

Cheyre, Cristobal; Leyden, Benjamin T.; Baviskar, Sagar; Acquisti, Alessandro

Working Paper

The Impact of Apple's App Tracking Transparency Framework on the App Ecosystem

CESifo Working Paper, No. 10456

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Cheyre, Cristobal; Leyden, Benjamin T.; Baviskar, Sagar; Acquisti, Alessandro (2023) : The Impact of Apple's App Tracking Transparency Framework on the App Ecosystem, CESifo Working Paper, No. 10456, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/279205>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Impact of Apple's App Tracking Transparency Framework on the App Ecosystem

Cristobal Cheyre, Benjamin T. Leyden, Sagar Baviskar, Alessandro Acquisti

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

The Impact of Apple's App Tracking Transparency Framework on the App Ecosystem

Abstract

We study the impact of the implementation of Apple's App Tracking Transparency (ATT) framework on the Apple App Store ecosystem. We use comprehensive data on every app available in both the Apple App Store and Google Play Store ecosystems in the eighteen-month period around the implementation of ATT, and a difference-in-differences analysis to investigate whether the introduction of the privacy transparency framework affected the incentives for developers in the Apple ecosystem to create new apps, update their existing apps, or withdraw from the market. We also leverage data on the presence of Software Development Kits (SDK) in a select number of apps in each ecosystem to study how developers adapted specific functionalities in their products, such as the use of advertising platforms or payment systems. We find that the number of available apps in the Apple App Store ecosystem quickly recovers after an initial drop following the introduction of ATT. When analyzing the use of SDKs, we find a reduction in the use of Monetization and Ad Mediation SDKs, and an increase in the use of Authentication and Payments SDKs. Our results suggest developers did not withdraw from the market after ATT and instead adapted to operate under the conditions of a more protective privacy framework.

Cristobal Cheyre

*Cornell University / Ithaca / NY / USA
cac555@cornell.edu*

Benjamin T. Leyden

*Cornell University / Ithaca / NY / USA
leyden@cornell.edu*

Sagar Baviskar

*Carnegie Mellon University
Pittsburgh / PA / USA
sbaviska@andrew.cmu.edu*

Alessandro Acquisti

*Carnegie Mellon University
Pittsburgh / PA / USA
acquisti@cmu.edu*

May 2023

Working paper - comments welcome.

Click here for the latest version: <https://www.benleyden.com/att>

We thank Carol Lu, Yoonseock Son, and conference participants at the Research Roundtable on Regulating Privacy at George Mason University, WISE 2022 and the 21st IIOC for their constructive feedback. Acquisti, Baviskar, and Cheyre gratefully acknowledge the financial support from the Program on Economics & Privacy at George Mason University.

1 Introduction

Online user tracking—that is, the collection of users’ demographic, psychographic, and behavioral data—has proliferated in the mobile app ecosystems, as the \$154 billion (in the US) mobile advertising industry heavily relies on personal-level consumer data acquired through mobile apps to serve targeted advertisements (PwC and IAB, 2023). Online user tracking may be a necessity for app ecosystems to prosper: tracking allows apps to show personalized ads and sell user data to third parties (Mhaidli et al., 2019; Ekambaranathan et al., 2021), generating revenues that enable developers to create and make available apps to users at low (or zero) prices (Ribera, 2022). At the same time, increasing online tracking has led to widespread privacy concerns that have prompted regulatory bodies and private companies to implement interventions aimed at protecting consumers’ privacy. Assessing the costs and benefits of such interventions has proven difficult. In this study, we leverage the implementation of Apple’s App Tracking Transparency (ATT) framework to analyze how restricting the ability of apps to track users may (or may not) affect the survival of existing apps and the creation of new ones, and the tracking and monetization strategies used by app developers.

In June 2020 Apple announced that it was working on a new App Tracking Transparency (ATT) framework to be introduced in a future version of iOS. ATT changed Apple’s privacy and data collection policy that governs how end users choose whether an app can track their activity across other companies’ apps and websites for the purposes of advertising or sharing with data brokers. Under ATT, user tracking for targeted advertising switched from an opt-out to an opt-in basis, as apps on Apple devices must ask for (and obtain) permission before tracking users’ activity across other companies’ apps and websites. Despite strong opposition from technology firms that derive most of their revenue from advertising, including Facebook buying full-page ads in major newspapers attacking the framework (see Figure 4 in the appendix)¹, Apple announced on April 20, 2021 that ATT would be released as part of iOS 14.5 the following week.

The introduction of ATT was met with concern by parts of the online advertising industry. In order for online ads to be effective for advertisers and valuable (i.e., revenue-generating) for firms selling advertising space (such as online publishers and game developers), ads must be as precisely targeted to the correct audiences as possible. This requires precise user tracking. Before the implementation of ATT, app developers and data brokers were able to track users at a very granular level by leveraging the Identifier for Advertisers (IDFA) cross-app tracking identifier, third-party data, and data sharing agreements between companies across various apps. With the introduction of ATT, app developers and data brokers faced

¹<https://www.engadget.com/facebook-ios-14-privacy-changes-new-add-090000011.html>

more significant challenges to track users at the same level as they used to, especially using third-party data. As anecdotal reports suggest that only a minority of Apple users consented to tracking under ATT, fears were raised that ATT would cause major losses to companies that derive revenues from online advertising. In turn, and due to those losses, consumers may ultimately suffer from the disappearance of ad-reliant apps and services.

Whether such fears did materialize, however, is still an open research issue. We investigate whether and how the introduction of ATT affected the app ecosystem in terms of availability and frequency of updates of apps, consumers’ valuations of apps as measured by the number and average of customer ratings, and the use of software development kits (SDKs) by developers. Our analysis leverages information from a data provider that tracks the universe of apps on both the Apple App and Google Play stores to implement a difference-in-differences strategy.

Our results show that ATT led to a temporary reduction in the entry of new apps to the Apple ecosystem, but that the effect is short-lived as it dissipates a few months after the policy. There is no clear effect on exit, though developer effort via product updates falls after ATT is implemented. This pattern suggests that ATT did not reduce overall developer interest in the platform, but that firm investment may have fallen. Relatedly, we find that apps receive fewer and lower ratings compared to before ATT, an effect that appears to be driven by the response of incumbent apps to the policy change. Finally, in terms of SDK usage, we observe a reduction in the use of data-intensive advertising SDKs, and an increase in the use of authentication and payment SDKs.

2 Related Literature

Since the enforcement of ATT, researchers have sprouted an interest in analyzing and understanding the effectiveness and impact of ATT on apps and other stakeholders in the ecosystem. [DeGiulio et al. \(2021\)](#) study how mobile apps present tracking requests to users and evaluate the observed design patterns impact on users’ privacy, finding that opt-in authorizations are effective at enhancing data privacy, and that the effect of ATT requests is robust to most implementation choices. [Li and Tsai \(2022\)](#) look at how the inability of apps to use tracking for advertising reduced new downloads, affecting to a greater degree large rather than small apps. [Kesler \(2022\)](#) studies whether and how app developers changed their monetization strategies following the implementation of ATT, finding a small increase in the number of paid apps and apps that offer in-app purchases. [Kollnig et al. \(2022\)](#) analyzes the impact of ATT implementation on data brokers and app makers, finding that the new policy is effective in preventing the collection of the IDFA cross-app tracking identifier. As

a result, they observe that data brokers are facing higher challenges in tracking users, and apps are starting to collect device information that can be used to track users at a group level (cohort tracking) or identify individuals probabilistically (fingerprinting). [Sokol and Zhu \(2021\)](#) consider the implementation of ATT as an anti-competitive strategy and how it is harmful not only for a fair competitive market but potentially for end users as well. [Jones \(2022\)](#) considers ATT to be utilized as an overarching guide, along with European General Data Protection Regulation (GDPR), for US government when enacting comprehensive federal privacy law.

Some researchers have explored the importance of SDKs for the app development process. [Kim and Wagman \(2021\)](#) identify robust effects of SDK releases on app development and provides consumer surplus estimates associated with apps. [Alomar and Egelman \(2022\)](#) investigated the privacy compliance processes followed by developers of child-directed mobile apps. Their results suggest that most developers rely on app markets to identify privacy issues, they lack a complete understanding of the third-party SDKs they integrate, and they find it challenging to ensure that these SDKs are kept up to date and privacy-related options are configured correctly. It is thus important to analyze the effect that ATT had on the use of SDKs, as they embody the dichotomy between creating welfare gains by easing the development process, but often at the cost of privacy and security.

Before the introduction of ATT, several studies started investigating how the introduction of GDPR, which restricted the ability of app developers to track users, has affected the app ecosystem and other ad-supported businesses. The effects estimated by these studies vary widely. [Kollnig et al. \(2021\)](#) studies third-party tracking in mobile apps from before and after the introduction of the GDPR. Their findings suggest that there has been limited change in the presence of third-party tracking in apps and that the concentration of tracking capabilities among a few large gatekeeper companies persists. [Janssen et al. \(2022\)](#) estimate that GDPR led to the exit of about a third of apps in the Google Play store, and reduced entry by half. In contrast, [Lefrere et al. \(2022\)](#) study online publishers and find that, although they implemented changes following GDPR, there is no evidence that the regulation inhibited publishers' ability to produce content or generate user engagement. Out of a broad sample of news and media websites they analyze, they find that almost none of them exited the market after GDPR. [Momen et al. \(2019\)](#) analyze app behavior before and after GDPR and shows that the app/user privacy has moderately improved after the implementation of the regulation.

We contribute to this literature by analyzing how ATT affected the incentives for developers to participate in the ecosystem, whether the changes induced by ATT led to any change in valorization of apps by consumers, and whether it influenced the reliance of devel-

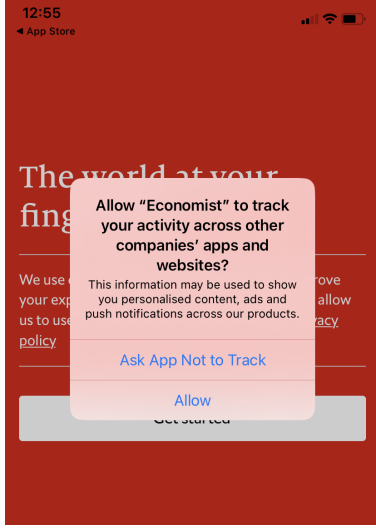
opers on third-party SDKs for incorporating functionalities in the app (such as behaviorally targeted advertising). Our context presents distinct advantages relative to studies that used GDPR to investigate the economic consequences of privacy initiatives. While GDPR has been marred by delayed and inconsistent enforcement, along with a lack of clarity over the parameters and methodologies for achieving its goals (Bygrave, 2017), ATT provides a clean, exogenous shock that app developers cannot avoid. Additionally, GDPR applies to all entities processing data from EU citizens, which makes it difficult to distinguish between treated and control entities. In contrast, ATT only applies to Apple devices, which makes it straightforward to compare the evolution of the Apple App Store ecosystem versus the Google Play Store ecosystem, which has not yet been affected by a similar policy, in order to study ATT’s effect.

3 Theoretical Framework

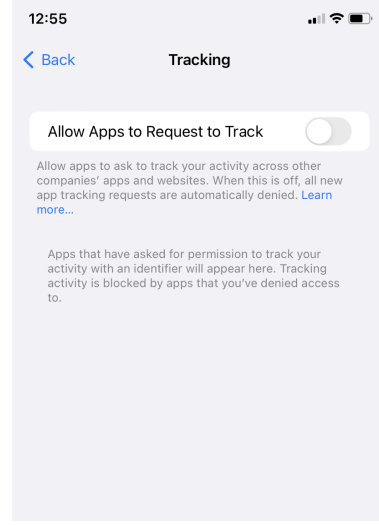
The introduction of Apple’s ATT effectively shifted consent for the collection of end-user data by apps, and the sharing of such information with third parties, from an opt-out to an opt-in basis on Apple devices. Starting with iOS 14.5, developers wishing to track users to serve them with behaviorally targeted advertising must present an app-tracking authorization request (see Figure 1a). If the authorization is not granted, the operating system does not share with the app the device’s IDFA, which prevents the app from accessing whatever information third parties may have about the user, and it further prevents the app from contributing information about the end-user to third parties. Moreover, a new privacy setting was incorporated in Apple’s operating system that allows users to opt out from even allowing apps to request permission to track (see Figure 1b). Anecdotal reports on the number of users that have chosen to opt-out vary, but they all suggest that a large share of users are choosing not to allow tracking.²

The changes introduced by ATT have the potential to significantly reduce the amount of user data apps (and the services they connect to) can leverage for advertising, which in turn may reduce the profitability of ad-supported apps, as content providers selling non-targeted advertising typically receive lower payments per impression (Sharma et al., 2019). Thus, if not enough users of an app consent to tracking, and the decline in revenue for the higher prevalence of non-targeted ads is significant, some developers may choose to stop investing

²For example <https://www.mediapost.com/publications/article/373613/privacy-update-att-idfa-opt-in-rate-at-25-overall.html> reports opt-in rates of about 25%. Other reports, such as <https://arstechnica.com/gadgets/2021/05/96-of-us-users-opt-out-of-app-tracking-in-ios-14-5-analytics-find/>, report 95% of users are opting out.



(a) App-tracking authorization request



(b) Privacy setting to prevent apps from requesting authorization to track

Figure 1: Apple’s Tracking Authorization Settings

in further developing and maintaining existing apps, or choose not to continue investing in the development of new apps. Notably, even if a particular app’s users consent to tracking, if enough users in the ecosystem are blocking tracking in general, there will be a broader ecosystem effect, in which even targeted advertising may become less valuable if the amount of information available for making inferences decreases and targeting becomes less precise (Lefrere et al., 2022).

We can expect developers of free, ad-supported apps to be affected by ATT in two different ways. As we have argued, they may start obtaining lower revenues from advertising as users start to decline requests to track and non-targeted ads become more prevalent. This can lead them to switch their monetization strategy to increase their reliance on in-app purchases, completely switch to a paid model, or abandon the ecosystem. Additionally, they may face increased difficulties in attracting new downloads, as one of the ways developers use to make their apps known and users to download them is to advertise in other apps, or in the app markets. If users can no longer be tracked to the same extent, it will be harder to identify users potentially interested in a new app and reach out to them through targeted advertising. If this leads to fewer installations, the app will have fewer users to serve ads to and will obtain lower revenues.

Determining the impacts ATT may have on end-users is more challenging. If ATT decreases the availability of new apps and/or the continued investment of developers with existing apps, consumers may suffer harm from the reduction of variety and/or quality. However, not all apps provide the same surplus to users and so a reduction in variety might


not affect consumer surplus in a meaningful way. Thus, to the extent that a privacy initiative such as ATT reduces the availability of apps, it is important to consider the characteristics of the apps exiting the market, or not entering the market. Characterizing apps that abandon the market is easy, as their characteristics are observable. In fact, in the case of the GDPR, [Janssen et al. \(2022\)](#) show that after the regulation became effective many apps exited the market, but, for the most part, it was apps that were not being used or were not valued by users. Determining the attributes of apps that do not enter because of the new policy is more challenging. [Janssen et al. \(2022\)](#) argue that the quality of new apps is difficult to predict ex-ante, and therefore that GDPR, to the extent that it discourages entry, will discourage the entry of high and low-quality apps to a similar extent. Thus, a loss of entry is likely to be more costly to consumers than a corresponding increase in exit. However, if developers do have a reasonable ex-ante notion of the likelihood of success of an app, or the risks involved in the development, we should expect the lost entrants to be lower quality on average. In this case, the consumer welfare losses associated with entry and exit might be more closely aligned. Of course, in addition to the losses due to product variety and developer investment, we also need to consider the additional value consumers may receive through the additional level of privacy offered by ATT.

To measure the impacts of ATT on developers and users, we analyze the evolution of the availability of apps in the Apple ecosystem, relative to the Android ecosystem, before and after the implementation of ATT. To examine if ATT changes the incentives developers have for investing on creating and maintaining apps, we examine whether ATT has affected the rate of entry, exit, and update frequency. While we cannot directly measure demand and consumer surplus, we rely on the number of ratings and average ratings received by apps as a proxy of demand and quality. Thus, to study if ATT has affected the quality and/or the valuation consumers have of apps, we examine if ATT had an effect on the number of ratings received by apps over time, and the average rating they receive. Finally, to determine to what extent ATT has changed the way apps rely on end-user data for monetization, we examine how their use of Software Development Kits (SDKs) typically associated with the collection and use of user data has changed after ATT.

4 Data

We use data from Airnow, which collects data on all apps in the Apple App Store and the Google Play Store, including information such as the name and genre of each app, its developer, its entry and exit dates, the number of reviews and average ratings it received over time, among others. Since 2015, Airnow has collected daily data on more than 19

Figure 2: Number of Free Apps in the Apple App Store and Google Play Store Markets



million apps. Additionally, for a smaller number of apps, Airnow provides information on the software developer kits (SDKs) each app uses. SDKs are development tools that allow developers to include features in their apps in a simple and standardized way. They are typically provided by third parties that have an interest in the developer including their features in their apps. For example, advertising technology platforms, such as Google and Facebook, provide SDKs to allow developers to easily include advertising in their apps. To collect data on the SDKs used by developers, Airnow periodically downloads a number of apps from both ecosystems and analyzes the SDKs included in them. They classify the SDKs in a number of categories based on the functionality of the SDK including monetization, ad mediation, payments, authentication, and others.

We restrict our analysis to those apps that set a price of \$0. In Table 1, we present summary statistics for key variables used in our analysis, separated by Apple’s App Store and the Google Play Store. Across the two platforms, we see just over 7 million products. Entry and exit on each platform roughly balance out during our sample period. In terms of SDK usage by developers, we see that Monetization SDKs are by far the most used, followed by Authentication SDKs. Ad Mediation technologies, which would be mostly directly impacted by the introduction of ATT are similar across platforms.

One month before the implementation of ATT, there were 1.6 million apps in the Apple App Store, and 3.3 million apps in the Google Play Store. Figure 2a shows how the number of *free*, active apps in each ecosystem has varied widely over the years, and has apparently been affected by different events over time, the most notable of them being the implementation of the EU GDPR that led to a decline in the number of active apps in both ecosystems. By 2020, the number of active apps in both ecosystems seem to have stabilized. Considering the frequent fluctuations each ecosystem experiences over time, which may be driven by

Table 1: Summary Statistics

	App Store		Google Play Store		N
	Mean	SD	Mean	SD	
Log Entry Count	7.19	1.12	8.60	0.83	91,149,638
Log Exit Count	7.66	0.78	8.62	0.84	91,149,638
Update	0.10	0.30	0.07	0.26	91,149,638
Log New Ratings	0.19	0.78	0.42	1.11	91,149,638
Avg. New Rating	3.98	5.47	3.70	93.89	15,543,166
Monetization	1.14	2.01	1.23	1.55	8,491,149
Ad Mediation	0.16	0.88	0.27	1.48	8,491,149
Authentication	0.47	0.75	1.14	0.90	8,491,149
Payments	0.12	0.46	0.58	0.74	8,491,149

extraneous confounding effects, we focus our analysis in the 18 months period around the implementation of ATT. Figure 2b shows the number of active apps in each ecosystem from 6 months before to 6 months after ATT. A drop in the number of active apps in the Apple ecosystem immediately after ATT became effective is observed, however, that drop was followed by a recovery that brought the number of active apps to above the pre-ATT level.

5 Empirical Analysis

We use a difference-in-differences (DiD) framework to study the effect of ATT on outcomes related to the state of the Apple App Store ecosystem including the entry, exit, and update frequency of apps, as well as the effect of the framework on outcomes experienced by apps, such as the number of new ratings they receive and changes to their average ratings, and how it influenced the functionalities included by developers in their apps through the use of SDKs. As the implementation of ATT was imposed on developers by Apple without much advance notice, we consider it as an exogenous event. Thus, examining the evolution of the Apple App Store, which was affected by ATT, versus the Google Play Store, which wasn't affected by a similar policy during the period we study, gives us an opportunity to estimate the causal effect of imposing a policy restricting the use of personal data on apps responses and the outcomes they experienced, as well as on the evolution of the ecosystem in general. Note that while the framework had been announced in advance, developers didn't know when it would be implemented and enforced. Effectively, developers received only a single week's advance notice of the implementation date for ATT via the release of iOS 14.5.

5.1 Market Dynamics: Entry, Exit, Updates, and Ratings

We first analyze how ATT influenced the dynamics of the market by analyzing its impact on entry, exit, and the frequency of app updates among free apps on the platform.³ To study the effects of ATT on entry and exit, we calculated genre-level entry and exit counts on each platform. We then estimate the difference-in-differences model

$$Y_{p,g,t} = \alpha_1 PostATT_t \times Apple_p + \delta_g + \mu_{g,t} + \epsilon_{g,t}, \quad (1)$$

where p indicates a platform, g indicates a genre (i.e., product category), and t indicates a month. Given a platform-genre-level entry or exit count $C_{p,g,t}$, we construct our outcome variables as $Y_{p,g,t} = \log(C_{p,g,t} + 1)$. $PostATT_t$ is equal to one if period t occurs on or after the release of ATT, and $Apple_p$ is equal to one if p indicates Apple’s App Store. Thus, α_1 is the coefficient of interest, measuring the degree to which entry or exit on Apple’s platform responds to the introduction of ATT. Finally, δ_g and $\mu_{g,t}$ are genre and month fixed effects. Months are defined as four-week intervals relative to the release of ATT.

For updating frequency, and the remaining outcomes considered in this paper, we leverage our full, app-level panel data set. Specifically, we estimate

$$Y_{i,t} = \beta_1 Post ATT_t \times Apple_i + \omega_i + \mu_t + \epsilon_{i,t} \quad (2)$$

where $Y_{i,t}$ is an indicator for whether app i updated in month t . Similar to Equation (1), β_1 is the coefficient of interest, and ω_i and μ_t are app and month fixed effects.

In column (1) of Table 2, we see that entry declined on the App Store relative to the Google Play Store following the introduction of the ATT policy. This is consistent with the fact that developers’ abilities to collect revenue through targeted advertising and other consumer tracking technologies was hampered by the new policy. However, panel (a) of Figure 3 shows that this finding is entirely driven by an initial (and large) decline in entry over the first two months of the policy, after which entry rates return to normal.

In column (2) of Table 2, we fail to find evidence of an effect on exit behavior on the App Store following ATT. While exit does occur on the platform, as documented in Table 1, it is common, particularly among top-performing apps, for a product to remain available for sale on the store long after the developer has stopped actively developing the product. This is because it’s free to keep an app on the store as long as the developer maintains an active developer account (which cost \$99/year for the App Store, and a one-time \$25 fee for the

³The analysis presented here considers only those apps that charge an upfront price of \$0 throughout the entire sample period. In unreported analysis, we find that the results of this paper are not sensitive to the inclusion of paid apps.

Table 2: Impact of ATT on Entry, Exit, and Updates

	Log Entry Count (1)	Log Exit Count (2)	Update (3)
After ATT x Apple	-0.1689*** (0.0607)	0.0288 (0.0360)	-0.0058*** (0.0001)
Platform-Genre FE	✓	✓	
Period FE	✓	✓	✓
App FE			✓
R ²	0.91317	0.96656	0.37519
Observations	817	817	91,149,638

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Google Play Store). Given this, looking at update frequency can be informative about the level of effort developers are engaged in on a platform, which encompasses the rate of quality improvements as well as whether a product is being actively maintained. Moreover, [Leyden \(2023\)](#) has previously shown that developers’ updating behavior is sensitive to the design and policies of a platform.

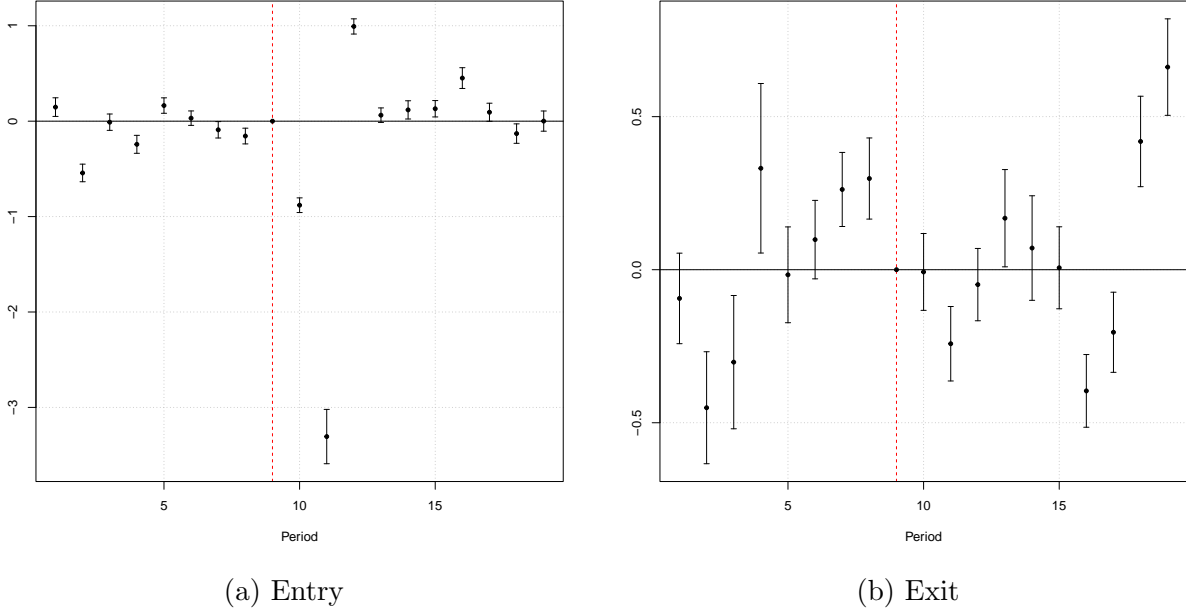
We present results for updating frequency in column (3) of Table 2. We see that updating falls slightly in response to ATT. While updating slightly declines on average, we find that this effect varies in both magnitude and sign across genre, as documented in column (1) of Table 9 where we re-estimate Equation (2) while allowing β_1 to vary by genre.⁴

We next turn to analyzing whether users’ valuation of apps changed after ATT. While we would like to examine the number of downloads on each platform and the consumer surplus users obtain from those apps, that information is not observable. Instead, we examine the effect of ATT using the number of new ratings received by apps each month, a proxy for downloads, and the average score of those ratings, a measure of product quality.

Theoretically, the relationship between ATT and consumers’ demand and valuation of apps is difficult to untangle. From one side, if consumers value transparency and control in-app privacy choices, and if ATT has been effective in reducing the exposure to invasive advertising to those that dislike it, the policy may lead to an increase of users’ valuations, and thus of ratings. From the other side, if ATT leads developers to stop updating then the

⁴The set of genres, or product categories, varies slightly different between the Apple and Google ecosystems. For this analysis, we create a mapping across platforms.

Figure 3: Entry and Exit Over Time



quality of apps may be affected. We find in Table 3 that the reduced effort by developers (as evidenced by the declining update frequency we observe in Table 2) dominates, as apps in the Apple ecosystem receive fewer and lower ratings after ATT. Additionally, apps are more likely to receive no ratings in a particular month following ATT.

We next study whether the number of ratings and average ratings received by apps that enter before vs. after ATT are different. We consider this in two steps. First, we re-estimate Equation (2) for these outcomes using only those apps that were released prior to the implementation of ATT. We present these results in Table 4, where we see that similar to the result in the overall sample, the number of ratings and the average of any new ratings for existing apps both declined following the release of ATT, and the likelihood of an app receiving no new ratings increased.

Second, we consider how this policy affected app ratings in the first month an app was available for sale. To do so, we compute the number of ratings each app received during its first full month listed, the average rating they receive in that same period, and whether they receive no ratings during the first month. Note that as each app is only observed for its first month of listing we cannot use app fixed effects. Instead, we use fixed effects by the interaction of platform and genre.

We present the results of estimating this augmented version of Equation (2) with this sample in Table 5. The results show that apps released *after* ATT received more ratings

Table 3: Effect of ATT on the Number of Ratings and Average Rating

	Log # New Ratings (1)	Avg. Rating (2)	No Rating (3)
After ATT x Apple	-0.0400*** (0.0002)	-0.2063* (0.1089)	0.0277*** (0.0001)
App FE	✓	✓	✓
Period FE	✓	✓	✓
R ²	0.79977	0.10083	0.59820
Observations	91,149,638	15,543,166	91,149,638

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Table 4: Effect of ATT on the Number of Ratings and Average Rating for Existing Apps

	Log # New Ratings (1)	Avg. Rating (2)	No Rating (3)
After ATT x Apple	-0.0398*** (0.0002)	-0.2124* (0.1108)	0.0276*** (0.0001)
App FE	✓	✓	✓
Period FE	✓	✓	✓
R ²	0.80783	0.10059	0.60101
Observations	85,243,409	14,904,632	85,243,409

The results in this table use a subsample of all apps on the App Store and Google Play Store that were released prior to the implementation of ATT.

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Table 5: Effect of ATT on the Number of Ratings and Average Rating in the First Full Month of Sale

	Log # New Ratings (1)	No Rating (2)	Avg. Rating (3)
After ATT x Apple	0.0934*** (0.0028)	-0.0532*** (0.0010)	-0.0047 (0.0191)
Platform-Genre FE	✓	✓	✓
Period FE	✓	✓	✓
R ²	0.05348	0.04475	0.00458
Observations	2,251,796	2,251,796	374,562

The results in this table use a subsample that consists of the first full month each app was available for sale.

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$

in the first month than those released before, and the probability that a new app attracts no ratings during its first month is lower after ATT. Thus, in light of the findings in the previous tables, we see that faced with the new requirements of ATT, developers seem to be choosing not to update their existing apps as frequently, and as existing apps are not updated, the interest on them appear to wane and users value them lower as they become outdated. And so the findings in Table 3 are primarily driven by incumbent apps, rather than by new entrants (who may have factored ATT into their original business model and the decision to enter).

5.2 Developers' Use of Software Developer Kits

In addition to studying how ATT may affect the dynamics of the app market, we are interested in investigating if the new policy had any effect on how developers construct their apps, as this can indicate how developers adjust their business model in response to ATT. We analyze whether ATT led to any change in the software developer kits (SDKs) used by developers. SDKs are third-party tools that developers use to include specific functionalities. We are particularly interested in analyzing the reliance of developers on Monetization SDKs (used to monetize apps through advertising), Ad Mediation SDKs (used to allocate ad impressions to the ad platform offering the best price), Authentication SDKs (that allow apps to let users log in using credentials from different platforms), and Payment SDKs (used

for securely processing card payments).⁵

Table 6: Examples of Software Developer Kits by Category

Monetization	Ad Mediation	Authentication	Payments
Facebook Audience Network	AdMob Mediation Adaptor	Facebook Login	Mastercard CBP
Google AdMob	ironSource Mediation Adaptor	Firebase Auth	Stripe
ironSource		Google Sign In	Square
SupersonicSDK		Validator	Venmo

We present examples of each category of SDK in Table 6. Monetization and Ad Mediation SDKs both rely heavily on user data, as that allows the connected ad platforms to serve precisely targeted advertisements. Thus, if a user does not consent to sharing data with third parties, the utility of these two categories of SDKs is hampered. In contrast, if apps after ATT start relying more on direct sales, we expect to see an increase in the use of Payments SDKs.

In Table 7, we present the results of estimating Equation (2) for each of these four SDK categories. We find that, as expected, Monetization and Ad Mediation SDKs decline in usage, while there’s a small increase in the take-up of Payments SDKs. The largest effect of the four is the decline in Monetization SDKs, which also represents the most-used category overall (see Table 1). These results show a clear effort by developers to reduce their reliance on data collection and targeting technologies, and possibly a shift to other forms of monetization.

In addition to the effects on ad and payment SDKs, we see evidence of an increase in the use of Authentication SDKs. Ex-ante, it isn’t clear what to expect in this case, as developers may be inclined to require customer logins in order to begin collecting and leveraging first-party data (which is permissible under ATT restrictions). However, the most prominent authentication SDKs are managed by companies managing large ad networks, like Facebook and Google, and so a developer’s attempts to rid themselves of their reliance on these companies may result in a reduction in the use of authentication SDKs. Ultimately, we find that developers become more likely to authenticate their users, which may be part of a broader strategy to better monetize users through first-party data collection and use.

This conforms, for the most part, to the expectations explained above. It is interesting to look at how the use of SDKs changes across genres of apps that depend on advertising that uses third-party tracking to different degrees. We present these results in Table 8, where we see that (similar to our earlier discussion of Table 9), responses vary in both magnitude and

⁵Within these SDKs it is worth highlighting the difference between Monetization and Ad Mediation SDKs. While both aim to deliver advertising, they do so in different ways. Monetization SDKs typically rely on a single platform to deliver ads, while the goal of Ad Mediation SDKs is to connect to multiple platforms to attempt to allocate an ad impression for the highest possible price.

Table 7: Impact of ATT on SDK Usage

	Monetization (1)	Ad Mediation (2)	Authentication (3)	Payments (4)
After ATT x Apple	-0.0250*** (0.0017)	-0.0021** (0.0008)	0.0054*** (0.0006)	0.0059*** (0.0004)
App FE	✓	✓	✓	✓
Period FE	✓	✓	✓	✓
R ²	0.88731	0.86644	0.89485	0.90313
Observations	8,491,149	8,491,149	8,491,149	8,491,149

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$

sign across genres. Of particular interest are the Gaming and Shopping genres. Gaming is the genre that is associated with the greatest use of Monetization and AdMediation SDKs. In contrast, Shopping apps use fewer of both. The results in Table 8 show a large reduction in both Monetization and Ad Mediation SDKs in Games, and an *increase* in the use of both in the Shopping category. One possible explanation for this is that under ATT, developers face no restriction in their use of first-party data. Shopping apps, perhaps more than any other company, will have access to troves of relevant, first-party data on their users, namely, sales data. As a result, these developers may have found it profitable to leverage the sudden change in their relative position in the user data ecosystem to their advantage. Notably, we also see a large increase in the use of Payments SDKs in the gaming category, as that category shifts towards other sources of revenue.

6 Conclusions

A frequent concern with the implementation of regulations and policies that limit the collection and sharing of user data is that they may harm the availability of free, ad-supported content, services, and applications if their creators can no longer effectively monetize their work. In this paper, we analyze how the implementation of Apple’s ATT framework has influenced the entry, exit, and update dynamics in the Apple App Store ecosystem, the quality of the products on the ecosystem, and the features included by developers in their apps through the use of SDKs.

Our analysis is based on a difference-in-differences framework that compares apps in the

Table 8: Impact of ATT on SDK Usage by Genre

	Monetization (1)	Ad Mediation (2)	Authentication (3)	Payments (4)
After ATT x Apple x Genre = books	-0.0188** (0.0092)	0.0265*** (0.0037)	0.0249*** (0.0039)	0.0330*** (0.0018)
After ATT x Apple x Genre = business	0.1684*** (0.0039)	0.0438*** (0.0017)	0.0494*** (0.0018)	0.0092*** (0.0011)
After ATT x Apple x Genre = education	0.1054*** (0.0045)	0.0340*** (0.0019)	0.0382*** (0.0019)	0.0185*** (0.0011)
After ATT x Apple x Genre = entertainment	-0.0758*** (0.0089)	0.0010 (0.0037)	-0.0310*** (0.0032)	0.0128*** (0.0014)
After ATT x Apple x Genre = finance	0.2736*** (0.0041)	0.0575*** (0.0018)	0.0428*** (0.0019)	0.0293*** (0.0010)
After ATT x Apple x Genre = food&drink	0.0653*** (0.0063)	0.0254*** (0.0024)	-0.0793*** (0.0039)	-0.0641*** (0.0037)
After ATT x Apple x Genre = games	-0.9545*** (0.0088)	-0.2725*** (0.0045)	-0.0437*** (0.0017)	0.0465*** (0.0008)
After ATT x Apple x Genre = graphic&design	-0.0098 (0.0395)	0.0252*** (0.0090)	0.0434*** (0.0143)	0.0166*** (0.0056)
After ATT x Apple x Genre = health&fitness	0.0766*** (0.0054)	0.0291*** (0.0027)	-0.0246*** (0.0027)	0.0082*** (0.0014)
After ATT x Apple x Genre = lifestyle	0.0783*** (0.0044)	0.0312*** (0.0020)	-0.0050** (0.0022)	-0.0130*** (0.0014)
After ATT x Apple x Genre = maps&navigation	-0.0341** (0.0138)	-0.0114* (0.0063)	-0.0015 (0.0062)	-0.0337*** (0.0045)
After ATT x Apple x Genre = medical	0.1569*** (0.0070)	0.0280*** (0.0026)	0.0577*** (0.0032)	-0.0005 (0.0022)
After ATT x Apple x Genre = music	-0.1086*** (0.0110)	-0.0126*** (0.0042)	0.0097** (0.0041)	0.0343*** (0.0021)
After ATT x Apple x Genre = news&magazine	-0.0086 (0.0066)	0.0216*** (0.0025)	0.0207*** (0.0033)	0.0338*** (0.0013)
After ATT x Apple x Genre = photo&video	-0.0566*** (0.0114)	0.0272*** (0.0054)	0.0044 (0.0043)	-0.0016 (0.0027)
After ATT x Apple x Genre = productivity	0.0844*** (0.0041)	0.0300*** (0.0018)	0.0647*** (0.0018)	0.0236*** (0.0010)
After ATT x Apple x Genre = shopping	0.1738*** (0.0061)	0.0410*** (0.0023)	-0.1015*** (0.0032)	-0.1236*** (0.0036)
After ATT x Apple x Genre = social	0.0171* (0.0094)	0.0362*** (0.0042)	-0.0399*** (0.0036)	0.0077*** (0.0019)
After ATT x Apple x Genre = sports	0.0006 (0.0073)	0.0128*** (0.0031)	-0.0235*** (0.0042)	0.0075*** (0.0021)
After ATT x Apple x Genre = travel	0.1097*** (0.0068)	0.0410*** (0.0026)	-0.0169*** (0.0032)	-0.0626*** (0.0028)
After ATT x Apple x Genre = weather	-0.1856*** (0.0134)	-0.0024 (0.0037)	0.0895*** (0.0053)	0.0285*** (0.0020)
App FE	✓	✓	✓	✓
Period FE	✓	✓	✓	✓
R ²	0.88900	0.86662	0.89495	0.90325
Observations	8,491,149	8,491,149	8,491,149	8,491,149

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Apple App Store with apps in the Google Play Store before and after the implementation of ATT. As Apple’s new framework was an exogenous shock that only affected developers in the Apple ecosystem, and that significantly affected the ability of apps to collect user data and share it with third parties for the purpose of advertising, this should be a good setting to estimate the causal effect of restricting the use of user data on the availability and characteristics of apps.

Our analysis suggests that, contrary to common concerns around the potential impact of similar policies, ATT did not have a negative long-run effect on the availability of apps, although we do observe a decrease in the number of updates to apps in the Apple App Store, which could be interpreted as developers losing interest in the platform. Instead, developers strategically adapted their efforts to the new conditions imposed by ATT. Examining the number of ratings received per month by existing apps vs new apps, we find evidence that ATT, and developers’ responses to it, have resulted in a decrease in the number of ratings received by existing apps and the score of such ratings. This may be related to developers’ decisions to update their apps less often. We find that new apps introduced after ATT seem to receive more ratings than apps introduced before, and don’t see the same decline in average ratings. Additionally, the number of new apps that receive no ratings in their first month has decreased since ATT.

We also examined how developers adapted their use of SDKs after ATT. Our result show that there is a decrease in the use of SDKs that rely on sharing of data with third parties for advertising, and an increase in the use of SDKs related to first-party data collection and monetization. We also observe an increase in the use of SDKs related to payment services, which suggests some apps may be looking for additional sources of revenue.

Overall, our results suggest that ATT has not significantly affected the availability of apps in the Apple ecosystem. It did, however, induce a response by developers who have adapted to the new framework. While the decrease in updates, and the apparent decrease valuation of users of existing apps, is an unwelcome outcome, the overall health of the platform seems strong.

References

- Alomar, N. and Egelman, S. (2022). Developers say the darnedest things: Privacy compliance processes followed by developers of child-directed apps.
- Bygrave, L. A. (2017). Data protection by design and by default : Deciphering the eu’s legislative requirements. *Oslo Law Review*, 4(2):105–120.
- DeGiulio, A., Lee, H., and Birrell, E. (2021). “ask app not to track”: The effect of opt-in tracking authorization on mobile privacy. In Saracino, A. and Mori, P., editors, *Emerging Technologies for Authorization and Authentication*, pages 152–167, Cham. Springer International Publishing.
- Ekambaranathan, A., Zhao, J., and Van Kleek, M. (2021). “money makes the world go around”: Identifying barriers to better privacy in children’s apps from developers’ perspectives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, New York, NY, USA. Association for Computing Machinery.
- Janssen, R., Kesler, R., Kummer, M. E., and Waldfogel, J. (2022). Gdpr and the lost generation of innovative apps. Technical report, National Bureau of Economic Research.
- Jones, C. (2022). The ios 14.5 update a game changer in federal privacy law.
- Kesler, R. (2022). The impact of apple’s app tracking transparency on app monetization. *Available at SSRN 4090786*.
- Kim, J.-H. and Wagman, L. (2021). The value of technology releases in the apple ios app.
- Kollnig, K., Binns, R., Van Kleek, M., Zhao, J., Lyngs, U., Tinsman, C., and Shadbolt, N. (2021). Before and after GDPR: Tracking in mobile apps. *Internet Policy Review: Journal on Internet Regulation*, 10(4):1–30.
- Kollnig, K., Shuba, A., Kleek, M. V., Binns, R., and Shadbolt, N. (2022). Goodbye tracking? impact of iOS app tracking transparency and privacy labels. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM.
- Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., and Acquisti, A. (2022). Does privacy regulation harm content providers? a longitudinal analysis of the impact of the gdpr. *Available at SSRN 4239013*.
- Leyden, B. T. (2023). Platform design and innovation incentives: Evidence from the product ratings system on apple’s app store. *CESifo Working Paper*.
- Li, D. and Tsai, H.-T. (2022). Mobile apps and targeted advertising: Competitive effects of data exchange. *Available at SSRN 4088166*.
- Mhaidli, A. H., Zou, Y., and Schaub, F. (2019). ”we can’t live without Them!” app developers’ adoption of ad networks and their considerations of consumer risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 225–244, Santa Clara, CA. USENIX Association.

- Momen, N., Hatamian, M., and Fritsch, L. (2019). Did app privacy improve after the gdpr? *IEEE Security & Privacy*, 17(06):10–20.
- PwC and IAB (2023). Internet advertising revenue report, fy 2022.
- Ribera, A. (2022). Trading off the orchard for an apple: the ios 14.5 privacy update. *Journal of European Competition Law & Practice*, 13.
- Sharma, P., Sun, Y., and Wagman, L. (2019). The differential effects of privacy protections and digital ad taxes on publisher and advertiser profitability. *Available at SSRN 3503065*.
- Sokol, D. D. and Zhu, F. (2021). Harming competition and consumers under the guise of protecting privacy: An analysis of apple’s ios 14 policy updates.

Appendix

Apple vs. the free internet

Apple plans to roll out a forced software update that will change the internet as we know it—for the worse.

Take your favorite cooking sites or sports blogs. Most are free because they show advertisements.

Apple's change will limit their ability to run personalized ads. To make ends meet, many will have to start charging you subscription fees or adding more in-app purchases, making the internet much more expensive and reducing high-quality free content.

Beyond hurting apps and websites, **many in the small business community say this change will be devastating for them too, at a time when they face enormous challenges.** They need to be able to effectively reach the people most interested in their products and services to grow.

Forty-four percent of small to medium businesses started or increased their usage of personalized ads on social media during the pandemic, according to a new Deloitte study. Without personalized ads, Facebook data shows that the **average small business advertiser stands to see a cut of over 60% in their sales for every dollar they spend.**

Small businesses deserve to be heard.
We're standing up to Apple for our small business customers and our communities.

Get the full story at fb.com/ApplePolicyUpdate

FACEBOOK
    

Figure 4: Facebook campaign against Apple's ATT

Table 9: Impact of ATT by Genre

	Update (1)	Log # New Ratings (2)	Avg. Rating (3)	No Rating (4)
After ATT x Apple x Genre = books	0.0003 (0.0005)	-0.0440*** (0.0008)	-0.1687 (0.1124)	0.0318*** (0.0006)
After ATT x Apple x Genre = business	-0.0090*** (0.0003)	-0.0260*** (0.0003)	-0.2857** (0.1110)	0.0239*** (0.0002)
After ATT x Apple x Genre = education	-0.0011*** (0.0003)	-0.0370*** (0.0004)	-0.2117* (0.1106)	0.0294*** (0.0003)
After ATT x Apple x Genre = entertainment	0.0007* (0.0004)	-0.0478*** (0.0007)	-0.2129* (0.1116)	0.0311*** (0.0004)
After ATT x Apple x Genre = finance	-0.0049*** (0.0007)	-0.0154*** (0.0009)	-0.1529 (0.1117)	0.0199*** (0.0005)
After ATT x Apple x Genre = food&drink	-0.0101*** (0.0004)	-0.0157*** (0.0005)	-0.1929* (0.1080)	0.0123*** (0.0003)
After ATT x Apple x Genre = games	0.0003 (0.0002)	-0.0855*** (0.0005)	-0.1768 (0.1103)	0.0425*** (0.0003)
After ATT x Apple x Genre = graphic&design	-0.0708*** (0.0054)	-0.0425*** (0.0100)	-0.3245 (0.2060)	0.0016 (0.0043)
After ATT x Apple x Genre = health&fitness	-0.0078*** (0.0005)	-0.0428*** (0.0006)	-0.2008* (0.1084)	0.0297*** (0.0004)
After ATT x Apple x Genre = lifestyle	0.0000 (0.0004)	-0.0330*** (0.0004)	-0.2263** (0.1105)	0.0265*** (0.0003)
After ATT x Apple x Genre = maps&navigation	0.0004 (0.0010)	-0.0104*** (0.0011)	-0.2802** (0.1312)	0.0157*** (0.0008)
After ATT x Apple x Genre = medical	-0.0084*** (0.0007)	-0.0273*** (0.0007)	-0.2009* (0.1112)	0.0241*** (0.0005)
After ATT x Apple x Genre = music	0.0034*** (0.0006)	-0.0475*** (0.0009)	-0.1937* (0.1135)	0.0348*** (0.0007)
After ATT x Apple x Genre = news&magazine	-0.0011 (0.0007)	-0.0486*** (0.0008)	-0.0894 (0.1263)	0.0245*** (0.0005)
After ATT x Apple x Genre = photo&video	-0.0044*** (0.0007)	-0.0761*** (0.0012)	-0.2805* (0.1473)	0.0420*** (0.0008)
After ATT x Apple x Genre = productivity	-0.0116*** (0.0003)	-0.0305*** (0.0003)	-0.2288** (0.1107)	0.0237*** (0.0002)
After ATT x Apple x Genre = shopping	-0.0266*** (0.0005)	-0.0509*** (0.0007)	-0.2193* (0.1135)	0.0317*** (0.0004)
After ATT x Apple x Genre = social	-0.0231*** (0.0007)	-0.0487*** (0.0010)	-0.2792** (0.1113)	0.0327*** (0.0006)
After ATT x Apple x Genre = sports	-0.0021*** (0.0007)	-0.0298*** (0.0009)	-0.2696** (0.1160)	0.0235*** (0.0006)
After ATT x Apple x Genre = travel	0.0011** (0.0005)	-0.0023*** (0.0006)	-0.2817** (0.1144)	0.0123*** (0.0004)
After ATT x Apple x Genre = weather	-0.0285*** (0.0020)	0.0514*** (0.0041)	-0.0993 (0.1188)	0.0105*** (0.0019)
App FE	✓	✓	✓	✓
Period FE	✓	✓	✓	✓
R ²	0.37524	0.79981	0.10083	0.59824
Observations	91,149,638	91,149,638	15,543,166	91,149,638

Robust standard errors in parenthesis (clustered by app and time).

*** $p < 0.01$, * $p < .05$, * $p < 0.1$