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

Authors are listed alphabetically. We are grateful to numerous colleagues and coworkers for comments on past iterations of this work and to Nadav Tadelis for detailed comments on an earlier draft. All errors remain our own. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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
JEL No. D22,D40,L10,L59,M38

ABSTRACT

Digital advertising, which uses consumer data to target ads to users, now accounts for most of global ad expenditures. Privacy concerns have prompted regulations that restrict the use of personal data. To inform these policy debates, we develop an equilibrium model of advertising and market structure to analyze the impact of privacy regulation on market outcomes. We test the model's predictions using the launch of Apple's App Tracking Transparency feature, which created a natural experiment that limited the use of consumer data. Leveraging data from all U.S. advertisers on Meta combined with offline administrative data, we find that reductions in digital ad effectiveness led to decreases in investments in advertising, increases in market concentration, and increases in product prices. These effects are economically meaningful in magnitude and suggest potential harms to both firms and consumers from privacy regulation.

Daniel Deisenroth
Meta Platforms
1 Hacker Way
Menlo Park, CA 94025
dbdeisenroth@meta.com

Utsav Manjeer
Meta Platforms
1 Hacker Way
Menlo Park, CA 94025
utsavmanjeer@meta.com

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

Steven Tadelis
Haas School of Business
University of California, Berkeley
545 Student Services Building
Berkeley, CA 94720
and NBER
stadelis@haas.berkeley.edu

Nils Wernerfelt
Kellogg School of Management
2211 Campus Drive
Evanston, IL 60208
nils.wernerfelt@gmail.com

1 Introduction

Digital advertising has become the primary way through which businesses reach and attract customers, its popularity fueled by using detailed user data to match consumers to ads. A recent survey of small and medium sized businesses in the US found that 82% report digital advertising is 'crucial' for the success of their business.[1] Instead of buying ad slots based on broad demographics (e.g., TV or radio ads), online ads can be targeted to consumers who visited a firm's website, put an item in an online shopping cart, or even who "look like" consumers who behaved in these ways.

At the same time, the use of detailed consumer data has engendered concerns from privacy advocates (e.g., Zuboff (2023)), which have stirred both regulators and firms to limit the use of personal data online. Several past regulations and product changes were aimed at limiting the collection and use of user data (e.g., GDPR, CCPA, Apple's App Tracking Transparency feature), and more loom (e.g., updated versions of the GDPR in India and Brazil, deprecation of third party cookies in Chrome). Collectively, these changes have–and will–impact millions of firms and billions of dollars in ad expenditures (or "ad spend").

Like with most policies, there are winners and losers. Individual users are the intended winners; these policies claim to protect them. It's unclear, however, how much users benefit because consumer behavior seems inconsistent with a strong preference for privacy—a tension often called the "privacy paradox" (Athey et al., 2017). The clear losers are platforms that rely on collection and use of user data to offer digital advertising services (Hackett and Harty, 2021). But how do these policies impact firms' abilities to reach potential customers? And how does that impact the markets in which these firms operate?

To shed light on these questions, we develop an industry model of firms who advertise to increase demand, and study the impact of a deterioration in the ad effectiveness. We then exploit a major shock to the collection and use of user data to test the model's implications. Specifically, the introduction of Apple's App Tracking Transparency (ATT) feature, which

[1] https://www.statista.com/statistics/1388013/importance-digital-advertising-success-smbs-usa/

limits the use of user data, offers a natural experiment through which to evaluate the impact of stronger privacy protections on outcomes of firms who use digital advertising.

We begin by describing the institutional details of advertising on Meta and how it was impacted by ATT in Section 2. Next, inspired by the seminal "endogenous sunk-costs" approach of Sutton (1991), we develop an equilibrium model of advertising and market structure in Section 3. Our model's comparative statics show that a negative shock to advertising effectiveness—the consequence of ATT—leads to reductions in advertising by consumer-facing firms, exit of some firms, and increases in product market prices for consumers. We test the model's predictions using data—described in Section 4—on the universe of U.S. advertisers on Meta from early 2020 until early 2023, paired with data from the Bureau of Labor Statistics on prices and firm counts.

Specifically, as described in Section 5, we leverage historical user data from Facebook (FB) and Instagram (IG) to identify exogenous variation in the loss of user data across both advertisers and industries on Meta. Namely, the share of users for whom data was lost depends on the share of iOS versus Android users that an advertiser wishes to target. This allows us to employ a difference-in-differences approach to test our model's predictions regarding the impact of ATT on our economic outcomes of interest, both on- and off-Meta's platform. We rank advertisers and industries by the share of users for whom data was lost and then compare outcomes for advertisers and industries who were more versus less impacted. Our analysis relies on standard assumptions that we discuss, and we offer several robustness checks to buttress our analyses. Though our data are significantly more detailed, this general identification strategy is similar to what recent related work on ATT has utilized (e.g., Aridor et al. (2024); Bian et al. (2023))

Our outcome data on the number of firms and product prices are available at the quarter-6 digit NAICS level. To map our on-platform data to the corresponding level, we use Meta's internal classification of advertisers by industry-vertical. By taking a weighted aggregation described in Section 5, we are able to leverage the on-platform Meta data to derive estimates of ATT exposure at the 6 digit NAICS level. With this crosswalk, we can then complete our analysis of market concentration and prices.

Section 6 presents our empirical analyses that confirm our model's predictions. On-platform results show that advertisers did indeed respond to the loss of user data: Advertisers that were more affected by ATT (that is, had a higher share of users for whom important data were lost) shifted away from campaigns that rely more on this data. Through the end of our data window, early 2023, we see the magnitude of substitution increase. Further, we find evidence of reductions in advertising on both the extensive and intensive margins and, at the industry level, we find more affected industries have higher ad spend per advertiser.

Turning to off-platform product markets outcomes, we find that industries that were more affected by ATT saw a 1.1% decrease in the overall number of establishments versus the less affected industries. Back of the envelope calculations imply that ATT eliminated roughly 91,000 establishments in the U.S. alone. We find the industries more impacted by ATT saw a 2.9% increase in prices relative to the less impacted industries. This is a substantial increase, similar to current inflation rates (though we use the PPI instead of the CPI). Finally, we find that the adverse effects of ATT are largely borne by smaller scale advertisers, with larger ones remaining relatively unaffected. This finding is consistent with several past studies on the effects of privacy changes that find disproportionate effects on smaller firms (e.g., Wernerfelt et al. (2022), Korganbekova and Zuber (2023), Campbell et al. (2015))

Our paper contributes to the literature in two main ways. First, to our knowledge, we provide the first large-scale, empirical evidence on how digital advertising can affect market structure and product prices. Given how widespread digital advertising has become, understanding these effects is timely for both economists and policymakers. This contribution builds on two literatures: the first focused on measuring the effects of digital advertising, and the second on understanding how advertising may affect market structure and prices.

On the former, one of the upsides of digital advertising over traditional channels has been the improved ability to measure advertising effectiveness (Goldfarb and Tucker, 2011). Much of this literature to date has focused on measuring short term outcomes that advertisers directly care about, such as cost per incremental customer or cost per ad click (e.g., Blake et al. (2015), Gordon et al. (2023), and Tadelis et al. (2023); see also Lewis and Rao (2015)). Separately, there are numerous anecdotes and case studies of individual companies—particularly those who sell direct to consumers—whose creation and growth was largely enabled by digital advertising (e.g., Rangan et al. (2021), Goldstein and Malone (2022)).[2] To our knowledge, however, there has been no systematic, large-scale analysis of the broader effects of digital advertising on market structure and prices.

On the latter, there is a long literature in economics on advertising, starting as early as Marshall (1890).[3] Advertising has since been modeled through several different lenses, while several decades ago there was a flurry of research focusing on the interplay between advertising and market concentration (e.g., Dixit (1980), Fudenberg and Tirole (1984), Schmalensee (1983), and Sutton (1991)). Empirical evidence on the connection between the two has also started to emerge; for example, Ellickson (2007) and Bronnenberg et al. (2011) apply Sutton's endogenous sunk-cost model to supermarkets and CPG products, respectively.[4] A relative strength of these papers is their in depth treatment of their product markets; a relative strength of our work is that we can speak to the breadth of impact of digital advertising by comparing across millions of advertisers and all 6 digit NAICS industries in the US.

Our paper's second contribution is to provide the largest scale evidence to date on the importance of collecting and using consumer data in advertising for firms. This data is increasingly policy relevant as impending (and substantial) product and regulatory changes will directly affect advertisers' abilities to use it. By focusing on the effects of Apple's ATT policy, our findings relate to the literature that has analyzed the effects of privacy regulation on different firm outcomes. Wernerfelt et al. (2022) provide evidence that user data may substantially reduce the cost of acquiring incremental customers, but do so in a partial equilibrium, experimental setting; given the many actors and stakes involved, understanding

[2]Bronnenberg et al. (2022) analyze the growth of the craft beer industry in the U.S. and speculate a possible role for digital advertising. Given their data the authors cannot explicitly test this hypothesis, but our work was partly inspired by this paper.

[3]See Bagwell (2007) for an overview of the economics literature on advertising.

[4]There have also been several empirical papers that analyze effects of advertising on pricing. See, for example, Milyo and Waldfogel (1999), Kwoka (1984), and Benham (1972).

effects of such policies in the field is an important contribution of our work. (Alcobendas et al., 2023) measure the impact of removing tracking cookies by leveraging non-experimental data and a structural model to consider the effects of losing this data on advertiser surplus.

The closest paper to ours is Aridor et al. (2024) who use a similar identification strategy to analyze the effects of ATT on advertising effectiveness, ad budget allocation across platforms, and firm revenue. Consistent with our results, they find relatively sizable, negative effects of ATT across these outcomes. Their work focuses on e-commerce firms and does not explore the impact of ATT on market structure and product prices paid by consumers. Others have studied the app market narrowly and found that the introduction of ATT led to a decrease in iOS app downloads and new app introductions (Li and Tsai (2022)), as well as a move towards subscription based monetization models (Cheyre et al. (2023)).

The recent wave of papers analyzing the impact of ATT builds off previous studies of the impact of the GDPR. For example, Demirer et al. (2024) study how the GDPR increased cloud storage data costs to firms; Janssen et al. (2022) document how the GDPR led to substantial decreases in consumer surplus, but only in the market for apps; and Goldberg et al. (2024) and Aridor et al. (2023) document how the GDPR affected website traffic and revenue. (More broadly, see Dubé et al. (2024) for a recent overview of unintended effects of privacy regulation.) We complement this literature by analyzing a set of key business outcomes, including advertising spending and campaign objectives on the one hand, as well as the number of firms and product prices on the other. Our study also covers all U.S. industries, and by virtue of leveraging internal data at Meta, we have access to both an economically meaningful sample of advertisers and one that avoids many (but not all) selection concerns.

2 Institutional Details

Millions of businesses use digital advertising platforms to reach and engage with customers. Compared to advertising on traditional media, digital advertising platforms such as FB, Google, or LinkedIn enable businesses to target ads to relevant consumers at a substantially lower cost by incorporating information from a large set of signals and data into complex ad delivery algorithms. In this section, we provide relevant, high level institutional background on how digital advertising works today.

Firms use ads to achieve not only sales, but also awareness, sign-ups for memberships, and more. For example, an advertiser may wish to increase "likes" on their FB page to signal that they are widely liked by FB users. Given a budget and target audience, Meta could distribute the ad spend uniformly across the target audience, but in all likelihood, 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 business' FB page. Meta could continue to spend the advertiser's budget on men and women, but that would neither be good for the advertiser, nor for women as they would see ads they're not interested in. Hence, by knowing the advertiser's objective, Meta could train a computational model to predict which users would help meet that objective based on its own data by shifting the advertiser's budget to focus on men, improving outcomes for both users and advertisers. This process is commonly referred to as "delivery optimization" by ad platforms.

It is common to refer to data that is directly gathered by advertising platforms (e.g., a user's past queries on Google, a user's past "likes" on FB, or a user's resume on LinkedIn) as "first party data" or "on-platform data." We next describe how data generated "offsite" is shared with advertising platforms and used to better target objectives. For a more detailed account see Wernerfelt et al. (2022), which this section borrows heavily from.

2.1 Digital Advertising and Third Party Data

In the aforementioned example, page likes are observed directly by Meta as they occur on-platform. Many outcomes that advertisers care about–for example, purchases–happen off platform. In isolation, advertising platforms would have no way of observing these off-platform actions and thus could not optimize ad delivery for objectives such as sales or sign-up registrations on the advertiser's website. This is where "third party data" (3PD) matters for ad delivery. 3PD consists of data that were generated outside the platform; for example, a business that advertises on FB can share information with Meta about purchases and other related activity by users on the advertiser's own websites and apps (e.g., adding an item to the shopping cart, signing up for a newsletter, a purchase, etc.).

Given that the objective of many advertising campaigns are actions that occur off the advertising platform, advertisers commonly use technologies that allow user data to be tracked and sent back to the focal platform for use in ad delivery. At Meta, this is commonly done with a "Meta Pixel" or the "Meta Software Development Kit" (SDK). An advertiser who wants to drive sales on their website can install a pixel (a piece of software) that logs actions–e.g., when customers make a purchase–and sends them to Meta. Meta then acquires an off-platform signal that it can use in its prediction model. Hence, advertisers benefit from installing pixels on their websites as these improve their advertising effectiveness. (SDK's perform similar functionality mobile apps.)

Without the 3PD signals ad delivery can only be optimized for on-platform signals. To the extent that the on-platform signals (e.g., "likes") are a good proxy for the off-platform signals, there may be little loss of ad effectiveness. However, Wernerfelt et al. (2022) found that optimizing for on platform outcomes greatly diminishes ad effectiveness when the advertiser's objective is an off platform outcome. The distinction is important as many looming policy and product changes are considering limiting the 3PD ecosystem, and it is unclear what this will mean for firms who rely on digital advertising. Apple's App Tracking Transparency (ATT) policy was the largest shift in the 3PD 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. ATT is operationalized through an Apple designed pop-up prompt that is displayed when the user installs and opens an app for the first time after upgrading to iOS 14.5. As shown in Figure 1, users can either "Ask App Not to Track," which disables 3PD, or "Allow" the app to record and share 3PD. All apps using the iOS platform must offer this prompt to users; failure to comply will lead to the removal of the app from Apple's App Store.

Figure 1: Apple's App Tracking Transparency User Interface

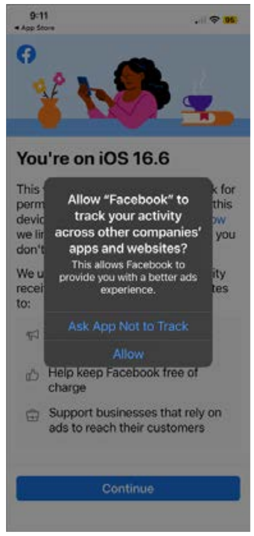

Apple's stated aim was to "improve transparency and empower users" with greater control over how their data was collected and used by third-party platforms and developers.[5] Prior to the implementation of ATT, users could choose to restrict 3PD information sharing across apps by using the "Limit Ad Tracking" setting. The introduction of ATT, however, required apps to obtain explicit permission from the user to enable 3PD information sharing.

Only if a user decides to allow the collection and use of their 3PD, can applications observe the user's activity across any third party applications and websites. The launch of ATT caused a large share of iOS users to not allow the use of their 3PD on FB and IG (and other ad platforms). We call these "ATT impacted users" because Meta was unable to leverage their off-platform data to deliver ads. Figure 2 shows that there was a gradual adoption of iOS 14.5 and subsequent releases with ATT prompts overtime, with slower adoption close to the immediate release of iOS 14.5 which accelerated in June 2021.

Among other things, ATT impacts 3PD for digital advertising through an "Identifier for Advertisers" (IDFA), a unique identifier assigned to each iOS device.[6] When users visit advertisers' websites or apps, advertisers can share information on their activity with Meta

[5] Apple 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/`.

[6] Users can reset the IDFA by going to the "Privacy & Security" settings, selecting "Tracking", and toggling "Allow Apps to Request to Track" off and on. Each off and on toggle resets the IDFA.

Figure 2: Adoption of iOS 14.5 and subsequent versions with ATT prompts

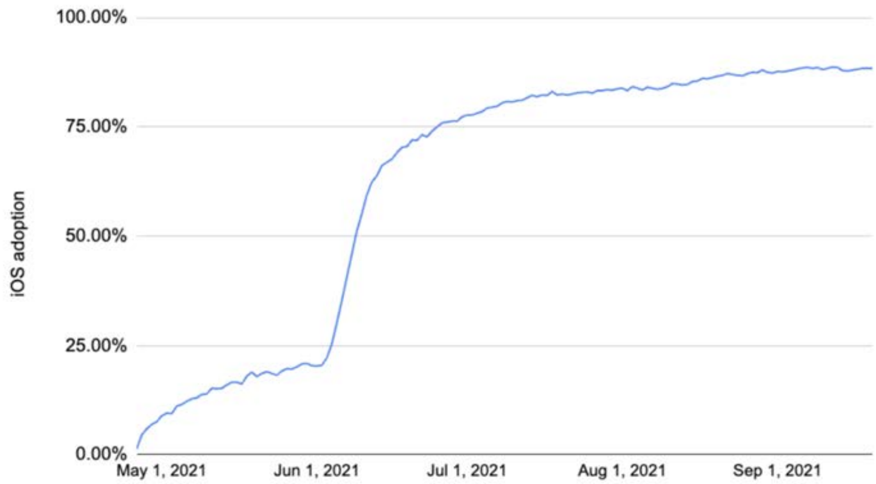

Notes: This figure shows the gradual adoption of iOS 14.5 and subsequent versions overtime, with an accelerated adoption in June 2021. Source: Gupta Media: `https://lookerstudio.google.com/u/0/reporting/3d5dda40-37ea-4b9f-bd91-bb8df8e12620/page/aDUJC?s=kTs6iab_AhQ`

through data sharing tools such as the aforementioned Pixel. Meta can then link the actions on the advertisers properties to FB and IG users through the IDFA. In practice, for ATT impacted users, advertisers are unable to observe the user's IDFA and therefore cannot share the actions of these users with Meta and other digital advertising platforms.

To structure our empirical analysis of how ATT will impact advertisers and market outcomes, we turn to our theoretical framework, which models the launch of ATT as a deterioration in the effectiveness of advertising within an industry equilibrium model.

3 Theoretical Framework

Inspired by the seminal work of Sutton (1991), we view advertising, which raise 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).

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. In essence, for any price p, a higher level of aggregate advertising increases the quantity demanded at that price. Note that we 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. For simplicity, we adopt the following functional form of demand,

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

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,

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

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:

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

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 (1) into the second stage profit function yields,

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

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:

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

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 (2) with the equilibrium solution given in (3) to obtain,

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

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

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

Proposition 3.1. *A unique free entry equilibrium exists if and only if* $e^2 R^4 \geq 32k$. *Furthermore, as advertising effectiveness* e *increases, the equilibrium number of firms* N^* *(weakly) increases, while the equilibrium price* p^* *decreases.*

Proposition 3.1 is proven in Appendix A and can in part be described graphically. Consider a numerical example in which $R = 10$ and $k = 40$, and imagine a comparative static analysis in advertising effectiveness, e. From Proposition 3.1, a necessary condition for the existence of a free entry equilibrium is $R^4 e^2 \geq 32k$, which basically means that entry costs (k) must be low relative to the strength of the potential market (R) and advertising effectiveness (e). 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.

Figure 3: Free Entry Equilibrium Number of Firms

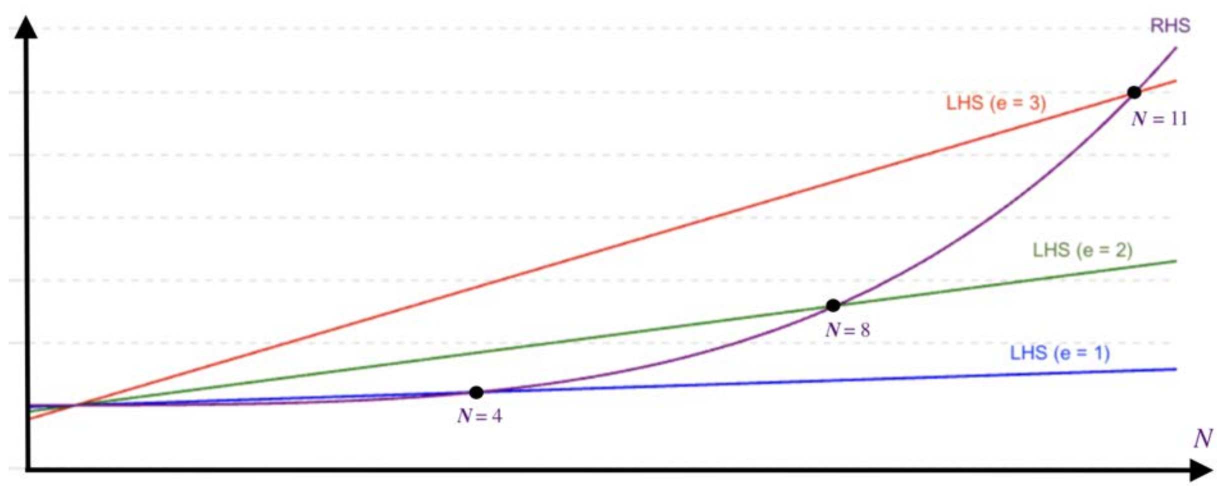

To see how changes in e impact the equilibrium number of firms, Figure 3 plots the convex RHS function $k(N+1)^4$ with $k = 40$, together with the linear LHS $e^2 R^4$ with $R = 10$ at three different values of advertising effectiveness: $e = 1$, $e = 2$, and $e = 3$. As the plot demonstrates, there 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.

Price *decreases* in advertising effectiveness because the increase in competition from entry more than compensates for the expanded market demand due to increased advertising. Indeed, calculations in Appendix A 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 in our model. Hence, the analysis above offer three testable predictions: Following ATT,

1. Firms will exit the market ($N^* \downarrow$)
2. Total advertising will decrease ($A^* \downarrow$) while surviving advertisers invest more ($a^* \uparrow$)
3. Product prices will increase ($p^* \uparrow$)

4 Data

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

4.1 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.

4.1.1 Ad-User Interactions

For our identification strategy, we use data from billions of ad-user interactions to compute the share of ads being delivered to ATT impacted users across individual advertisers and industries over a three year period from 2020 to 2023. A major strength of our data is that it includes the universe of all Meta users in the U.S.

For each ad impression on FB or IG, we observe the ATT status of every user to whom the ad was delivered on a given date, as well as the cost accrued to the advertiser for the ad impression. With these, we measure ATT impacted user rates for advertisers and the industries they belong to as the share of ad spend directed towards ATT impacted users at specific points in time in the post-ATT period.

For example, imagine there are 100 advertisers: 50 in industry A, 30 in industry B, and 20 in industry C. For each advertiser we measure the fraction of ad-spend that was directed towards ATT impacted users, and then aggregate these fractions at the industry level. If, say 62% of ad-spend in industry A was directed at ATT impacted users, while the fractions for industries B and C were 42% and 56% respectively, then this creates a ranking from most- to least-impacted industry vis-a-vis how impacted targeting is due to ATT impacted users: Industry A was most impacted (62%), next industry C (56%), and last industry B (42%). This variation is key to our identification strategy as explained in Section 5.

4.1.2 Outcomes

Industry-level. Meta internally classifies all advertisers into 25 aggregated industries, such as "Retail," "Ecommerce," and "Professional Services," which are further segmented into 200 industries. For instance, advertisers classified as "Retail" are further classified into "Apparel and Accessories," "Beauty," and others. These industries are the level at which on-platform data is matched with off-platform industry classifications, as described below. For each industry, we collect monthly data on the number of active advertisers and the total ad-spend on FB or IG ad campaigns. Advertisers naturally enter and exit Meta over time so we will refer to reductions in the count of advertisers as "net exit" from the platform.

Advertiser–level. We fix a sample of advertisers and track their on-platform behavior over time. For tractability, we focus our primary analyses on all U.S. advertisers active on FB or IG in mid 2020. Our choice of the sample period is motivated by two factors. First is to have a sufficiently long pre-ATT time period. 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 in early 2020. With this, we construct an advertiser-month level panel comprising a rich set of behavioral outcomes for approximately 1.4 million advertisers.

We observe whether or not an advertiser was active in a given month, which allows us to measure advertiser exit post-ATT. Second, conditional on being active in a given month, we collect data on advertisers' *campaign strategies*, i.e., how many campaigns the advertiser ran that were optimized to achieve "lower–funnel" outcomes like sales, as well as "higher-funnel" outcomes like brand awareness. Third, we measure advertisers' total ad spend on FB and IG in a given month. Finally, we incorporate a set of advertiser characteristics including size, industry, and location in various sections of the paper, and provide details accordingly.

4.2 U.S. Industry Data

We leverage publicly available BLS data that are reported at the industry level using the North American Industry Classification System (NAICS). U.S. industries are assigned a 2-digit NAICS code, and subsequent sub-industries have additional digits in their NAICS codes. For our analyses we use the most granular available classification for each outcome.[7]

Prices. We obtain monthly data on prices from the BLS' Producer Price Index (PPI). The PPI data is available at the 6-digit NAICS code-month level and covers all manufacturing

[7]The most granular industry-level data is at the 6-digit NAICS level. For instance, Agriculture, Forestry, Fishing and Hunting is assigned a NAICS code 11. This is further disaggregated into Crop Production (NAICS code 111), Oilseed and Grain Farming (NAICS code 1111), Soybean Farming (NAICS code 11111), Soybean Farming (NAICS code 111110), Oilseed (Except Soybean) Farming (NAICS code 111120).

industries, as well as about 70% of service industries. The PPI measures the average change over time in the selling prices for all U.S. domestic producers of goods and services.[8]

Firm Count. We leverage quarterly data on the number of existing business establishments from the BLS' Quarterly Census of Employment and Wages (QCEW). This rich dataset covers all private sector establishments and we use the 6-digit NAICS industry codes at the state-quarter level. Table 1 presents pre-ATT summary statistics for our main variables just prior to the ATT announcement. Panel A describes our on-platform industry data with 200 Meta defined industries over 9 months pre-ATT ($N = 1,800$), panel B describes our individual advertiser data, and panel C presents the summary statistics for our off-platform outcomes. For the off-platform data on prices, we observe 51 industries for 11 months in the pre-ATT period ($N = 561$), and for data on the number of establishments, we observe 5,153 state-industry units for 4 quarters ($N = 20,612$).

4.2.1 Matching Meta's Industry Verticals to BLS industries

Linking on-platform opt-out rates to off-platform data requires a systematic mapping between Meta's internal industry classification and the NAICS industries. This mapping is facilitated by the availability of detailed descriptions of the types of business establishments included within each NAICS industry[9] along with Meta's systematic internal categorizations of advertisers into industries. As a first step, we manually match Meta's internal industry classifications to the NAICS industries based on these descriptions. As a second step, we validate our matching via hundreds of random advertiser-level audits confirming individual advertisers are indeed in the assigned NAICS industry.

There are two additional nuances to our matching method. First, the process depends on the granularity of the off-platform data. For prices and firm counts that are available at the 6-digit NAICS code, we frequently match multiple NAICS industries to a single on-platform

[8]This is in contrast to other price metrics such as the Consumer Price Index (CPI) which measures the change in prices over time of a basket of goods and services consumed by a typical U.S. urban household—given our focus on outcomes at the industry level, PPI is a more appropriate outcome.

[9]For example, description of the Direct Property and Casualty Insurance Carries NAICS industry can be found at `https://www.naics.com/naics-code-description/?code=524126`

Table 1: Summary Statistics

	Mean	SD	10th Percentile	25th Percentile	Median	75th Percentile	90th Percentile	N
Panel A: On-Platform Industry Data								
Post-ATT Opt-out Shares	49.011	6.620	41.049	44.493	49.168	53.310	57.341	200
Log(No. Total Advertisers)	6.235	2.129	3.520	4.836	6.420	7.661	8.576	200
Log(Total Ad Spend)	15.176	2.074	12.350	14.057	15.729	16.422	17.307	200
Panel B: Individual Advertiser Data								
Entertainment and Media	0.161	0.368	0	0	0	0	1	1,450,817
Retail	0.127	0.333	0	0	0	0	1	1,450,817
Ecommerce	0.115	0.319	0	0	0	0	1	1,450,817
Professional Services	0.172	0.378	0	0	0	0	1	1,450,817
Publishing	0.037	0.189	0	0	0	0	0	1,450,817
Consumer Packaged Goods	0.038	0.191	0	0	0	0	0	1,450,817
Small Advertiser	0.966	0.181	1	1	1	1	1	1,450,817
Large Advertiser	0.034	0.181	0	0	0	0	0	1,450,817
No. Campaigns	10.798	77.029	1	1	3	8	18	1,450,817
Ad Spend	5420.98	103295.51	10.00	30.00	116.31	509.45	2269.07	1,450,817
Age (days)	781.608	896.479	16.000	63.000	435.000	1211.000	2106.000	1,450,817
Share Lower Funnel Campaigns	0.373	0.253	0.08	0.167	0.333	0.556	0.769	1,450,817
Share Upper Funnel Campaigns	0.627	0.253	0.231	0.444	0.667	0.833	0.92	1,450,817
Panel C: Off-Platform Industry Data								
Log(Number of Establishments)	5.460	2.050	2.770	4.130	5.490	6.790	8.070	5,153
Log(Price-PPI)	5.091	0.345	4.745	4.888	5.104	5.247	5.356	51

Notes: The table presents pre-ATT summary statistics, including the mean, standard deviation, the 10th, 25th, 50th, 75th and the 90th percentiles for our sample. Panel A presents describes our on-platform industry data, Panel B describes our on-platform advertiser data, and panel C describes our off-platform industry data

Figure 4: Example Matching of NAICS industries to Meta Sub Verticals

industry. In a few cases, we also match multiple on-platform industries to a single NAICS 6-digit industry. Figure 4 shows examples of both kinds of matches.

Second, NAICS codes are updated across North America every five years to reflect appropriate changes in the economic landscape. The most recent code updates were implemented in 2022, which falls within the timeframe of our analysis.[10] While many of the updates were simply changes in the NAICS code number itself, with no changes in the underlying composition of the industry, some reclassifications resulted in a set of industries being either merged together or split further.[11] Crucially, for industries affected by these changes, the relative contributions of previous industries in the newly re-classified industries is unavailable.[12] To maintain consistency throughout the period of our analysis, we perform similar reclassifica-

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

[11]For instance, industries coded as Anthracite Mining (NAICS: 212113), Bituminous Coal and Lignite Surface Mining (NAICS: 212111), and Bituminous Coal Underground Mining (NAICS: 212112) in the 2017 NAICS code classification were reclassified into Underground Coal Mining (NAICS: 212115) and Surface Coal Mining (NAICS: 212114)

[12]For instance in the reclassification example provided in footnote 13, while it is clear that all of Bituminous Coal and Lignite Surface Mining contributes to Surface Coal Mining and all of Bituminous Coal Underground Mining contributes to Underground Coal Mining, it unclear what share of Anthracite Mining contributes to Underground Coal Mining and Surface Coal Mining

tions in the pre-2022 data; if a set of NAICS industries were merged in the post-2022 data after reclassification, we merge the same NAICS industries in the pre-2022 data.

5 Empirical Strategy

We employ a difference-in-differences (DiD) framework to estimate the causal impact of ATT's data sharing restrictions on key business outcomes, both for advertisers on Meta's ads platforms as well as for firms across the wider U.S. economy. Given that Apple's ATT was rolled out globally at the same time, our empirical strategy exploits variation in the extent to which different advertisers and industries are more versus less impacted by the ATT shock.

5.1 ATT Impacted User Rates and Treatment Assignment

As described in Section 4.1.1, for every advertiser and every industry we calculate the share of FB and IG ad spend for ads that were delivered to ATT impacted users, henceforth referred to as "ATT impact." Naturally, advertisers and industries with a greater ATT impact are more impacted by the ATT policy change.[13] Because advertisers' strategies and target audiences may have endogenously changed post-ATT, we measure ATT impact as the proportion of ad spend for ads delivered *before* ATT to target users that did not opt in to data sharing *after* ATT. This ensures that we are able to capture the degree to which advertisers can no longer receive offsite signals from their optimal set of (pre-ATT) targeted users.

Advertiser-level treatment assignment. For each advertiser in our sample, we estimate ATT impact at two distinct points in time: August 2021 and March 2023.[14] We define advertisers in the top 10% of ATT impact as our primary treatment group, and advertisers in the bottom 10% as our control group. We find that of the advertisers in the top decile of ATT impact shares in August 2021, 85 percent are also in the top decile in March 2023. Similarly, we find that 93 percent of advertisers in the bottom decile in August

[13] Since ATT does not impact web browsers on computers, advertising and advertising platforms such as Meta may still receive offsite signals from users who engage on the web.

[14] We are able to observe these at fewer points in time compared to the industry level opt-out shares due to computational constraints (e.g.: there are about 200 industries vs 1.5 million advertisers in our sample).

Figure 5: ATT Impact Rates for More vs. Less Impacted Advertisers

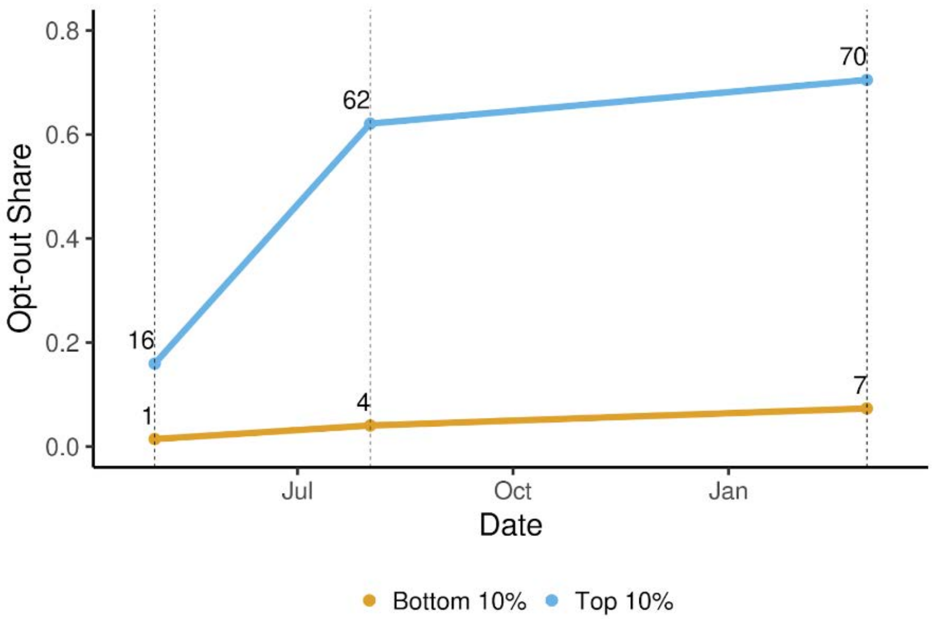

Notes: This figure plots the ATT impact rate for the 10% most vs. least impacted advertisers. For each advertiser, we calculate the share of the pre-ATT ad spend delivered to post-ATT impacted users at three distinct points in time: May 2021, August 2021, and March 2023.

2021 are also in the bottom decile in March 2023. Taken together, these imply high rates of stability in our sample and consequently, reasonably valid assignment of treatment.

Figure 5 shows that there was a differential uptake in ATT between the more and less affected advertisers with the average ATT impact for advertisers in the top 10% reaching 70 percent, while those for advertisers in the bottom 10% are only 7 percent by March 2023.

Industry-level treatment assignment. We estimate ATT impact for each industry based on user-level opt-out rates at seven distinct points in time in the post-ATT period. Similar to the individual advertiser analysis, for our on-platform analyses we define industries in the top 10% of ATT opt-out rate as treated and those in the bottom 10% as control. We find that across the different points in time, the set of most and least impacted industries remain remarkably stable.[15] Additionally, there is a sizable gap in ATT impact shares across these industries; as of April 2023—the latest date in our analysis—the average ATT impact

[15]These include the months 5/21, 8/21, 12/21, 3/22, 12/22, 3/23, and 4/23.

share for the 20 most impacted industries is about 60.3 percent, whereas for the 20 least impacted industries, it is only about 37.5 percent. For our analyses using off-platform data, which is more sparse, we define industries above the median ATT impact as our primary treatment group and those below the median as our control group.

Table 2 presents the summary statistics for our main variables split by treatment and control groups. For our on-platform industry sample in Panel A, we see slight differences between treatment and control industries where the former have slightly more advertisers and consequently more ad spend. Panel B presents our individual advertiser data for a range of variables including the share of advertisers that fall into the five largest industries by total revenue share, advertiser age in days, advertiser size[16], total advertising spend in a given month, the total number of campaigns they ran in a given month, and the proportion of those campaigns that were optimized for lower- and upper-funnel goals. For individual advertisers, the magnitudes of a range of characteristics are relatively similar across the treatment and control groups, including those pertaining to advertising strategy as well as the distribution of advertisers across a range of industries. The latter enables us to mitigate concerns that our results are driven by a small set of industries facing idiosyncratic shocks. Finally, Panel C shows that our off-platform outcomes are similar in the pre-ATT period across the treatment and control industries.

5.2 Empirical Specification: The Difference-in-Difference Model

Our main specification takes the following standard form:

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

where i represents an industry (or advertiser) and t a period. Y_{it} is the outcome of interest (e.g., ad-spend A^* or number of firms N^* for industry i in period t). $Treat_i$ is an indicator variable which equals 1 for treated industries (or advertisers) and equals 0 for control industries

[16]We use Meta's internal classification that defines an advertiser as large if they are a top 100 global business brand, or a top 200 U.S. advertising brand categorized by Ad Age, or spend at least \$5 million annually on advertising or marketing, or spend at least \$500,000 on advertising on Meta's platforms

Table 2: Summary Statistics by Treatment and Control Groups

	Mean	SD	P10	P25	Median	P75	P90	N	Mean	SD	P10	P25	Median	P75	P90	N
	Top 10%								Bottom 10%							
Panel A: On-Platform Industry Data																
Post-ATT Opt-out Shares	60.25	2.19	57.53	58.74	60.12	61.37	62.88	20	37.47	3.07	33.70	36.66	37.44	40.12	40.59	20
Log(No. Total Advertisers)	7.55	1.44	5.85	6.43	7.61	8.57	9.27	20	4.36	2.41	1.36	3.22	3.90	5.78	7.58	20
Log(Total Ad Spend)	16.82	1.25	14.91	16.41	16.97	17.53	18.09	20	13.37	2.98	10.62	12.53	13.13	15.13	17.00	20
Panel B: Individual Advertiser Data																
Entertainment and Media	0.251	0.434	0	0	0	1	1	113,311	0.158	0.365	0	0	0	0	1	126,235
Retail	0.143	0.351	0	0	0	0	1	113,311	0.186	0.389	0	0	0	0	1	126,235
Ecommerce	0.125	0.33	0	0	0	0	1	113,311	0.138	0.345	0	0	0	0	1	126,235
Professional Services	0.166	0.372	0	0	0	0	1	113,311	0.086	0.280	0	0	0	0	0	126,235
Publishing	0.035	0.184	0	0	0	0	0	113,311	0.073	0.261	0	0	0	0	0	126,235
Consumer Packaged Goods	0.052	0.222	0	0	0	0	0	113,311	0.030	0.172	0	0	0	0	0	126,235
Small Advertiser	0.966	0.18	1	1	1	1	1	113,311	0.991	0.093	1	1	1	1	1	126,235
Large Advertiser	0.034	0.18	0	0	0	0	0	113,311	0.009	0.093	0	0	0	0	0	126,235
No. Campaigns	10.401	51.551	1	1	3	7	19	113,311	14.630	98.122	1	2	4	10	24	126,235
Ad Spend	12542	137218	10	28	98	546	5524	113,311	1065	16029	5	17	70	273	913	126,235
Age (days)	729	872	16	52	393	1104	2009	113,311	492	723	12	17	135	729	1527	126,235
Share Lower Funnel Campaigns	0.466	0.284	0.111	0.200	0.429	0.720	0.879	113,311	0.338	0.249	0.056	0.125	0.286	0.500	0.721	126,235
Share Upper Funnel Campaigns	0.534	0.284	0.121	0.280	0.571	0.800	0.889	113,311	0.662	0.249	0.279	0.500	0.714	0.875	0.944	126,235
	Above Median								Below Median							
Panel C: Off-Platform Industry Data																
Log(Number of Establishments)	5.480	2.070	2.830	4.110	5.480	6.810	8.130	2,577	5.44	2.03	2.71	4.13	5.51	6.78	8.01	2,576
Log(Price-PPI)	5.185	0.377	4.866	5.020	5.189	5.300	5.378	25	5.012	0.315	4.644	4.760	5.035	5.244	5.356	25

Notes: The table presents pre-ATT summary statistics, including the mean, standard deviation, the 10th, 25th, 50th, 75th and the 9th percentiles for our sample split by treatment and comparison groups. Panel A presents describes our on-platform industry data, Panel B describes our on-platform advertiser data, and panel C describes our off-platform industry data (or advertisers). Similarly, $Post_t$ equals 1 for post-ATT time periods and 0 otherwise. Finally, θ_t and λ_i are time and unit fixed effects, respectively, and ϵ_{it} is the error term. The coefficient of interest β_1 measures the impact of ATT on industries (or advertisers).

In January 2021, Meta publicly communicated that it will be forced to comply with Apple's ATT requirements, including having to display the ATT opt-out prompt that would facilitate users' opting out of third party data sharing with Meta's apps. The communication–also delivered via email to all advertisers–explicitly warned that the "App Tracking Transparency framework will have hard-hitting implications across targeting, optimization, and measuring campaign effectiveness." As advertisers on Meta can access information to gauge the share of their ad impressions originating from iOS devices, and therefore the extent to which their advertising may be impacted, an announcement of such significance could affect advertiser behavior in anticipation of ATT. This is in spite of the fact that the actual implementation of ATT, beyond a subset of beta test users, was not until the iOS 14.5 update in April 2021.

To adequately account for these potential effects on advertisers, for our on-platform analyses we construct $Post_t$ to equal 1 in the period after December, 2020 (or, after quarter-4 of 2020), as an indication of a period where the effects of ATT have started to take place, at least in the world of digital advertising. For all analyses involving off-platform outcomes, we continue to define $Post_t$ as the period starting in May 2021 (or, quarter 2 of 2021), corresponding to Apple's formal implementation of ATT, and incorporating the fact that we can expect a delay in the spillover of on-platform effects to off-platform outcomes.

Finally, we note that our estimates compare the effect of ATT on the most versus least affected industries, where the least impacted industries were still impacted at meaningful levels. Hence, our estimates are likely to be lower bounds for the counterfactual of comparing industries fully affected by ATT to a world in which ATT was not implemented.

Parallel trends. For the difference-in-differences setup to be valid, the parallel trends assumption must be satisfied. In our setting, this assumption implies that in the absence of ATT, outcomes for the treatment and comparison groups would have evolved along similar trajectories over time. To provide evidence on parallel trends, we employ a fully flexible difference-in-differences event study model with time varying treatment effects.

We also conduct several robustness tests. For our industry level analysis, we vary the sample by comparing different ranges of most affected industries, as well as dropping the Consumer Packaged Goods industry, which was most affected by COVID-19, from our analysis. For our individual advertiser results, we test the robustness of our estimates to variations in the sample as well as to a separate sample of advertisers active in early 2020; we also test an alternative empirical strategy, the Generalized Synthetic Control method (Xu, 2017).

6 Results

6.1 Preliminary Results

We start by analyzing whether ATT led to any detectable shifts in advertiser behavior. Intuitively, if the ATT shock was not substantial enough to yield detectable changes in advertiser behavior, we may worry about any effects we find (or not) in our primary analyses. In particular, ATT does not affect digital advertising effectiveness uniformly: some types of campaigns are more reliant on tracking data than others. In particular, the campaigns whose objectives relate to "lower funnel" (off-platform) actions, such as sales or app downloads, are likely more affected than the campaigns whose objectives relate to "upper funnel" (on-platform) actions, like increasing the number of "likes" or brand-awareness.

We test this hypothesis using advertiser level data on the shares of lower and upper funnel ad campaigns, conditional on the advertiser being active. Our difference-in-differences regression estimates imply a 1.4 percentage point reduction in ad campaigns optimized for lower funnel optimization goals among the most affected advertisers, relative to the least affected advertisers. Figure 6 plots the event study estimates and shows a clear shift away from lower funnel ad campaigns toward upper funnel ad campaigns. Furthermore, the divergence between the share of lower and upper funnel optimized ad campaigns increases in

Figure 6: The Impact of ATT on the Share of Lower Funnel Campaigns

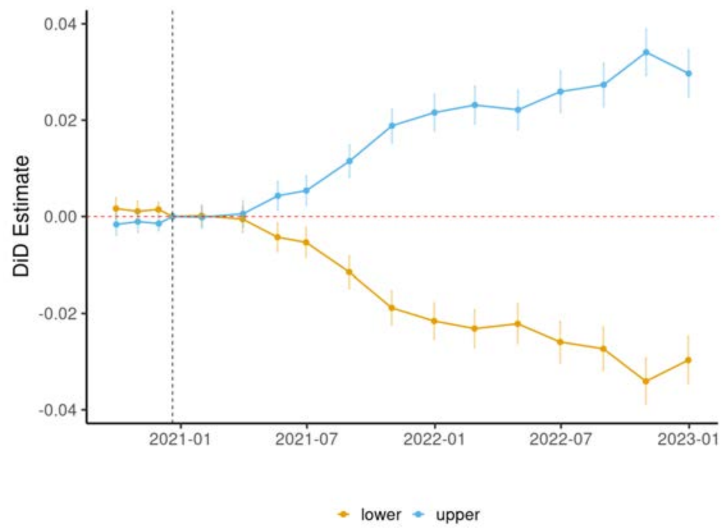

Notes:The figure presents the event study estimates of the impact of ATT on the share of ad campaigns optimized for lower and upper funnel goals along with the 95% confidence interval.

the medium-to-long term, consistent with the gradual increase in the difference in opt-out shares between the most and the least affected advertisers presented above in Figure 5.

We carry out two main robustness checks. First, we consider various samples by expanding the thresholds for treatment and comparison groups ranging from the top/bottom 10% to 50%. Appendix figure C.1 presents these coefficients, which are consistent with our main estimates. Second, we take a distinct sample of advertisers that were active in early 2020 and consider an alternative empirical strategy, the Generalized Synthetic Control (Xu, 2017) and show in Appendix Figure D.2 that our estimates are robust to this alternative empirical strategy and sample. We note again that this analysis conditions on being an active advertiser.

6.2 Effects of ATT on Advertising: A^* and a^*

Our model predicts that deteriorating advertising effectiveness should lead to less advertising in the market, i.e., a reduction in A^*, due to the reduction in the number of advertisers. At the same time, those remaining advertisers will advertise more, so a^* increases. To test these hypotheses we start by evaluating the number of active advertisers across industries: do the industries that experienced a greater reduction in advertising effectiveness see a greater reduction in the number of advertisers?

Figure 7: The Impact of ATT on the Number of Advertisers

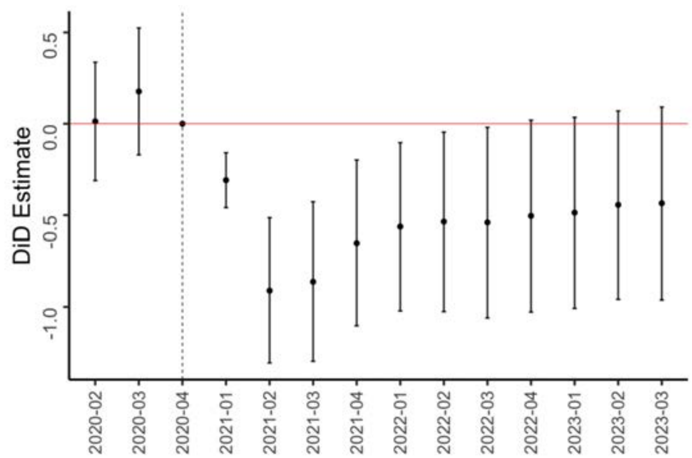

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertiser net exit at the industry sub-vertical level. The analysis uses user-ad data, aggregated at the industry level from Q2 of 2020 till Q2 of 2023.

Figure 7 displays our difference-in-differences estimates comparing industries in the top versus bottom 10% of ATT impact, which are consistent with our model's predictions. Namely, immediately after the ATT announcement, there is a 66 percent increase in net exit on FB and IG among industries more affected by ATT relative to those less affected with an overall average effect of 49.2%. Table 2 presents the difference-in-differences coefficients for net exit as well as the proportion of lower funnel ad campaigns. Since the number of distinct industries in our sample is relatively small, we test the significance of the estimates using permutation methods similar to those in Abadie et al. (2010) where we plot the distribution of estimated effect sizes and compare our estimated effect to that distribution. This means that for every 1% reduction in tracking data, 1.6% fewer firms advertise.

We note that these orders of magnitude appear broadly in line with other studies of ATT and advertising effectiveness (e.g., Wernerfelt et al. (2022); Aridor et al. (2024); Cecere and Lemaire (2023)). Each of these look at slightly different outcomes, samples, and settings, and their results are consistent with the order of magnitude and direction of effects that we observe: ATT had a substantial effect on advertising effectiveness, and with that came a substantial reduction in the number of firms who advertise.

We also conduct the analysis at the individual advertiser level, consisting of advertisers from all industries. Comparing advertisers above and below median ATT impact, we find that ATT led to a 3.4 percentage point increase in the probability of exit from advertising on Meta platforms among the most affected advertisers (Appendix E.1).

Table 3: The Impact of ATT on Advertiser Behavior

	Individual Advertiser Analysis	Industry Analysis
	(1)	(2)
	Share Lower Funnel Campaigns	Net Exit
DiD Estimate	-0.014***	-0.677**
	(0.0004)	(0.287)
Fixed Effects	Advertiser + Date	Industry + Date
R2	0.784	0.942
N	1,021,752	1600

Notes: The table presents the regression coefficients of the impact of ATT on the share of lower funnel ad campaigns (column 1) and industry net exit (column 2). Standard errors for column 1 are clustered at the advertiser level and at the industry level for column (2). ***p < 0.01, **p < 0.05

As robustness checks, we first consider variations in the analysis sample. Specifically in Appendix Figures C.2 and C.4, we vary the most and least affected industries from 10% to 50% (median) and show that regardless of the sample chosen, the coefficients remain large and negative, while it is significant in samples with a larger difference in average opt-out rates (for example comparing top and bottom 10% of the most and least affected industries). Second, we consider alternative levels of industry aggregation on Meta's platforms and find similar results (Appendix E.2). Last, we drop industries most acutely affected by Covid-19, those in the Consumer Packaged Goods industry, and re-run the analysis to find that our results remain robust (Appendix D.1)

Next, we turn to the effect of ATT on total industry ad spend (A*). To this, we turn to our individual advertiser sample. Figure 8 shows the event study plots and, while there are small pre-trends in this analysis, there's a clear drop in ad spend post ATT of roughly 28%

Figure 8: The Impact of ATT on Advertiser Ad Spend

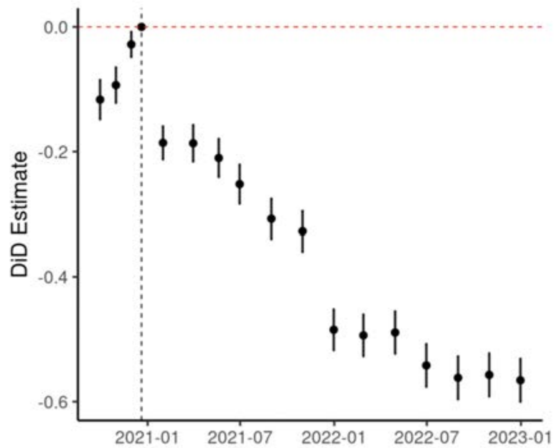

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertising spend at the advertiser level. The analysis uses user-ad data, aggregated at the advertiser level.

which is significant, despite the pre-trends being in the opposite direction of the post-period trends. As we detail in Appendix B, small advertisers were significantly more acutely affected compared to large advertisers and saw statistically distinguishable reductions in ad spend.[17]

Last, we turn to the effect of ATT on ad spend for the advertisers who remain active (a*). To study this, we consider advertising spend per active advertiser at the industry level. Figure 9 presents the event study estimates and the regression results show a 47.1 percent increase in ad spend per advertiser, consistent with our model.

6.3 Effects of ATT on the Number of Firms: N*

We now turn to our off platform results. For these, our analysis occurs at the industry-quarter level. To provide greater precision for our estimates, we present our off-platform results comparing industries above and below the median ATT opt-out rates.

[17]Note that there is a second downward jump in aggregate advertising spend in January 2022. We conjecture that this is due to adjustments of advertising budgets for the year after ATT was launched. Many companies set annual budgets for year t in Q4 of year $t-1$, and as companies learned throughout 2021 that ATT reduced ad effectiveness, they would naturally reduce 2022 budgets.

Figure 9: Ad Spending per Advertiser using Industry Level Data

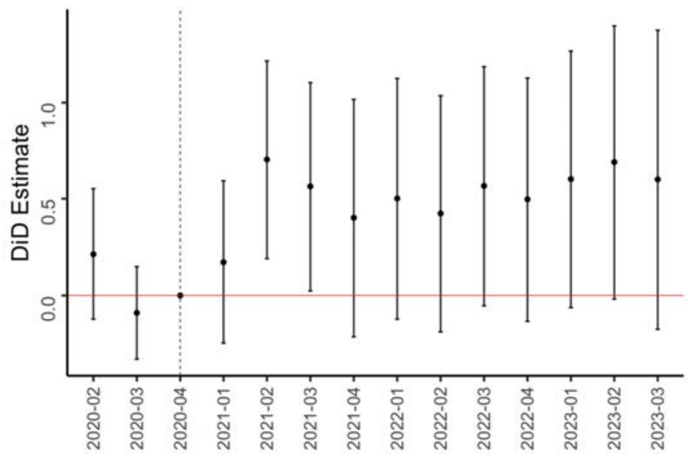

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertiser ad spend concentration using data at the industry level. Ad spend concentration is defined as ad spend per advertiser. The analysis uses user-ad data, aggregated at the industry level from Q2 of 2020 till Q2 of 2023.

The model predicts that ATT should result in a reduction in the equilibrium number of firms in a given product market. Unlike the data on business registrations and prices, the number of business establishments is available at the industry-state level. To compute opt-out shares at the state-industry level, we leverage the ATT impact as well as the state-level locations of all advertisers from our advertiser level sample. For each advertiser, we obtain the state in which the registered admin of the advertiser's ad account on Meta's platforms is located. We then aggregate the ATT impact across all advertisers in a given industry located within that state, resulting in a measure of ATT impact for approximately 5,200 state–industry pairs and run the following regression:

$$Y_{sit} = \beta_0 + \beta_1(Treat_{si}xPost_t) + \theta_t + \lambda_i + \gamma_s + \epsilon_{it}.$$

We include all state–industry pairs in the data to compare differences across pairs with above vs. below median ATT impact. Figure 10 present the event study estimates for this analysis. Relative to state–industry pairs below median of ATT impact, those above the median experienced a 1% decline in the total number of active business establishments,

Figure 10: The Impact of ATT on the Count of Business Establishments

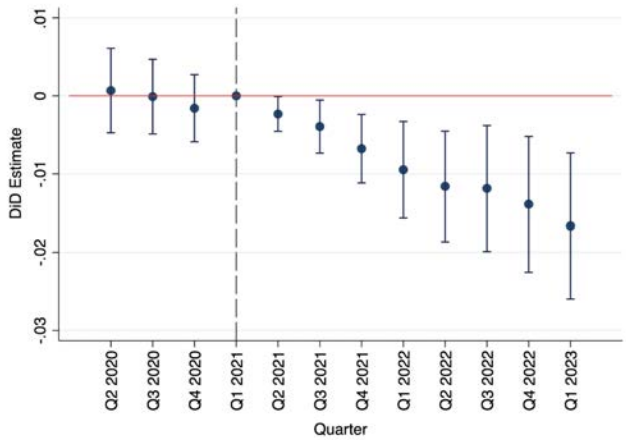

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on the total number of establishments using BLS data.

equating to roughly 91,000 establishments (see column 2 of table 4). Consistent with the preceding analysis we check the robustness of our estimates to variations in the analytical sample and show that our estimates remain robust (Appendix Figure C.6).

Table 4: The Impact of ATT on Number of Establishments and Prices

	(1) Number of Establishments	(2) Prices
DiD Estimate	-0.011***	0.029*
	(0.005)	(0.016)
Fixed Effects	State + Industry + Date	Industry + Date
R2	0.997	0.986
N	61792	700

Notes: The table presents the DiD coefficients on the impact of ATT on the off-platform number of establishments (column 1) and off-platform prices (PPI) (column 2). Data on the number of establishments is at the state-industry-quarter level and we include state, industry, and time fixed effects. Data on prices is at the industry-time level and we include industry and time fixed effects. ***$p < 0.01$, *$p < 0.1$.

Figure 11: The Impact of ATT on Prices

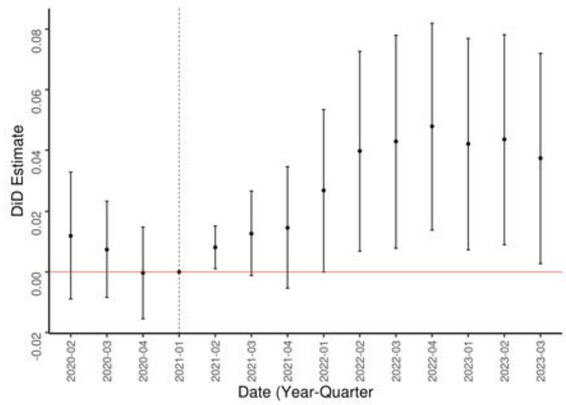

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on prices using PPI as a measure for prices.

6.4 Effects of ATT on Prices: p*

Figure 11 presents the event study estimates and shows an upward trend in prices post-ATT for industries more impacted by ATT relative to those less impacted. On average, ATT led to a 2.9% increase in prices in industries more- compared to less-impacted industries. We nevertheless check the robustness of our results to varying samples in Appendix Figure C.7. We note the event study coefficients starting mid-2022 show statistically significant effects of 4-5% higher prices, suggesting that we are capturing a delayed but stark causal effect.

7 Concluding Remarks

As the digital advertising ecosystem evolves, what data can or cannot be used by advertisers remains a central policy question for both platforms and regulators alike. Despite the push by regulators to further curb the ability of advertisers to collect and use data to target consumers, there is minimal evidence on how such policies may affect online and offline outcomes for advertisers and consumers. To our knowledge, we provide the first, large-scale evidence of how restricting a major category of online data for advertisers (third party tracking data) affected not only advertising behavior, but also firm entry, exit, and product market prices.

Our results suggest harms not only to firms, but also to consumers who face more concentrated markets with prices that are at least 3.4% higher than they would be in lieu of the ATT feature. We add the necessary caveat that a complete welfare analysis of such data restrictions is beyond the scope of our data and analyses. Nonetheless, our results raise a red flag for further restricting data use for digital advertising. Interestingly, the GDPR has also been documented to have adverse impacts on industry outcomes (Johnson, 2024).

As mentioned in the introduction, 82% of surveyed small and medium sized businesses in the U.S. report that digital advertising is 'crucial' for their success. Previous literature has found that privacy laws disproportionately have hurt smaller scale businesses (Wernerfelt et al., 2022; Korganbekova and Zuber, 2023). Solving for a model with asymmetric firms is beyond the scope of our paper, yet one expects that larger firms are more robust to shocks than smaller firms. Hence, the negative advertising shock created by ATT should impact smaller firms more than larger firms. Indeed, Aridor et al. (2024) find that small businesses were most adversely impacted by ATT, a fact that we too confirm in our data. Appendix B shows a detailed breakdown of our results above into large versus small advertisers and shows that small advertisers were indeed significantly more adversely impacted by ATT.

Related to the fact that small businesses are the most adversely impacted by data sharing restrictions is that niche businesses that rely on specific targeting may suffer the most. As a result, publishers of niche online content will have fewer advertisers and may not be able to sustain their own businesses (Breanna, 2023). Adding these observations to the documented privacy paradox mentioned in the introduction suggests that the pros of data sharing restrictions touted by many are not only hard to empirically measure and document, but the cons to businesses and consumers are meaningful. There is a need for future research to pick apart these pros and cons in meaningful and measurable ways to further inform the debate over privacy regulation.

References

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505.

Alcobendas, M., Kobayashi, S., Shi, K., and Shum, M. (2023). The impact of privacy protection on online advertising markets. *Available at SSRN 3782889*.

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*.

Aridor, G., Che, Y.-K., and Salz, T. (2023). The effect of privacy regulation on the data industry: Empirical evidence from GDPR. *RAND Journal of Economics*, 54(4):695–730.

Athey, S., Catalini, C., and Tucker, C. (2017). The digital privacy paradox: Small money, small costs, small talk. *NBER Working Paper No. 23488*, No. w23488.

Bagwell, K. (2007). The economic analysis of advertising. *Handbook of industrial organization*, 3:1701–1844.

Benham, L. (1972). The effect of advertising on the price of eyeglasses. *Journal of Law and Economics*, 15(2):337–352.

Bian, B., Pagel, M., and Tang, H. (2023). Consumer surveillance and financial fraud.

Blake, T., Nosko, C., and Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1):155–174.

Breanna, A. (February 28, 2023). What the well-meaning critics of online advertising are missing—and how they could hurt the communities they're trying to protect. *Fortune*.

Bronnenberg, B., Dubé, J.-P., and Joo, J. (2022). Millennials and the takeoff of craft brands: Preference formation in the us beer industry. *Marketing Science*, 41(4):710–732.

Bronnenberg, B. J., Dhar, S. K., and Dubé, J.-P. H. (2011). Endogenous sunk costs and the geographic differences in the market structures of cpg categories. *Quantitative Marketing and Economics*, 9:1–23.

Campbell, X., Goldfarb, A., and Tucker, C. (2015).

Cecere, G. and Lemaire, S. (2023). Have i seen you before? measuring the value of tracking for digital advertising. *http://dx.doi.org/10.2139/ssrn.4659963*.

Cheyre, C., Leyden, B. T., Baviskar, S., and Acquisti, A. (2023). The impact of apple's app tracking transparency framework on the app ecosystem. *Available at SSRN 4453463*.

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. (1980). The role of investment in entry-deterrence. *The economic journal*, 90(357):95–106.

Dubé, J.-P., Bergemann, D., Demirer, M., Goldfarb, A., Johnson, G., Lambrecht, A., Lin, T., Tuchman, A., Tucker, C. E., and Lynch, J. G. (2024). The intended and unintended consequences of privacy regulation for consumer marketing: A marketing science institute report. *Available at SSRN 4847653*.

Ellickson, P. B. (2007). Does sutton apply to supermarkets? *The RAND Journal of Economics*, 38(1):43–59.

Fudenberg, D. and Tirole, J. (1984). The fat-cat effect, the puppy-dog ploy, and the lean and hungry look. *The American Economic Review*, 74(2):361–366.

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):325358.

Goldfarb, A. and Tucker, C. E. (2011). Online advertising. *Advances in Computers, Volume 81*.

Goldstein, J. and Malone, K. (2022). Tech giants and tiny dogs.

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.

Hackett, R. and Harty, D. (October 22, 2021). Apple's ad changes wiped $142 billion off snap, facebook, and other online ad giants. *Fortune*, https://fortune.com/2021/10/22/apple-snap-facebook-earnings-google-twitter-pinterest-ad-tracking.

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. (2024). Economic research on privacy regulation: Lessons from the gdpr and beyond. In Goldfarb, A. and Tucker, C., editors, *The Economics of Privacy*. University of Chicago Press.

Korganbekova, M. and Zuber, C. (2023). Balancing user privacy and personalization. *Working paper*.

Kwoka, J. E. (1984). Advertising and the price and quality of optometric services. *American Economic Review*, 74(1):211–216.

Lewis, R. A. and Rao, J. M. (2015). The unfavorable economics of measuring the returns to advertising. *The Quarterly Journal of Economics*, 130(4):1941–1973.

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

Marshall, A. (1890). *Principles of economics: unabridged eighth edition*. MacMillan and Co., London.

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.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies*, 90(5):2555–2591.

Rangan, V. K., Corsten, D., Higgins, M., and Schlesinger, L. A. (2021). How direct-to-consumer brands can continue to grow. *Harvard Business Review*, 99(6):100–109.

Schmalensee, R. (1983). Advertising and entry deterrence: an exploratory model. *Journal of political Economy*, 91(4):636–653.

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. *NBER Working Paper No. 31201*.

Wernerfelt, N., Tuchman, A., Shapiro, B., and Moakler, R. (2022). Estimating the value of offsite data to advertisers on meta. *SSRN Working Paper*.

Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1):57–76.

Zuboff, S. (2023). The age of surveillance capitalism. pages 203–213.

Appendix

This Appendix includes six parts. In Section A we prove our main proposition and derive detailed comparative statics on advertising; in Section B we analyze effects for small versus large scale advertisers; in Section C we analyze robustness of our results to different samples; in Section D we explore robustness to other specifications; in Section E we conduct two ad hoc robustness exercises; finally, in Section F we provide ATT impact statistics by industry.

A Equilibrium and Comparative Statics

Proof of Proposition 3.1. To prove existence and uniqueness, evaluate equation (4) at $N = 0$: the LHS is negative while the RHS is zero. Now evaluate equation (4) at $N = 1$: the two sides would equal each other if and only if $e^2R^4\frac{1}{2} = 16k$. Furthermore, the slope of the LHS is an increasing linear function of N, e^2R^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 $e^2R^4 = 32k$, and multiplying each side by $\frac{1}{2}$ yields $e^2R^4\frac{1}{2} = 16k$. This implies that when $e^2R^4 = 32k$, the LHS and RHS meet only at $N = 1$, which is the unique free-entry equilibrium. If $e^2R^4 < 32k$, then the LHS is less than the RHS for any value of N and hence the two sides of 4 never cross, implying no free entry equilibrium. It follows immediately that if $e^2R^4 > 32k$, 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 $N > 1$. That N^* increases in e 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. Last, 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 (1) above. Because N^* is (weakly) increasing in e, it follows immediately that p^* (weakly) decreases in e because it decreases in N^*. □

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 3.1 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, $e_H(N) = e_L(N+1)$.

The solution to the advertising stage in equation (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 e increases *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 increases 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.

Figure A.1 plots the comparative statics described above. with a simulated example that sets $R = 10$ and $k = 40$, letting e vary. Using equation 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 (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.1: Comparative Statics for a^* and A^*

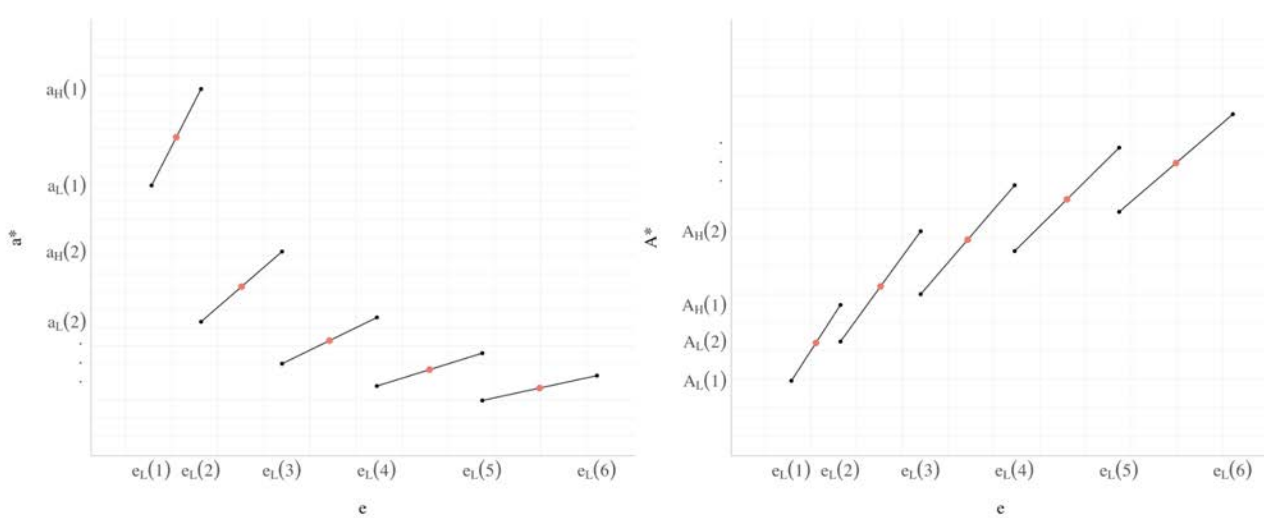

Notes: The respective midpoints are shown in red.

As the Figure A.1 shows, 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 pattens: as advertising becomes less effective so as to reduce the number if equilibrium forms N^* (a nontrivial decrease in e), then total advertising A^* decreases while the level of advertising of remaining firms a^* increases. We take the view that ATT was indeed a nontrivial decrease in e.

B Differential Effects for Small vs Large Advertisers

Our theoretical model assumes symmetric firms, and solving for a model with asymmetric firms is beyond the scope of this paper. However, if one sensibly assumes that larger firms are more robust to shocks than smaller firms, we would expect the negative advertising shock created by ATT to impact smaller firms more than larger firms. Hence, we posit that ATT may especially reduce ad effectiveness for small advertisers, who have fewer alternative advertising outlets compared to larger businesses. Existing literature suggests that small businesses are more heavily reliant on digital and social media advertising—and therefore are likely to be more impacted by ATT (Moorman (2022); Peck (2022)). Furthermore, previous literature has found that privacy laws disproportionately have hurt small businesses (Wernerfelt (2022), Korganbekova and Zuber (2023)). With over 99% of businesses operating in the U.S. classified as small and medium businesses (U.S. Small Business Administration, 2023), the subsequent analysis in this section may be of particular importance from a policy perspective.

We adopt Meta's internal classification of advertiser by size. An advertiser is classified as large if they are a top 100 global advertising brand, or a top 200 U.S. advertising brand categorized by Ad Age, or spend at least $5 million annually on advertising, or spend at least $500,000 annually on advertising on Meta's platforms. All other advertisers are classified as small advertisers. Consistent with the broader US economy, approximately 97% of the advertisers in our advertiser-level dataset are classified as small advertisers.

We are unable to systematically obtain information on the number of employees and annual revenues—variables traditionally used to determine business size—for the millions of small businesses using Meta's advertising platforms. However, we expect a high correlation between large and small advertisers vs large and small businesses because Meta's internal classification is based to some extent on ad spend.

Advertising Strategy. Figure B.1 presents the event study estimates of the impact of ATT on the share of lower funnel ad campaigns by advertiser size. We observe a consistent shift away from lower to upper funnel ad campaigns across both small and large advertisers. We note that the shift is larger for large advertisers; in practice, very few small advertisers engage in any upper funnel advertising at baseline, meaning it may be a less appealing option to substitute to than it is for large advertisers.

Figure B.1: The Impact of ATT on Advertising Campaign Strategy by Advertiser Size

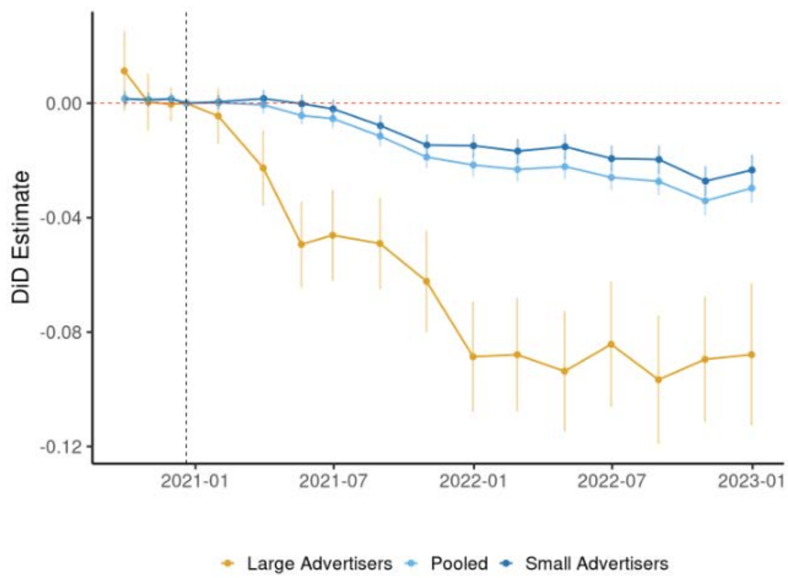

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertiser advertising strategy using individual advertiser data split by advertiser size. The analysis uses user-ad data, aggregated at the advertiser level.

Effects on Extensive Margin of Advertising, Firm exit. Figure B.2 presents the event study estimates of the impact of ATT on industry level net exit and the advertiser level probability of being active on Meta's platform split by small and large advertisers. A clear pattern emerges from these figures. The adverse impact of ATT on advertiser net exit is driven almost entirely by increased net exit among small advertisers. Our industry level difference-in-difference estimates suggest a 51% higher net exit among small advertisers in the most affected industries relative to small advertisers in the least affected industries. In contrast, the DiD estimate for large advertisers indicates a decrease in net exit in the most affected industries, compared to the least affected industries, although this effect is small and not statistically significant. Nevertheless, we can rule out equality across the estimates for small vs large advertisers (p-value = 0.03).

Figure B.2: The Impact of ATT on Net Exit and the Probability of Being Active by Advertiser Size

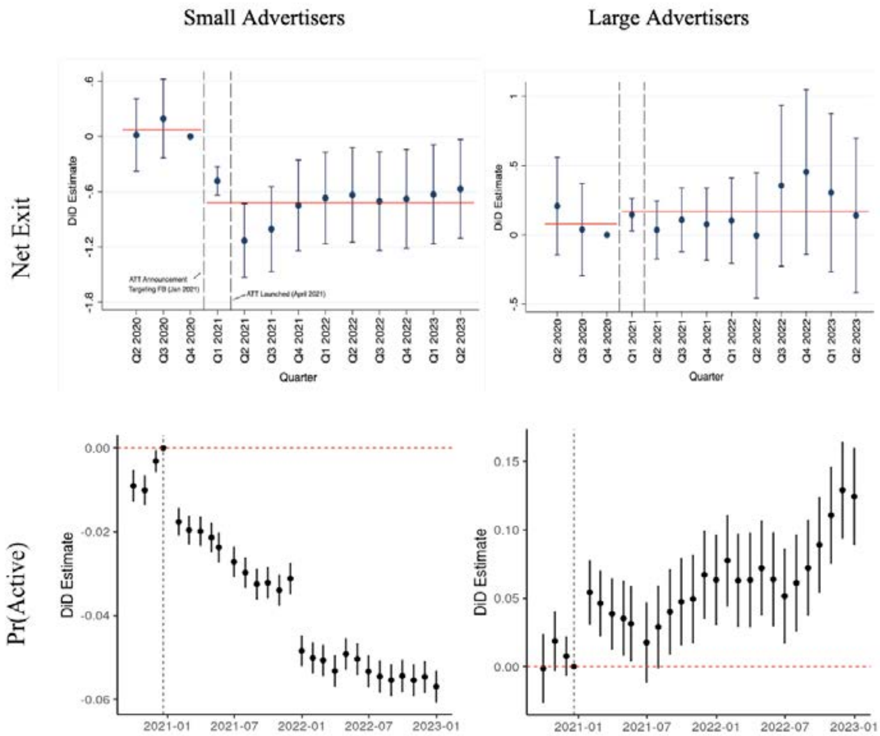

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertiser exit and net exit split by advertiser size. The first row of figures uses industry level data while the second row uses data at the advertiser level.

The individual advertiser level analyses corroborate these patterns. The impact of ATT on advertiser exit is driven entirely by small advertisers, who are more reliant on offsite signals and thereby more impaired by ATT's data sharing restrictions. ATT leads to a 3.5 percentage point increase in the probability of exit for small advertisers. In contrast, large advertisers experience a 5.6 percentage point *decrease* in the probability of exit. To provide further credence to our results, we present corresponding analyses where we categorize the size of advertisers in our sample based only on quartiles of lifetime advertising spend on FB and IG. These results are presented in Figure B.3 below. Both figure illustrate that the adverse effects of ATT were concentrated among the smaller advertisers.

Figure B.3: The Impact of ATT on the Probability of Exit by Historical Advertiser Spend

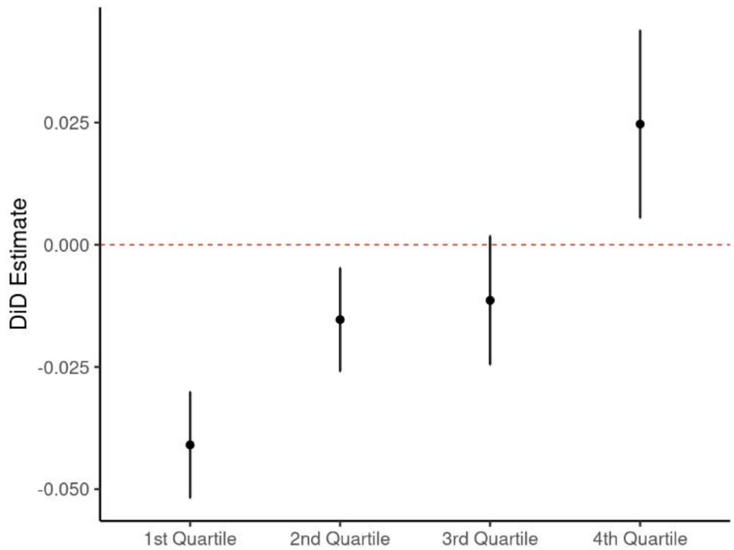

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertiser exit split by advertiser size quartiles of lifetime advertising spend on FB and IG.

Effects on Intensive Margin of Advertising, Firm ad spend. Additionally, in Appendix figure B.4, we show consistent patterns in ad spend whereby ad spend for most affected small advertisers decreases by 29%. Although the positive average effect size for large advertisers is less reliable given that the pre-ATT trends trend in the same direction as the post-ATT trends, we do not find any evidence of decreasing spend for these advertisers. Table B.1 presents the regression coefficients and summarizes our results on the heterogeneous effect of ATT by advertiser size. The difference in effects across small vs large advertisers is intuitive. As ad effectiveness declines, particularly for lower funnel ads that rely more heavily on offsite signals, small advertisers exit the market. Large advertisers, on the contrary, are better positioned to shift towards upper funnel ad campaigns, where they also face lower prices per ad impression that remain relatively stable over time. As a result, they are less likely to exit the market.

Figure B.4: The Impact of ATT on Ad Spend by Advertiser Size

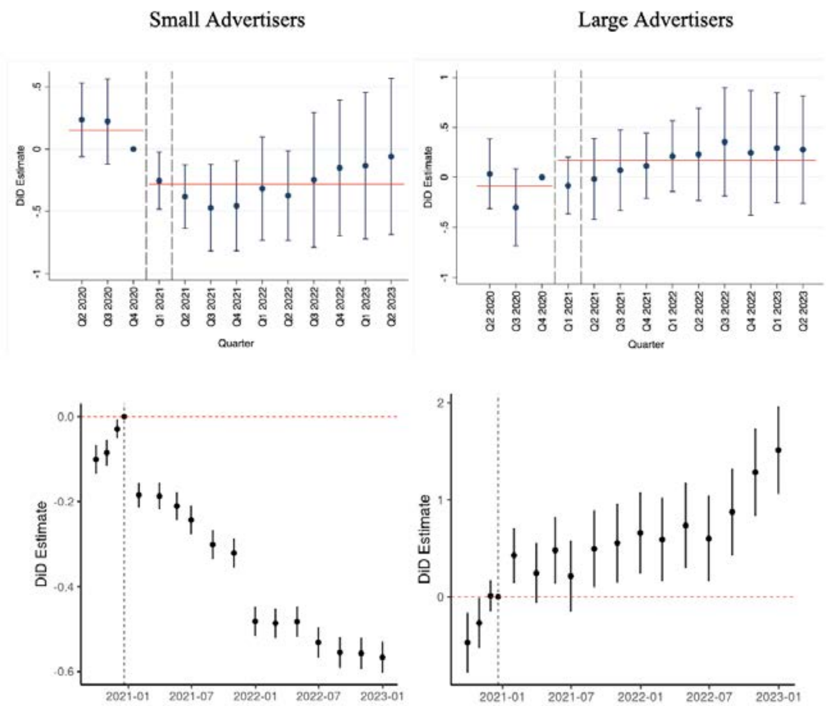

Notes: The figure presents the event study estimates along with the 95% confidence intervals of the impact of ATT on advertiser ad spend split by advertiser size. The first row of figures uses industry level data while the second row uses data at the advertiser level.

Table B.1: Difference in Differences Estimates of ATT by Advertiser Size

	(1) Pooled Sample	(2) Small Advertisers	(3) Large Advertisers
Panel A - Industry Level Analysis			
Net Exit	-0.677**	-0.791**	0.090
	(0.287)	(0.318)	(0.275)
Ad Spend	-0.211	-0.429	0.252
	(0.277)	(0.273)	(0.292)
Panel B- Advertiser Level Analysis			
Share Lower Funnel Ad Campaigns	-0.014***	-0.010***	-0.061***
	(0.001)	(0.001)	(0.007)
Pr(Active)	-0.034***	-0.035***	0.056***
	(0.003)	(0.004)	(0.007)
Ad Spend	-0.337***	-0.339***	0.850***
	(0.050)	(0.048)	(0.150)

C Robustness to Variations in the Analysis Sample

While our main analysis considers top and bottom 10% most and least affected advertisers and industries for on-platform analysis, and above below median for the off-platform analysis, we test the robustness of our estimates to variations in the analysis sample. Specifically, we vary the comparison sample from top and bottom 10% to 50% for each of our main results. Table C.1 below summarizes these results

Table C.1: Impact of ATT by Variations in the Analytical Sample

Sample (Top/bottom)	**10%**	**20%**	**30%**	**40%**	**50%**
Share of Lower Funnel Campaings	-0.0136***	-0.0157***	-0.0152***	-0.0132***	-0.0112***
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
Pr(Exit)	-0.0339***	-0.0208***	-0.0161***	-0.0143***	-0.0122***
	(0.0036)	(0.004)	(0.003)	(0.003)	(0.003)
Advertiser Ad Spend	-0.337***	-0.206***	-0.158***	-0.134***	-0.1107***
	(0.011)	(0.008)	(0.007)	(0.006)	(0.005)
Net Exit	-0.677**	-0.693**	-0.297	-0.310	-0.241
	(0.287)	(0.234)	(0.202)	(0.180)	(0.163)
Industry Ad Spend	0.466**	0.668***	0.483***	0.356**	0.306**
	(0.245)	(0.219)	(0.180)	(0.162)	(0.141)
Number of Establishments	-0.018	-0.031***	-0.021***	-0.150***	-0.011*
	(0.011)	(0.007)	(0.006)	(0.005)	(0.005)
Prices	0.0385	0.0343	0.0312	0.0275	0.0292*
	(0.031)	(0.023)	(0.019)	(0.018)	(0.016)

Figure C.1: The Impact of ATT on the Share of Lower Funnel Optimized Ad Campaigns by Variations of the Analytical Sample

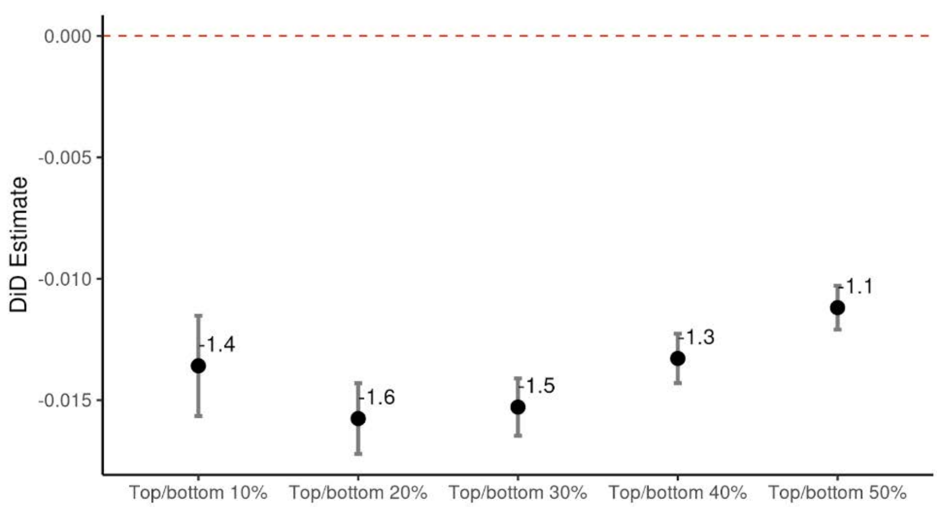

Notes: The figure presents the DiD coefficients of the impact of ATT on the share of ad campaigns optimized for lower funnel goals. Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

Figure C.2: The Impact of ATT on the Advertiser Probability of Exit by Variations of the Analytical Sample

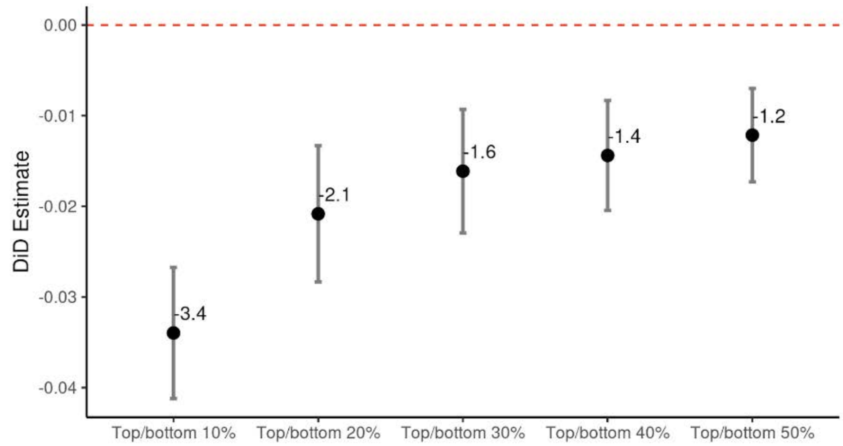

Notes: The figure presents the DiD coefficients of the impact of ATT on the advertiser probability of exit. Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

Figure C.3: The Impact of ATT on Advertiser Ad Spend by Variations of the Analytical Sample

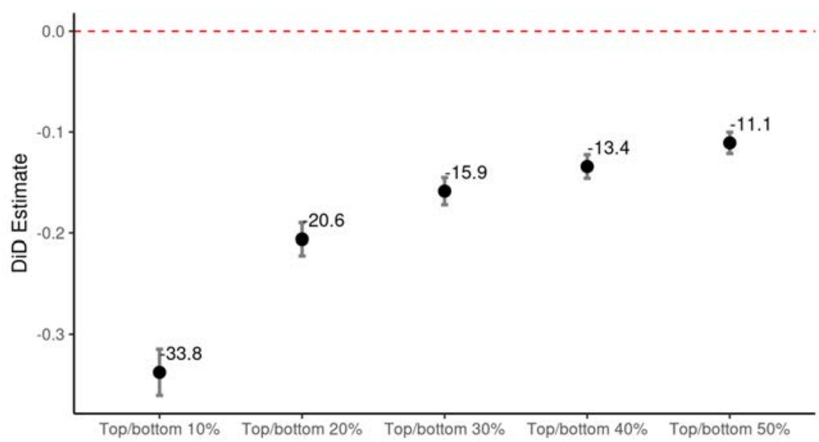

Notes: The figure presents the DiD coefficients of the impact of ATT on advertiser ad spend. Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

Figure C.4: The Impact of ATT on Industry Net Exit by Variations of the Analytical Sample

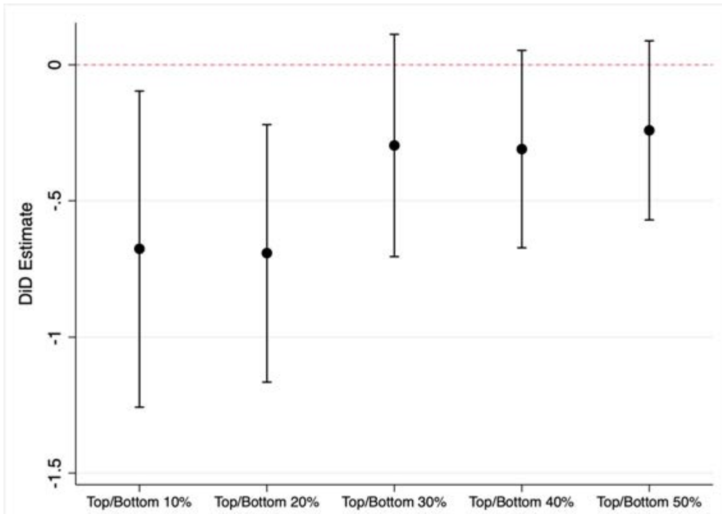

Notes: The figure presents the DiD coefficients of the impact of ATT on the industry Net Exit. Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

Figure C.5: The Impact of ATT on Industry Ad Spend per Firm by Variations of the Analytical Sample

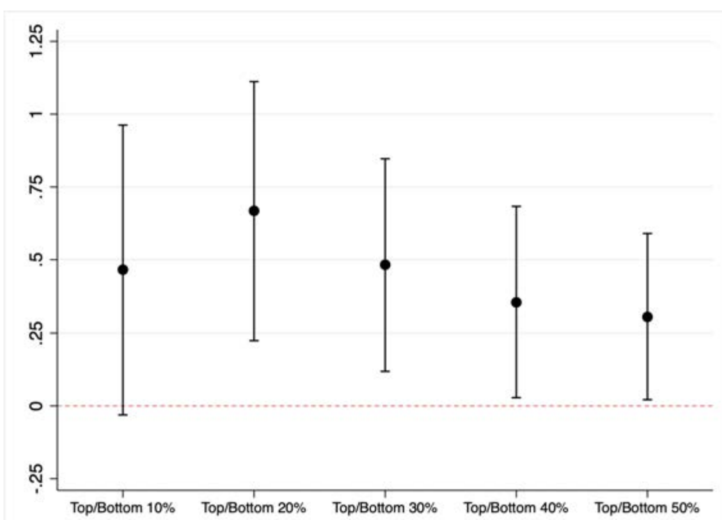

Notes: The figure presents the DiD coefficients of the impact of ATT on the industry ad spend per firm. Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

Figure C.6: The Impact of ATT on off-platform number of establishments by Variations of the Analytical Sample

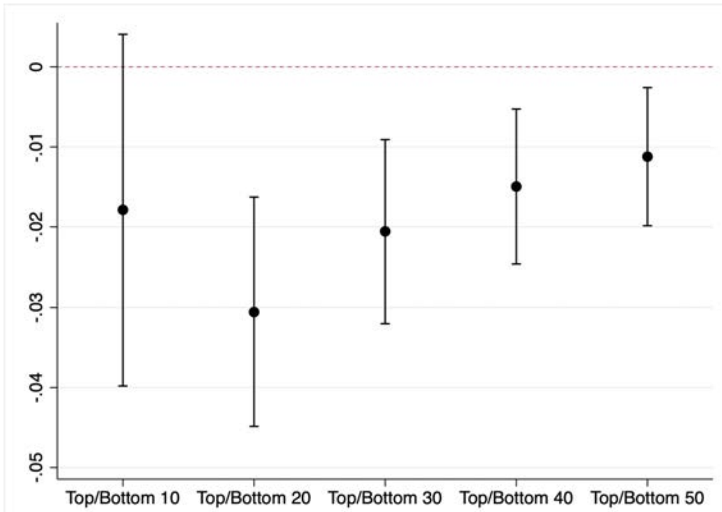

Notes: The figure presents the DiD coefficients of the impact of ATT on off-platform number of establishments. Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

Figure C.7: The Impact of ATT on Off-platform Prices by Variations of the Analytical Sample

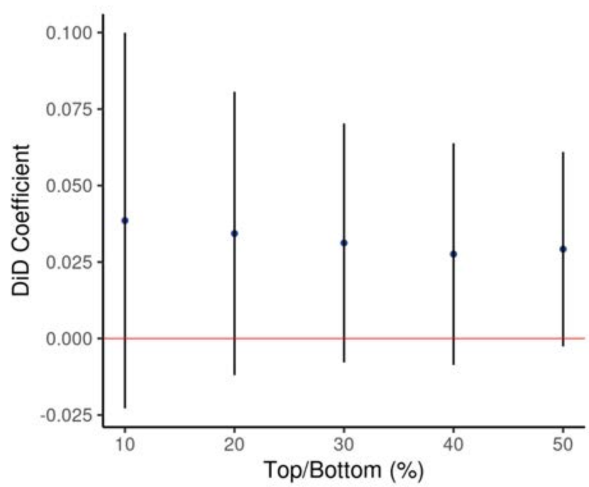

Notes: The figure presents the DiD coefficients of the impact of ATT on off-platform price (PPI). Each point represents a separate regression which include the top/bottom n% of the most affect advertisers ranging from 10% to 50%

D Various Specifications and Samples

D.1 Alternative Specifications

In this subsection, we test the robustness of our results to alternative specifications, such as using as reference points dates closer to the ATT implementation rather than announcement for our on-platform analysis, and dropping various industries.

Figure D.1: The Impact of ATT by various specifications

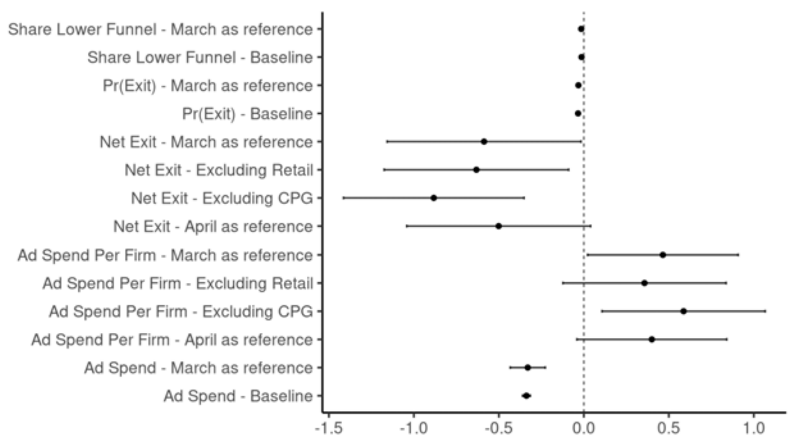

Notes: The figure represent DiD coefficients from various regression specifications using both advertiser-level and industry-level data. Advertiser level outcomes include share of lower funnel campaings, the probability of exit and ad spend, while the industry level outcomes include net exit and ad spend per firm.

Table D.1: Robustness of Estimates to Various Specifications

Outcome	Specification	Coefficient
Pr(Exit)	Baseline	-0.0339
		(0.0012)
	March as reference	-0.0319
		(0.004)
Share Lower Funnel	Baseline	-0.01358
		(0.001)
	March as reference	-0.0154
		(0.003)
Ad Spend	Baseline	-0.337
		(0.011)
	March as reference	-0.330
		(0.052)
Net Exit	Excluding CPG	-0.883
		(0.271)
	Excluding Retail	-0.632
		(0.277)
	March as reference	-0.587
		(0.291)
	April as reference	-0.501
		(0.276)
Ad Spend Per Firm	Excluding CPG	0.587
		(0.245)
	Excluding Retail	0.357
		(0.245)
	March as reference	0.465
		(0.226)
	April as reference	0.400
		(0.225)

D.2 Synthetic Control Estimates

For our on-platform advertiser analysis, we observe a small violation of parallel trends. While the pre-trends trend in direction opposite to the post-ATT trends, we nevertheless check the robustness of our estimates in two ways. First, we consider a distinct sample of advertisers active in Early 2020 and re-run the synthetic control analysis described in Xu (2017). This method extends the canonical synthetic control method to allow for multiple treatment units and allows for time varying coefficients. Consistent with our main results, we detect a 1.6 percentage point decrease in the share of lower funnel optimized ad campaigns (see fig D.2) and 5.2 percentage point increase in the probability of exit for the most affected advertisers relative to the least affected advertisers (see figure D.3).

Figure D.2: The Impact of ATT on the Share of Lower Funnel Campaigns

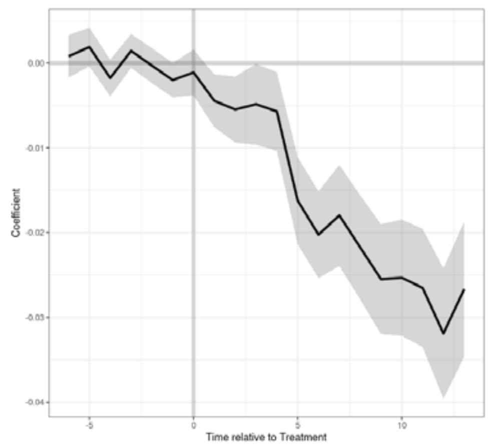

Notes: The figure plots the average effect of ATT on the share of lower funnel ad campaigns for the treated advertisers relative to the synthetic control created using the procedure outlined in Xu (2017).

Figure D.3: The Impact of ATT on the Probability of Exit

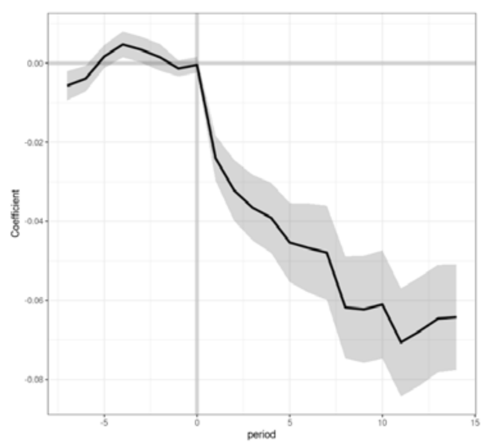

Notes: The figure plots the average effect of ATT on the probability of being active for the treated advertisers relative to the synthetic control created using the procedure outlined in Xu (2017).

Second, we implement the Honest DiD approach by Rambachan and Roth (2023) on our advertiser data and find that our first post-period estimates are robust to allowing for the violation of parallel trends up to twice as large as the max violation in the pre-treatment period (see figure D.4)

Figure D.4: Implementing HonestDiD on Advertiser Probability of Exit

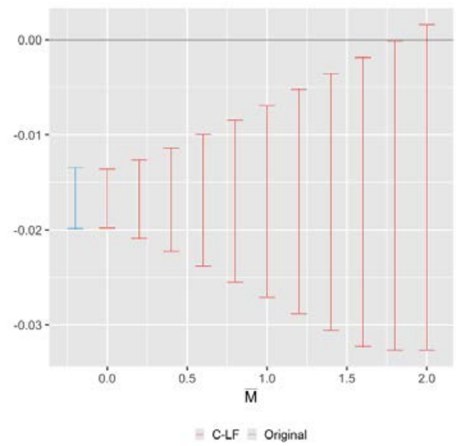

Notes: The figure implements the HonestDiD approach outlined in Rambachan and Roth (2023) on the probability of exit at the individual advertiser level. The figure shows that the post ATT estimates are robust to allowing for the violation in the pre-ATT parallel trends to be twice as large.

E Additional Robustness Checks

E.1 Advertiser Probability of Exit

Our findings suggest that advertisers most affected by ATT are more likely to stop advertising compared to advertisers that are less affected by ATT. Figure E.1 presents the main event study estimates. We observe a discontinuous increase in the probability of exit among advertisers in the treatment group, relative to the comparison group, after ATT comes into effect. Consistent with our results on ad spend, the event study illustrates an increase in the intensity of the impact of ATT over time. The difference-in-differences estimate implies a 3.4 percentage point higher probability of exit among the top 10% most affected advertisers, relative to the bottom 10% least affected advertisers.

Figure E.1: The Impact of ATT on Advertiser Probability of Exit

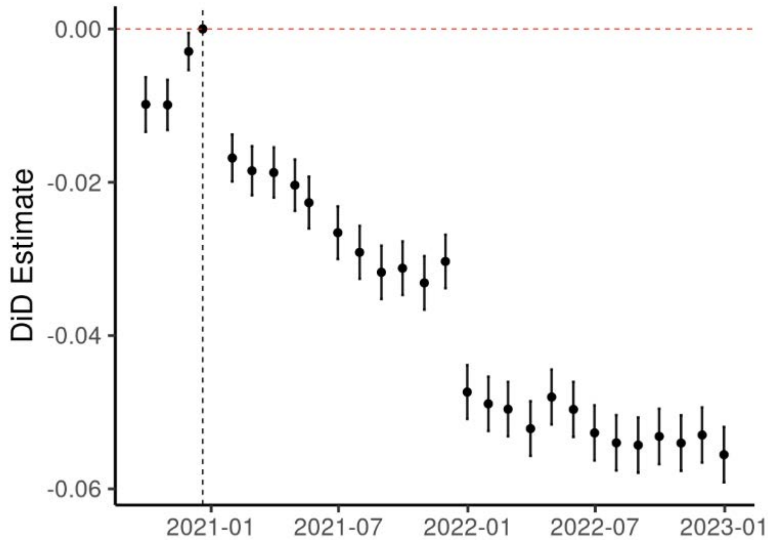

Notes: The figure plots the average effect of ATT on the probability of being active using advertiser level data.

E.2 Net Exit- Aggregated Industry Analysis

We consider an alternative approach where our level of analysis is an aggregated industry, rather than an industry, on Meta's platforms. As we note above, Meta classifies all on-platform advertisers into 25 aggregated industries. We, therefore, compare advertiser net exit across the top 5 most impacted aggregated industries, relative to the bottom 5 least affected aggregated industries. Consistent with our primary findings, this analysis finds that ATT increases advertiser net exit by 43%. Figure E.2 below presents the event study for these estimates. Further, these results are also driven entirely by small advertisers, which experience a 51% higher net exit. Large advertisers, as with our benchmark industry-level results, experience an increase in net entry; and at the aggregated industry level, the increase in net entry is sizable and statistically significant.

Figure E.2: The Impact of ATT on Net Exit using Aggregated Industry Level Data

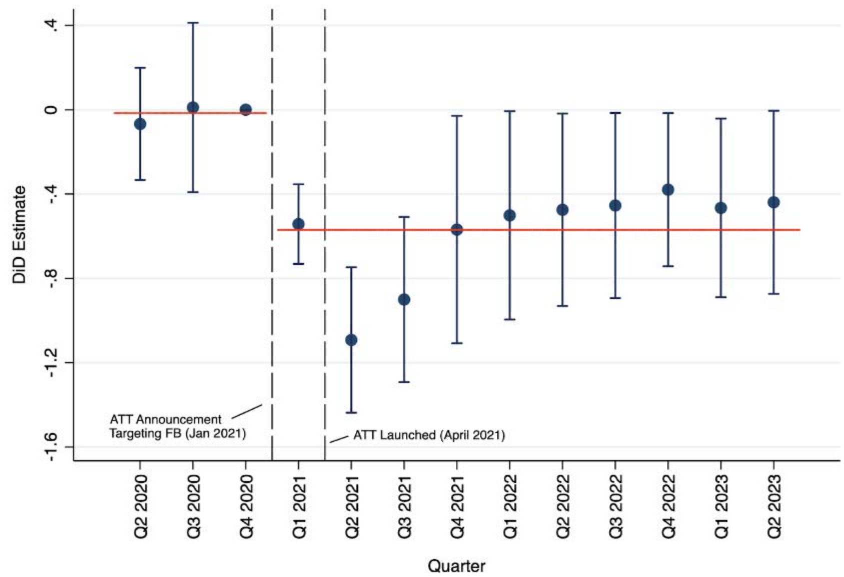

Notes: The figure plots the average effect of ATT on Net Exit using Aggregated Industry Level Data

F Industry Impact Rates

Table F.1: Impact rates for top and bottom 10% most impacted industries

Least Impacted Industries	Impacted Share	Most Impacted Industries	Impacted Share
Gaming - eSports	0.305	Consumer Packaged Goods - Food	0.573
Telecom - Telephone Service Providers	0.312	Retail - Footwear	0.574
Government - Government Owned Media	0.34	Retail - Department Store	0.575
Organizations and Associations - Religious	0.347	Professional Services - Fitness	0.579
Business to Business - File Storage, Cloud, and Data	0.366	Consumer Packaged Goods - Water, Soft Drink, and Beverage	0.581
Travel - Ride Sharing and Taxi Services	0.367	Restaurants - Quick Service	0.59
Government - Offices of Heads of States and Governments	0.368	Ecommerce - Apparel and Accessories	0.59
Ecommerce - Virtual Services	0.368	Ecommerce - Beauty	0.593
Telecom - Telecommunications Equipment and Accessories	0.369	Consumer Packaged Goods - Durable Household Goods	0.599
Business to Business - Logistics and Transportation	0.37	Consumer Packaged Goods - Non Durable Household Goods	0.6
Politics - Political Organizations and Campaigns	0.379	Consumer Packaged Goods - Personal Care	0.602
Utilities - Mining and Quarrying	0.386	Consumer Packaged Goods - Beauty	0.604
Gaming - Mobile Gaming	0.387	Retail - Home, Furniture, and Office	0.608
Automotive - Auto Resellers	0.4	Travel - Air	0.61
Professional Services - Legal	0.401	Restaurants - Coffee	0.611
Publishing - Career and Tech	0.402	Retail - Sporting	0.622
Technology - Social Media	0.403	Organizations and Associations - Arts, Heritage, and Education	0.625
Government - International Organizations	0.405	Consumer Packaged Goods - Baby	0.628
Government - Government Controlled Entities	0.41	Consumer Packaged Goods - Apparel and Accessories	0.632
Healthcare - Residential Care Centers	0.411	Retail - Apparel and Accessories	0.654