga, S. A. (1960). Imperfect markets through lack of knowledge. *Quarterly Journal of Economics*, 74(1):29–52.
- Raggio, D., Ganapaneni, S., Ocampo, D., and Worsfold, S. (2025). 2025 small business advertising trends report. Intuit SMB MediaLabs https://digitalasset.intuit.com/render/content/dam/intuit/sbseg/en_us/Blog/Downloadable-assset/Ebook/2025-small-business-advertising-trends-report.pdf.
- Shapiro, B. T. (2018). Positive spillovers and free riding in advertising of prescription pharmaceuticals: The case of antidepressants. *Journal of Political Economy*, 126(1):381–437.
- Stigler, G. J. and Becker, G. S. (1977). De gustibus non est disputandum. *American Economic Review*, 67(2):76–90.
- Sutton, J. (1991). *Sunk Cost and Market Structure*. MIT Press.
- Tadelis, S., Hooton, C., Manjeer, U., Deisenroth, D., Wernerfelt, N., Dadson, N., and Greenbaum, L. (2023). Learning, sophistication, and the returns to advertising: Implications for differences in firm performance. Technical report, National Bureau of Economic Research.
- Telser, L. G. (1964). Advertising and competition. *Journal of Political Economy*, 72(6):537–562.
- Wernerfelt, N., Tuchman, A., Shapiro, B. T., and Moakler, R. (2025). Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science*, 44(2):268–286.

Appendix: Theoretical Framework

This Appendix contains a model of advertising and market structure that generates comparative static predictions consistent with what we observe in the data. We include this model to highlight plausible theoretical assumptions and settings that would predict our results. Separately, our Online Appendix contains more details on our industry matching process and all the robustness checks detailed in Section 5.4.

Inspired by the seminal work of Sutton (1991), we view advertising, which raises the demand of consumers for a firm’s products, as an *endogenous sunk cost* that firms choose to invest, and in return it provides a more robust demand for the firm’s products. We develop a toy-model following the three-stage game introduced in Chapter 3 of Sutton (1991); the stages are as follows:

1. Many potential firms consider entering a focal good market at entry cost $k > 0$.
2. Given the entry of N firms, each firm i that entered chooses a level of advertising $a_i \geq 0$ at a symmetric cost function $\frac{1}{2}a_i^2$. Advertising by firms increases the size of the market, and demand is given by the linear inverse demand function $p(Q) = R - b(A, e)Q$ where R is the market “reservation” price above which no consumer purchases, Q is the aggregate amount of product sold to consumers, and $b(A, e) > 0$ is decreasing in total advertising by the firms, A , and in advertising effectiveness, e .²⁵ In essence, for any price p , a higher level of aggregate advertising increases the quantity demanded at that price. For simplicity, we adopt the following functional form of demand,

$$p(Q) = R - \frac{Q}{e \sum_{i=1}^N a_i}$$

²⁵We depart from Sutton (1991) by assuming that the total size of the market is determined by the amount of advertising that the firms engage in. In particular, when a firm invests in advertising it increases the size of the market for all firms. This captures the idea that when a firm’s ad piques a consumer’s interest, the consumer may search for similar products and eventually purchase from another seller.

3. After the choices of $\{a_i\}_{i=1}^N$ are fixed, the firms engage in Cournot competition. Each chooses quantity q_i , with a symmetric marginal cost of production normalized to $c = 0$.

We solve for a sequential equilibrium working backwards. In the third stage, both entry and advertising costs are sunk and each firm maximizes the following *ex post* profit function,

$$\pi_i(q_i, q_{-i}) = p(Q)q_i = \left(R - \frac{q_i + \sum_{j \neq i} q_j}{e \sum_{i=1}^N a_i}\right)q_i.$$

Taking the FOC yields the following optimality condition,

$$R - \frac{2q_i + \sum_{j \neq i} q_j}{e \sum_{i=1}^N a_i} = 0.$$

We focus on symmetric equilibrium and denote by q^N the equilibrium quantity of each of the N firms, which in turn yields the following solution for the third stage of the game:

$$q^N = \frac{eR \sum_{i=1}^N a_i}{N+1} \quad \text{and} \quad Q^N = \frac{eRN \sum_{i=1}^N a_i}{N+1}. \quad (\text{A.1})$$

Turning to the second stage of the game in which firms choose advertising levels a_i , each firm maximizes its *interim* profits given their rational expectations of the third-stage solution (entry costs are sunk). Substituting (A.1) into the second stage profit function yields,

$$\begin{aligned} \pi_i(a_i, a_{-i}) &= \left(R - \frac{Q^N}{e \sum_{j=1}^N a_j}\right)q^N - \frac{1}{2}a_i^2 = \left(R - \frac{\frac{eRN \sum_{i=1}^N a_i}{N+1}}{e \sum_{i=1}^N a_i}\right)\frac{eR \sum_{i=1}^N a_i}{N+1} - \frac{1}{2}a_i^2 \\ &= \frac{eR^2}{(N+1)^2} \sum_{j=1}^N a_j - \frac{1}{2}a_i^2 \end{aligned} \quad (\text{A.2})$$

and taking the FOC yields the following optimality condition,

$$\frac{R^2 e}{(N+1)^2} - a_i = 0.$$

Focusing on symmetric equilibria we denote by a^N the equilibrium advertising level of each of the N firms, which yields the following solution for the second stage of the game:

$$a^N = \frac{eR^2}{(N+1)^2} \quad (\text{and total advertising is equal to } A^N = \frac{eR^2 N}{(N+1)^2}). \quad (\text{A.3})$$

We are now ready to state the free entry condition of zero profits *ex ante* (before entry costs are sunk) to compute the number of firms who will enter in equilibrium. If N firms enter, then the profits anticipated by these N firms can be expressed by substituting advertising levels in (A.2) with the equilibrium solution given in (A.3) to obtain,

$$\begin{aligned} \pi_i^N &= \frac{eR^2}{(N+1)^2} \sum_{j=1}^N a_j - \frac{1}{2}a_i^2 - k = \frac{eR^2}{(N+1)^2} \cdot \frac{eR^2 N}{(N+1)^2} - \frac{1}{2} \cdot \left(\frac{eR^2}{(N+1)^2} \right)^2 - k \\ &= \frac{e^2 R^4}{(N+1)^4} \cdot \left(N - \frac{1}{2} \right) - k. \end{aligned}$$

Hence, we can write the zero-profit condition as,


$$e^2 R^4 \left(N - \frac{1}{2} \right) = k(N+1)^4. \quad (\text{A.4})$$

This quadric equation does not have an analytic solution, but our simple model admits the following useful result:

Proposition A.0.1. *A unique free entry equilibrium exists if and only if $e^2 R^4 \geq 32k$. Furthermore, the number of firms in a free-entry equilibrium is (weakly) increasing in e .*

Proof. To prove the first part of the Lemma, evaluate equation (A.4) at $N = 0$: the LHS is negative while the RHS is zero. Now evaluate equation (A.4) at $N = 1$: the two sides would equal each other if and only if $e^2 R^4 \frac{1}{2} = 16k$. Furthermore, the slope of the LHS is an increasing linear function of N , $e^2 R^4$, while the slope of the RHS is an increasing convex function of N , $4k(N+1)^3$. If at $N = 1$ the slopes of the LHS and RHS are equal, then $e^2 R^4 = 32k$, and multiplying each side by $\frac{1}{2}$ yields $e^2 R^4 \frac{1}{2} = 16k$. This implies that when $e^2 R^4 = 32k$, the LHS and RHS meet only at $N = 1$, which is the unique free-entry equilibrium. If $e^2 R^4 < 32k$, then the LHS is less than the RHS for any value of N and hence the two sides of A.4 never cross, implying no free entry equilibrium. It follows immediately that if

Figure A.1: Free Entry Equilibrium Number of Firms



$e^2 R^4 > 32k$, then there are two intersections of the two functions, one at a value $N < 1$ and one at a value $N > 1$. It is also easy to see that of these two intersecting points, the only one that is a stable free-entry equilibrium has $N > 1$. The second part of the proposition follows immediately from the fact that the slope of the LHS is increasing in e and hence, when e increases, the intersection of the two functions must occur at higher values of N . $\square$

Proposition A.0.1 can be seen graphically. Consider a numerical example in which $R = 10$ and $k = 40$, and imagine a comparative static analysis in advertising effectiveness, e . From Proposition A.0.1, a necessary condition for the existence of a free entry equilibrium is $R^4 e^2 \geq 32k$, and substituting $R = 10$ and $k = 40$, this becomes $e > \sqrt{0.128} \approx 0.358$. Hence, for any $e > 0.358$, there is a unique free entry equilibrium.

To see how changes in e impact the equilibrium number of firms, Figure A.1 plots the linear LHS $e^2 R^4$ with $R = 10$ at three different values of advertising effectiveness: $e = 1$, $e = 2$, and $e = 3$. It also plots the convex RHS function $k(N + 1)^4$ with $k = 40$. As the plot demonstrates, there is one solution at values of $N < 1$, which is not an equilibrium as one firm will always profit from entering under these conditions. The second solution occurs at values of $N > 1$, and in particular, the free entry equilibrium number of firms are $N = 4, 8$, and 11 for $e = 1, 2$, and 3 respectively.

Proposition A.0.2. *In the unique free entry equilibrium, as advertising effectiveness e increases, (i) price decreases; (ii) and total quantity produced increases.*

Proof. Equilibrium price is given by,

$$p^* = R - \frac{Q^N}{e \sum_{i=1}^N a_i} = R - \frac{\frac{eRN \sum_{i=1}^N a_i}{N+1}}{e \sum_{i=1}^N a_i} = \frac{R}{N+1},$$

where the second equality follows from the solution to the third stage of the game in equation (A.1) above. The comparative statics in Proposition A.0.1 establish that N is increasing in e , and it therefore follows immediately that p (weakly) decreases in e because it decreases in N . Turning to the total quantity produced,

$$Q^N = \frac{eRN \sum_{i=1}^N a_i}{N+1} = R^3 e^2 \frac{N^2}{(N+1)^3}.$$

Clearly, $R^3 e^2$ is increasing in e , but the derivative of $\frac{N^2}{(N+1)^3}$ in N is $\frac{N(2-N)}{(N+1)^3}$, which is negative for $N > 2$, making it difficult to sign this derivative in e . However, because an increase in e increases the incentives to expand the market, and prices are (weakly) decreasing in e , it must follow that total quantity is increasing in e . $\square$

That total quantity produced increases as advertising becomes more effective is anticipated because the market expands and more firms enter and compete to supply this market. Second, price *decreases* as advertising becomes more effective. This is due to the fact that the increase in competition from entry more than compensates for the stronger market demand. Last, calculations in ?? show that as advertising becomes more effective and the number of firms increase, each firm invests *less* in advertising while *aggregate* advertising by all firms increases. With a fixed number of firms, each firm would advertise *more* as advertising effectiveness increases due to the complementarity between advertising effectiveness and investment in advertising. In equilibrium, however, this market expansion invites more entry, and as the number of firms grows in response to more effective advertising (and hence, a larger market), there is more free-riding among the advertisers who enter.

By making advertising less effective, the impact of ATT is akin to a decrease in e our model. Hence, the analysis above offer three testable predictions: Following ATT,

1. Total advertising will decrease (A^* decreases)
2. Firms will exit the market (N^* decreases)
3. Product prices will increase (p^* increases)

Despite the fact that e is continuous, the comparative statics of a^* and A^* cannot be derived continuously because N^* is discrete, meaning that the marginal impact of e on a^* and A^* jumps discontinuously whenever the value of e crosses a threshold that changes N^* .

Recall from Proposition A.0.1 that N^* is increasing in e . Let $e_L(N)$ denote the lowest value of e for which $N^* = N$, and let $e_H(N)$ denote the highest value of e for which $N^* = N$. Hence, for any value of $e \in [e_L(N), e_H(N)]$, there will be exactly $N^* = N$ firms in equilibrium. Note that by definition, $e_H(N) = e_L(N + 1)$.

Recall that the solution to the advertising stage in equation (A.3) implies that, *fixing* N , a^N (the optimal choice of a , anticipating the third-stage Cournot game with N firms in the final stage) is increasing linearly in e with a slope $R^2/(N + 1)^2$. Hence, as long as we increase e *within* any of the intervals $[e_L(N), e_H(N)]$, N^* does not change and a^N increases linearly in e within a well defined interval $[a_L(N), a_H(N)]$. Furthermore, since $A^N = Na^N$, it too is increasing in e within a well defined interval $[A_L(N), A_H(N)]$. However, once e rises to $e_H(N) + \epsilon$ (which is equal to $e_L(N + 1) + \epsilon$), we have a new equilibrium with $N + 1$ firms. This discontinuous jump in N causes a discontinuous drop in the marginal value of ads, implying a discontinuous drop in both a^N and A^N .

Figures A.2 and A.3 plot the comparative statics described above. with a simulated example that sets $R = 10$ and $k = 40$, letting e vary. Using equation A.4 and setting a specific value of N , we compute the boundaries of the intervals $[e_L(N), e_H(N)]$ for each N . Then, using equation (A.3) we compute the intervals $[a_L(N), a_H(N)]$ and $[A_L(N), A_H(N)]$, which correspond to the increasing piecewise-linear best response functions for a^* and A^* as e varies within the defined intervals $[e_L(N), e_H(N)]$. Both a^N and A^N follow a “saw-tooth” function due to the forces described above.

Figure A.2: Comparative Statics for a^*





Figure A.3: Comparative Statics for A^*



As Figure A.2 and Figure A.3 show, along the piecewise linear parts of the functions, both a^* and A^* are increasing due to the increasing marginal returns to advertising. However, at the jump points, there is a drop in both a^* and A^* , creating local non-monotonicities. If, instead, we look at the midpoints of the intervals as a summary statistic of the values of a^* and A^* for different values of N^* , then we obtain clear monotonic patterns: as advertising becomes less effective so as to reduce the number of equilibrium forms N^* (a nontrivial decrease in e), then total advertising A^* decreases while the level of advertising of remaining firms a^* increases, consistent with what we observe in the data.

Online Appendix

In this Online Appendix, we provide details on our industry matching process as well as robustness checks detailed in Section 5.4.

OA.1 Matching Meta Industries to NAICS Codes

To merge our on platform and off platform data, we went through all the 6-digit NAICS codes and Meta’s 176 industries, and manually matched them based on their descriptions. In doing so, there were several kinds of possible matches; below we walk through examples of each kind and describe how we did the join.

One-to-one. Meta’s industry classification description for “Advertising and Marketing - Buying Agency” maps to a single NAICS industry, “NAICS 541830: Media buying agencies.” In this case, for our measure of ATT exposure for the BLS data we took the sum of all impressions from advertisers in Meta’s “Advertising and Marketing - Buying Agency” industry that went to iOS users and divided by their total number of impressions. The BLS data is also straightforward to integrate here as it lives on the same level.

One-to-many. Meta’s industry classification description for “Entertainment and Media - Movies” was mapped to five different NAICS industries (e.g., NAICS 512110: Motion Picture and Video Production and NAICS 512131: Motion Picture Theaters (Except Drive-Ins)). In this case, for our measure of ATT exposure, we aggregate the impressions on the Meta industry side as above. For the BLS data, we sum all the establishments across the five NAICS codes. This leaves us with one ‘synthetic’ vertical we call “Entertainment and MediaMovies” that has both the ATT exposure from advertisers in Meta’s ‘Entertainment and Media - Movies’ industry and BLS data from the five NAICS codes.

Many-to-one. Meta’s classification system has seven different periodical publishing industries (e.g., “News and Current Events,” “Beauty/Fashion,” and “Food”); these were mapped onto NAICS 513120: Periodical publishers. In this case, for the measure of ATT exposure, we take the total number of iOS impressions from advertisers in all of these industries and divide by their total number of impressions; for the BLS data for the new synthetic vertical,

we can take apply it directly.

Many-to-many. Finally, two of Meta’s internal industries “Education - Not-for-Profit Colleges and Universities” and “Education - For-Profit Colleges and Universities” map to three different NAICS industries. In this case, we again aggregate as described earlier across all advertisers in the Meta industries to derive the synthetic vertical’s measure of ATT exposure. We similarly aggregate the BLS data, creating one synthetic vertical that has the combined measures of both ATT exposure and BLS outcomes.

After conducting this matching process, we conducted manual checks where we took a sample of registered businesses under a given NAICS code, identified whether they advertised on Meta, and then confirmed their internal industry was the right one given our assignment. Below we provide a sample of more matches (the entirety would be dozens of pages); note how there are one-to-one and many-to-one matches:

Table OA.1: Examples of Meta Industries, 6-digit NAICS codes, and the merged industry definitions

Meta Industry	NAICS Description	Merged Industry Name
Travel, Railroads	NAICS 482111 Line-haul railroads	TravelRailroads
Travel, Railroads	NAICS 482112 Short line railroads	TravelRailroads
Travel, Hotel and Accommodation	NAICS 721110 Hotels (except casino hotels) and motels	TravelHotel and Accommodation
Travel, Hotel and Accommodation	NAICS 721191 Bed-and-breakfast inns	TravelHotel and Accommodation
Travel, Hotel and Accommodation	NAICS 721199 All other traveler accommodation	TravelHotel and Accommodation
Travel, Hotel and Accommodation	NAICS 721211 RV (recreational vehicle) parks and campgrounds	TravelHotel and Accommodation
Travel, Hotel and Accommodation	NAICS 721214 Recreational and vacation camps (except campgrounds)	TravelHotel and Accommodation
Travel, Tourism and Travel Services	NAICS 561520 Tour operators	TravelTourism and Travel Services
Travel, Auto Rental	NAICS 532111 Passenger car rental	TravelAuto Rental
Travel, Auto Rental	NAICS 532120 Truck, utility trailer, and RV (recreational vehicle) rental and leasing	TravelAuto Rental

One complication in the matching process is that NAICS codes are updated every five years across North America to reflect changes in the economic landscape. The most recent update occurred in 2022, which is within the timeframe of our analysis.²⁶ Many of these updates involved only changes to the NAICS code numbers, without altering the underlying industry definitions. However, some updates were more substantive. In certain cases, the NAICS revision process led to the reclassification of industries—either by merging multiple previously distinct industries into a single new category, or by splitting an existing industry into several new, more specific categories. For example, in the 2017 NAICS classification, Anthracite Mining (NAICS: 212113), Bituminous Coal and Lignite Surface Mining (NAICS: 212111), and Bituminous Coal Underground Mining (NAICS: 212112) were reclassified in 2022 as Underground Coal Mining (NAICS: 212115) and Surface Coal Mining (NAICS: 212114).

²⁶Additional details of the code updates may be found here: https://data.bls.gov/cew/apps/bls_naics/NAICS2022_graphics_PDF.pdf


A key challenge is that, for industries affected by these re-classifications, the relative contributions of the original industries to the new categories are not always available. For instance, in the example above, while all of Bituminous Coal and Lignite Surface Mining is now classified as Surface Coal Mining, and all of Bituminous Coal Underground Mining is now classified as Underground Coal Mining, it is unclear how Anthracite Mining is divided between the two new categories. To address these challenges and maintain consistency throughout the period of our analysis, we harmonize the data across years by applying the same re-classifications historically to the pre-2022 data. Specifically, if a set of NAICS industries were merged into a single category in the post-2022 data, we aggregate those same industries in the pre-2022 data.

Finally, we note that one of the new industries (Mining and Quarrying) had a substantially lower ATT exposure than any other new industry (outside $1.5 \times \text{IQR}$). Hence, we omitted that industry from our analyses, but including it does not change our results meaningfully.


OA.2 Additional percentile splits

In the main text, we compare outcomes for units that are above versus below median exposed to ATT. To explore sensitivity of our result, below for each of our main outcomes, we rerun our main specification comparing the top vs. bottom 10%, 20%, 30%, and 40% (including our above vs. below median as well for comparison). These are different estimands and the results in the main text may or may not carry over to these different populations (for example, those severely impacted may behave very differently than those who are simply right above the median). We find, however, that our comparative statics are quite consistent across the quantile cuts. Note that in the results for the number of establishments and prices, the standard errors often grow as expected when the sample size shrinks, but the point estimates remain relatively stable.


Figure OA.1: Effects of ATT on outcomes by quantile splits




(a) Ad Spend




(b) Active Advertisers




(c) Share Upper Funnel



(d) Ad Spend per Active Advertiser



(e) Number of Establishments



(f) Prices

Table OA.2: Effects by exposure contrast (Top vs. Bottom quantiles)

	Top/Bottom 50%	Top/Bottom 40%	Top/Bottom 30%	Top/Bottom 20%	Top/Bottom 10%
Ad Spend	−0.0741*** (0.0028)	−0.0907*** (0.0031)	−0.1277*** (0.0035)	−0.1870*** (0.0041)	−0.3245*** (0.0054)
Active Advertisers	−0.0313*** (0.0015)	−0.0383*** (0.0017)	−0.0511*** (0.0019)	−0.0681*** (0.0023)	−0.1282*** (0.0031)
Share Upper Funnel	0.0044*** (0.0013)	0.0057*** (0.0015)	0.0074*** (0.0018)	0.0087*** (0.0022)	0.0085*** (0.0031)
Ad Spend per Active Advertiser	0.0219*** (0.0027)	0.0318*** (0.0030)	0.0415*** (0.0036)	0.0617*** (0.0045)	0.0674*** (0.0069)
Number of Establishments	−0.0143*** (0.0042)	−0.0165*** (0.0047)	−0.0187*** (0.0054)	−0.0235*** (0.0070)	−0.0158 (0.0108)
Prices	0.0352* (0.0177)	0.0475** (0.0184)	0.0461** (0.0222)	0.0389 (0.0263)	0.0421 (0.0390)

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

OA.3 Leave-one-out verticals

A concern with our differences-in-differences specification may be differential trends by treatment assignment. To mitigate this, we reran all our main specifications below where for each outcome variable we drop each industry. We find our results are quite stable. Given the number of regressions here, we present results only in figure form – in each row, the listed industry is the one that is omitted (and as in the main text, ad spend, the number of establishments, and prices are all log transformed).

Figure OA.2: Leave-one-out Industry Analysis: Ad Spend




Figure OA.3: Leave-one-out Industry Analysis: Share upper funnel campaigns.




Figure OA.4: Leave-one-out Industry Analysis: Number of Active Advertisers.




Figure OA.5: Leave-one-out Industry Analysis: Ad Spend per Active Advertiser.




Figure OA.6: Leave-one-out Industry Analysis: Number of establishments.




Figure OA.7: Leave-one-out Industry Analysis: Prices.



OA.4 Varying measures of ATT exposure


There are several ways one could define a measure of ‘exposure’ to ATT. Below we provide both graphical (Figures OA.8-OA.10) and numerical (Table OA.3) evidence on how our results change under different definitions of our exposure metric. Specifically, we rerun our main analyses where we define exposure by:²⁷

- The share of pre-ATT impressions that went to iOS users (main text)
- The share of pre-ATT users who saw ads from the focal advertiser who were iOS users
- The share of pre-ATT ad spend that went to iOS users
- The share of pre-ATT ad spend that went to users who would opt out of tracking by 8/1/2021
- The share of pre-ATT ad spend that went to users who would opt out of tracking by 3/1/2023


Due to data retention issues, we could not run the last two exposure measures on the number of establishments and prices. Reassuringly, however, across the runs we do we find our results are quite consistent across different exposure definitions.

²⁷All these definitions are provided at the advertiser level; when the regressions required aggregation they were aggregated in the same manner as described in the main text. For consistency, these were all gathered based on the same date and sample as our main exposure metric.

Figure OA.8: Ad spend and activity levels.




(a) Ad Spend




(b) Active Advertisers

Figure OA.9: Share upper funnel and ad spend per active advertiser.




(a) Share Upper Funnel




(b) Ad Spend per Active Advertiser

Figure OA.10: Number of establishments and prices.



(a) Number of Establishments



(b) Prices

Table OA.3: Sensitivity of main results to alternative ATT exposure measures

	iOS Impressions	iOS Users	iOS Spend	Spend on users who opted out by 8/21	Spend on users who opted out by 3/23
Ad Spend	-0.07414*** (0.002774)	-0.07634*** (0.002774)	-0.07224*** (0.002774)	-0.07490*** (0.002773)	-0.08086*** (0.002775)
Active Advertisers	-0.03130*** (0.001514)	-0.03272*** (0.001514)	-0.03120*** (0.001514)	-0.03272*** (0.001514)	-0.03597*** (0.001515)
Share Upper Funnel	0.004396*** (0.001327)	0.004740*** (0.001327)	0.003958*** (0.001327)	0.004570*** (0.001328)	0.004658*** (0.001328)
Ad Spend per Active Advertiser	1.541*** (0.3276)	1.257*** (0.3276)	1.728*** (0.3277)	1.577*** (0.3281)	0.8831*** (0.3256)
Number of Establishments	-0.01427*** (0.004178)	-0.01887*** (0.004174)	-0.01600*** (0.004177)		
Prices	0.03517** (0.01767)	0.03486** (0.01769)	0.04296** (0.01736)		

Notes: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

OA.5 Continuous differences-in-differences

In the main text, we study heterogeneity in exposure to ATT using an above versus below median design. Under the assumption that high- and low-exposure units would have followed parallel trends absent ATT, this approach identifies post-intervention differences in average treatment effects. We emphasize this design because it provides a clear and interpretable way to assess how changes in digital advertising effectiveness affect market structure.

In this setting, another quantity of interest is the average change in outcomes when moving from no ATT exposure to the treatment dose actually received. Such a quantity may be particularly relevant for an ex post program evaluation of Apple’s policy. A complication is that ATT was a global shock: for example, the range of exposure for our off-platform industries spans from 0.4588 to 0.7574, so we observe neither untreated units nor units with close to zero exposure. Nevertheless, under the conditions of de Chaisemartin et al. (2025), it is possible to identify the sign of a weighted average of slopes (WAS) of units’ potential outcomes, which summarizes the average outcome change per unit of exposure relative to no ATT.²⁸ In this subsection, we implement their estimator and assess whether we can provide evidence on the sign of the WAS for our outcomes of interest.

The key identifying assumption underlying this approach is that, absent ATT, units with the lowest exposure and the full population would have followed the same average outcome trends. Although this assumption is fundamentally untestable, de Chaisemartin et al. (2025) propose a diagnostic pre-trends test that compares changes in outcomes for the least-exposed units to those for the full sample in periods prior to ATT. Implementing this test, we reject the null of parallel pre-trends in four of the six outcomes: ad spend, the number of active advertisers, and ad spend per active advertiser ($p < 0.001$). By contrast, we do not reject parallel pre-trends for share upper funnel or prices ($p=0.2917$ and 0.2379 , respectively). We therefore interpret the results of this exercise with caution, given this evidence of potential pre-trend violations – particularly relative to the above-below median analysis in the main text, where pre-trends appear more stable.

In Table OA.4 below we report the test statistic from de Chaisemartin et al. (2025) that is positive if and only if the WAS is positive (Theorem 2). We conduct inference via bootstrap, allowing us to test whether we find significant evidence of positive or negative WAS.

²⁸The authors also provide assumptions under which related estimands are point identified. We focus on the sign result given our interest in comparative statics and our view that the additional assumptions required for point identification are unlikely to hold in our setting (e.g., that the least-treated units and the full population experience the same effect at the lowest dose).

Table OA.4: WAS sign test results.

Outcome	Test statistic
Ad spend	-0.8295^{***} (0.1308)
Active advertisers	-0.4096^{***} (0.0746)
Share upper funnel	0.0227 (0.0273)
Ad spend per active advertiser	4.3104 (18.658)
Number of Establishments	-0.10626 (0.1190)
Prices	0.8709^{***} (0.2383)

Notes: de Chaisemartin et al. (2025)'s result is for a two-period panel; we use Q1 2021 and Q1 2022 as the pre- and post-periods, but running the analysis with other time periods yields similar results. As in the main text, ad spend, the number of establishments, and prices are log transformed.

We note that all of the signs of the WAS are consistent with those observed in our main text. For ad spend and advertising activity levels we find evidence of a significant and negative WAS; for prices it is significant and positive. For the number of establishments, we also find evidence of a negative WAS, though the significance level is 0.106 in a one-sided test. Our estimates for the share of upper funnel and the ad spend per active advertiser are directionally the same as in the main text but not statistically significant. We note that measurement error in the exposure variable would attenuate the estimated effects, and that the estimator's slow convergence rates may reduce power in smaller samples (e.g., for the number of establishments).


Although this estimand differs from our baseline differences-in-differences estimand, the conclusions are directionally consistent. Subject to the caveat of potential pre-trend violations, both approaches suggest that ATT reduced advertising spending and activity, increased prices, and led to (directionally) fewer establishments.

OA.6 Different panel from another pre-ATT date


Our main analysis was based on a panel of advertisers pulled on July 31, 2020. We used one day since computationally, to go through the universe of US advertisers on Meta and all their ad impressions was quite expensive. For robustness, however, we also repeated the exercise on July 30, 2020 and then reran our main analysis with that completely separate panel of advertisers.

We note that due to data retention issues, we were only able to analyze our on platform outcomes and our exposure metric here is the advertiser-level share of ad spend that went to users who could opt out of tracking by August 2021. Hence, it is different than the measure used in the main text, but as shown in Section OA.4 the results are similar to those in the main text. Below we see that for the panel of advertisers pulled in June and July, we observe quite similar outcomes across our four main on platform outcomes. This helps provide assurance that our results are not due to idiosyncratic aspects of our sample.


Figure OA.11: On platform results from different sample selection




(a) Ad Spend



(b) Active Advertisers



(c) Share Upper Funnel



(d) Ad Spend per Active Advertiser