Did Apple's App Tracking Transparency Framework Harm the App Ecosystem?*

Cristobal Cheyre† Benjamin T. Leyden‡ Sagar Baviskar§
Alessandro Acquisti¶

June 2024
Click here for the latest version.

Abstract

We study the impact of Apple's App Tracking Transparency (ATT) framework on the Apple App Store ecosystem. ATT restricted app developers' access to personal identifiers used to target ads. Promoted as a privacy-enhancing initiative, the change was controversial: various stakeholders, including Meta, criticized ATT and predicted it would harm the app ecosystem. We collect data on every app available in both the Apple App Store and Google Play Store in the eighteen months around the implementation of ATT. We use a difference-in-differences analysis to comprehensively investigate whether the introduction of ATT negatively affected the app ecosystem. We consider multiple possible downstream effects, including changes in the likelihood that developers in the Apple ecosystem create new apps, update their existing apps, adapt app functionalities (such as the use of advertising platforms or payment systems), or withdraw from the market, as well as changes in the number of ratings and average ratings—as proxies for changes in the quality of apps. 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. The effect on ratings is nuanced. For existing apps, ATT leads to a minimal decline in the number of ratings received and average ratings. In contrast, the number of ratings for new entrants increases slightly after ATT. We also analyze Software Development Kits (SDK) data for a select number of apps and find a reduction in the use of Monetization and Ad Mediation SDKs, as well as an increase in the use of Authentication and Payments SDKs. Our results suggest that, contrary to pessimistic predictions about the impact of ATT on the app ecosystem, by and large, developers did not withdraw from the market after ATT and instead adapted to operate under the conditions of a more protective privacy framework.

*We thank Carol Lu, Yoonseock Son, and seminar and conference participants at the Research Roundtable on Regulating Privacy at George Mason University, WISE 2022, the 21^{st} IIOC, IMT-BS, and APIOC 2023 for their constructive feedback. Acquisti, Baviskar, and Cheyre gratefully acknowledge the financial support from the Program on Economics & Privacy at George Mason University. Acquisti and Cheyre acknowledge support from the National Science Foundation through award 2237327/2237328/2237329 (Understanding the Impact of Privacy Interventions on the Online Publishing Ecosystem).

†Cornell University. Email: cac555@cornell.edu

‡Cornell University; CESifo. Email: leyden@cornell.edu

§Carnegie Mellon University. Email: sbaviska@andrew.cmu.edu

¶Carnegie Mellon University. Email: acquisti@andrew.cmu.edu.

1 Introduction

Online user tracking—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 relies heavily on individual-level consumer data to serve targeted advertisements (PwC and IAB, 2023). While app developers report that online user tracking and advertising are necessary for monetizing apps (Mhaidli et al., 2019; Ekambaranathan et al., 2021; Ribera, 2022), and thus for the ecosystem to prosper, the prevalence of online tracking has led to widespread privacy concerns and has prompted both regulatory bodies and private companies to implement interventions aimed at protecting consumers' privacy.

One of the most prominent of these interventions is Apple's App Tracking Transparency (ATT) framework, introduced by Apple in April 2021. ATT significantly restricted developer access to a critical user identifier used to target ads. This change was controversial. Facebook mounted a national campaign denouncing the change as a "forced software update that will change the internet as we know it—for the worse."[1] Facebook's reasoning was that the change would make it more difficult to run personalized ads. This, in turn, would force providers of free ad-supported content (such as app developers) "to start charging you subscription fees or adding more in-app purchases, making the internet much more expensive and reducing high-quality free content" (Moon, 2020). Whether these predictions have materialized remains an open debate.

Concerns about the potential negative consequences of privacy interventions are not new. When the European General Data Protection Regulation (GDPR) was enacted, it was predicted it could have calamitous consequences for the availability of online content.[2] More recently, in response to the Federal Trade Commission's request for comments on their Advance Notice of Proposed Rulemaking to consider a new "Trade Regulation Rule on Commercial Surveillance and Data Security," advertising technology firms and advertising industry associations made similar claims (Federal Trade Commission, 2022).

Multiple academic studies have investigated whether privacy interventions affect the profitability and effectiveness of advertising. Goldberg et al. (2024) find that after GDPR, there

[1] See Figure 4 in Appendix A (Moon, 2020)
[2] See, for example, Downes (2018).

is a decline in recorded page views and revenues for EU websites. Aridor et al. (2024) finds that after ATT, advertising on Facebook becomes more expensive and less effective. Kraft et al. (2023) finds that ATT reduced the amount of traffic, making trackable ad impressions more expensive. Johnson et al. (2023a) and Kircher and Foerderer (2023) find that Youtube's 2020 settlement with the FTC over violations of the Children's Online Privacy Protection Act, which eliminated behaviorally targeted advertising from videos for kids, led to a decrease in new child-oriented content on the platform. When taken in isolation, these results suggest that privacy interventions have the potential for large adverse effects on ad-supported content providers and, thus, on consumers who benefit from them. However, it may be unwarranted to conclude that privacy interventions will inevitably hinder entire online ecosystems from such evidence—as current analyses do not focus on the capacity of ecosystems to evolve and adapt to more protective privacy frameworks. With some exceptions (Lefrere et al., 2022; Janssen et al., 2022), few studies to date have investigated the downstream impacts of privacy interventions on online ecosystems as a whole.

In this study, we use the implementation of Apple's App Tracking Transparency (ATT) framework to analyze how restricting the ability of apps to track users affects the ad-supported app ecosystem and, ultimately, its users. We consider a comprehensive host of multiple possible downstream effects, such as the quantity and quality of apps available in the marketplace after the enactment of ATT. We first examine how the new framework affected the incentives developers in the Apple ecosystem face to create new apps, update and invest in their existing apps, or withdraw from the market. We then analyze how developers adapted the monetization functionalities in their products, such as the use of advertising platforms or payment systems, by leveraging information on the presence of Software Development Kits (SDKs) in a select number of apps. Finally, we analyze the number of ratings received by apps as a proxy for app consumption and the average rating received by apps as a proxy for consumers' valuation. We distinguish between the effect for apps introduced before and after ATT to explore whether apps introduced under the new framework—which were conceivably created with different levels of development investments and monetization strategies—are evaluated differently by consumers.

Our analysis leverages information from a data provider that tracks the universe of apps on the Apple App Store and Google Play Store. This allows us to estimate the effect of ATT on the app ecosystem using a difference-in-differences strategy. We find that ATT leads to a temporary reduction in the entry of new apps into the Apple ecosystem, but the effect is short-lived, as it dissipates a few months after the policy, as the number of available apps in the Apple App Store ecosystem quickly recovers. Firm exit increases, but only very slightly, and some developers appear to reduce their effort slightly, as measured by the frequency of product updates. However, the magnitude of these effects, measured using common tools such as Cohen's d and equivalence testing (Schuirmann, 1987), is minimal. Next, we observe a slight reduction in data-intensive, targeted advertising SDKs and a slight increase in authentication and payment SDKs, indicative of a moderate shift toward alternative, non-invasive forms of monetization. Finally, we analyze consumers' consumption and valuation of apps. We find that pre-existing apps (those introduced before ATT) receive slightly fewer and lower ratings compared to before ATT. In contrast, new entrants (apps introduced after ATT) perform better than those introduced before ATT. On average, they are less likely to receive no ratings and receive a larger number of ratings. We find no significant effects in terms of average rating.

Overall, our results suggest that, contrary to the expectation of strong negative effects suggested by industry, the new framework produces minor and nuanced effects. ATT did not substantially reduce overall developer interest in the platform, even if firm investment in existing apps may have moderately decreased. Developers seem to have adapted to the new framework, in part by changing functionalities within their apps. These changes do not seem to have affected consumer consumption or valuation of apps. Our results highlight the importance of focusing not only on the effects of privacy interventions on "direct" metrics such as the effectiveness of behaviorally targeted advertising. To understand the effects of privacy interventions, it is worthwhile to explore how firms adapt to privacy interventions and to analyze the net downstream effects of those changes on consumers.

2 Related Literature

A growing body of literature explores the implications of privacy regulations and industry self-regulatory efforts. One large wave of contributions followed the introduction of GDPR in May 2018, and another emerged in response to the implementation of ATT. Below, we focus first on the stream of literature that has focused on the GDPR's effect on online tracking and its effect on the availability and quality of online content.[3] Next, we focus on ATT-related studies.

The introduction of GDPR led to changes in online tracking practices, particularly among EU organizations. Studies that focus on short-term responses find a sharp reduction in the use of third-party cookies (which are typically used to track and identify users across websites) and HTTP requests related to tracking (Peukert et al., 2022). The long-term evidence, however, is mixed. Johnson et al. (2023b) find that online tracking returns to the same level as before GDPR, while Lukic et al. (2023) find that it remains at a lower level relative to arguably unaffected (i.e., non-EU) websites. Lefrere et al. (2022) also find that tracking levels remain below pre-GDPR levels for EU-based online publishers.

Reductions in tracking may not affect tracking firms equally. The largest tracking firms may gain market power as a result of privacy protections such as GDPR. Peukert et al. (2022) find that the largest vendor (Google) suffered relatively lower losses than smaller vendors and thus increased its market power. Lefrere et al. (2022), too, examine changes to online publishers' use of tracking vendors before and after GDPR. In their work, they distinguish how websites react differently to EU and US visitors and find that websites choose different tracking partners depending on the visitor's region. For EU visitors, websites use significantly fewer trackers than for US visitors. Moreover, for US visitors, EU websites continue to use some of the trackers that they stopped using for EU users after GDPR. This suggests websites react to privacy interventions strategically.

Aridor et al. (2023) find evidence of compositional changes in tracking for those users that consent to be tracked. They show that, after GDPR, there is a decrease in the number

[3] In addition to work on the provision of ad-supported online content, researchers have also considered the effects of privacy regulations on the consumption of online content. See, for example, Congiu et al. (2022); Schmitt et al. (2021).

of total tracking cookies used but an increase in tracking for those who consent to be tracked. Similarly, Godinho de Matos and Adjerid (2022) find that after GDPR, some users declined to be tracked. Still, the inability to keep tracking those users may be to some extent offset by other users that were more amenable to marketing and allowed the use of additional data types for those purposes.

Researchers have also studied the effect of GDPR on tracking practices in the context of mobile apps. Kollnig et al. (2021) find there have been limited changes in the presence of third-party tracking in apps and that the concentration of tracking capabilities among a few large gatekeeper companies persists. However, this does not mean privacy and data protection have not improved after GDPR. Momen et al. (2019) analyze app behavior before and after GDPR and show that user privacy has moderately improved after the implementation of the regulation. Finally, Warberg et al. (2023) examine privacy dialogs displayed by websites after GDPR, finding that over time there is an increase in the adoption of consent mechanisms (i.e., dialog boxes providing users with privacy choices), and within those dialog boxes, the option to decline tracking has become more explicit.

Closer to our focus are studies that have analyzed the availability and quality of online content and mobile apps. Lefrere et al. (2022) study EU and US online publishers and find that, although EU publishers implemented changes following GDPR (relative to US publishers), there is no evidence that the regulation inhibited their ability to produce content or generate user engagement. Moreover, they find that almost none of the publishers in their sample exited the market after GDPR. In contrast, 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, although changes in the Google Play Store that were not related to GDPR may have affected these dynamics Kollnig and Binns (2022).

Since the enforcement of ATT, researchers have been interested in understanding the effectiveness of ATT in reducing tracking and its impact on the app ecosystem and its users. The issues studied to some extent overlap with those previously analyzed for GDPR and reveal some similarities and differences. As we discuss in Section 3, this is not surprising, as GDPR and ATT have significant differences in terms of requirements, reach, and implementation.

Some studies have analyzed the effectiveness of the framework in effectively reducing tracking. Kollnig et al. (2022b) 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. However, as that data brokers are facing higher challenges in tracking users, 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). Kollnig et al. (2022a) compare privacy features between Apple and Android mobile devices and find that although there are privacy violations in both platforms, there is less advertising tracking in iOS devices. 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. The reduction in consumer tracking on iOS devices has resulted in some benefits to its users. Bian et al. (2024) find that after ATT, the rate of financial fraud decreases in regions with greater iOS penetration, which they attribute to ATT's tracking protection. Finally, research has shown that the limits that ATT imposes on tracking have affected advertising effectiveness. After ATT, publishers received lower revenues per impression shown to Apple users (Kraft et al., 2023). Aridor et al. (2024) show that targeting ads on Facebook based on off-platform data became more difficult after ATT, making it harder for merchants to acquire new clients.

Another set of contributions has focused on the effects of ATT on competition in the advertising market. Several studies have argued that ATT effectively shifts market power towards Apple, as while it blocks third-party tracking, it does not block Apple's ability to track users. Sokol and Zhu (2021) argue that ATT represents an anti-competitive strategy that is harmful not only for a fair competitive market but potentially for end users as well. Its features seem designed to reinforce Apple's infrastructural control over the platform (Woodward, 2023). Hoppner and Westerhoff (2021) have similar concerns and argue that the policy is limiting data-based competition in the advertising industry, and it may violate EU competition regulations. Empirically, Deisenroth et al. (2024) use data on Facebook advertisers to show that offline industries that were more affected by ATT saw a decrease in the number of firms and an increase in prices relative to less-affected industries.

There is some evidence that ATT has affected how apps are monetized. 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. In the case of gaming apps, Le Meur (2023) finds that after ATT, game app developers were able to adapt without moving completely away from the freemium model, although adapting was easier for larger developers that had access to more first party data. 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.

Our work is also related to research that explored the use of Software Developer Kits (SDKs) in 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. Jin et al. (2024) study how the introduction of GDPR affects Android app developers' use of SDKs. They find that relative to apps only available in the US, apps listed in the EU (either those only listed in the EU or those listed in both the EU and the US) reduce their use of SDKs. However, this effect is nuanced. EU-only developers reduce their reliance on SDKs from major providers that collect sensitive information while increasing their use of SDKs from major providers that collect non-sensitive data.

The key contribution of our study relative to related research consists in its focus on studying the impact of ATT on the complete universe of free apps, analyzing how developers adapted their products following the implementation of the framework, and evaluating the consequences of these changes for users. Unlike prior work, our results suggest that the expectation of large, welfare destroying consequences of privacy interventions are not necessarily warranted. In the case of ATT, its effect on developers and, ultimately, consumers, seems to have been more nuanced than ordinarily predicted, as developers adapted to operate under the new privacy framework.

3 Institutional Background and Theoretical Framework

The mobile app ecosystem has flourished under the ad-supported model, in which many apps are offered at no cost to consumers and are monetized through advertising and the sale of user data to third parties (Mhaidli et al., 2019; Ekambaranathan et al., 2021; Ribera, 2022). Thus, restricting the ability of mobile apps to access and monetize this data has the potential to significantly shift the incentives developers face for maintaining and developing new apps or may lead them to adopt different monetization strategies. If privacy interventions lead to a significant reduction in existing apps and depress the entry of new apps, they can negatively affect consumer surplus (Janssen et al., 2022). The introduction of ATT provides us with an expedient scenario to explore these questions.

In June 2020, Apple announced that it was working on a new consumer privacy framework to be introduced in a future version of its iOS mobile operating system. The App Tracking Transparency (ATT) framework changed the privacy and data collection policy that governs how Apple 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 a two-stage, opt-in basis. Despite strong opposition from technology firms that derive most of their revenue from advertising, Apple announced on April 20, 2021 that ATT would be released as part of iOS 14.5 the following week (Moon, 2020).

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. Under this original regime, consumers had the ability to opt out of sharing this identifier, but few did. ATT changed this opt-*out* model by implementing a two-stage opt-*in* process, under which consumers had to first opt into allowing developers to ask for permission to access the IDFA (see Figure 1a), and then, conditional on that allowance being granted, each developer would individually have to ask for and be granted permission to access the identifier by the user (see Figure 1b). As a result, app developers and data brokers faced more significant challenges to track users at the same level as they used to. Reports indicate that the number of users that have chosen not to opt-in could range from 75% to as high as 95% (Lukovitz, 2022; Axon, 2021). This sharp reduction in the number of users that can be tracked has the potential to significantly affect apps that are monetized through advertising.

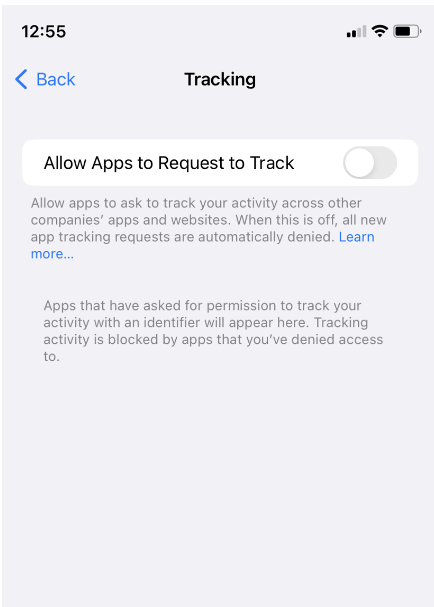

(a) Privacy setting to prevent apps from requesting authorization to track

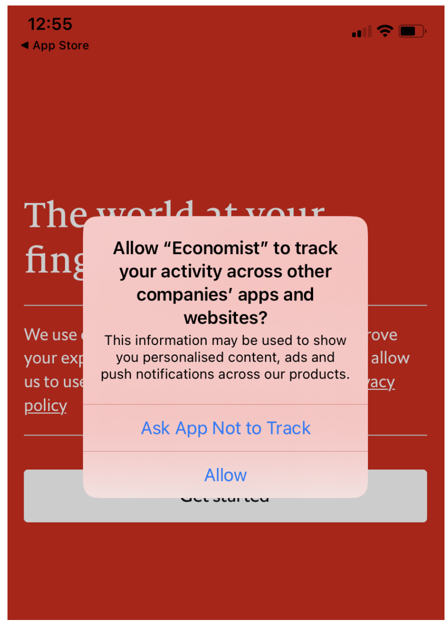

(b) App-tracking authorization request

Figure 1: Apple's Tracking Authorization Settings

It is worth comparing ATT and GDPR, as a series of prior studies have leveraged the implementation of GDPR to analyze the impact of privacy interventions. ATT presents some advantages over GDPR for the purpose of studying the impact of privacy interventions. 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 presents a relatively cleaner setting. Developers cannot avoid the requirements of ATT, and while the framework was first announced in June 2020 at Apple's Worldwide Developers Conference, developers only learned the date it would become effective with a week's notice through a blog post by Apple (Apple Developer, 2021). Additionally, as GDPR applies to all entities processing data from EU citizens, it is difficult to find a suitable control group of fully untreated 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. While the Google ecosystem has not yet been affected by a similar policy, there could still exist other confounding factors as both platforms are periodically introducing changes that may affect the app ecosystem. Thus, comparisons between the two ecosystems should be limited to relatively short periods of time around the event, and the potential of concurrent changes must be considered carefully.

A key difference between ATT and GDPR is the way their requirements are implemented. While in the case of GDPR, each entity must interpret the regulation and decide whether to implement changes independently, in the case of ATT, the requirements are enforced through Apple's iOS operating system by restricting, from the platform's side, access to the IDFA. Moreover, after an update, the new operating system diffuses quickly across devices. In the case of iOS 14.5, according to data included in Kraft et al. (2023), approximately 80% of devices had updated to the new operating system by the end of June 2021.

The changes introduced by ATT may significantly reduce the amount of user data (and the services they connect to) apps can leverage for advertising. In fact, prior work has found ATT was effective in preventing the collection of the IDFA cross-app tracking identifier Kollnig et al. (2022b) and that its opt-in requirement was robust to most developers' implementation of tracking requests (DeGiulio et al., 2021). This, 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). 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 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 consented to tracking, if enough users in the ecosystem were blocking tracking in general, the implementation of ATT would cause 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.

Additionally, ATT may affect the ability of developers to reach potential users. Targeted advertising, either in other apps or any other online setting, and search advertising are key conduits developers use to make their apps known and acquire new users (Li and Tsai, 2022). 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 is the case, apps may obtain fewer users than before ATT.

The two mechanisms outlined above may, together or independently, affect the revenues of developers of free, ad-supported apps. They may start earning lower revenues from advertising as users start to decline requests to track, and non-targeted ads become more prevalent. They may also see their revenues dwindle if they face difficulties in obtaining new users, and thus, their ad impressions inventory suffers. As a result of these changes, developers may increase the amount of advertising they include in their apps, try to replace lost advertising revenue by increasing their reliance on in-app purchases, completely switch to a paid model, or simply abandon the ecosystem.

Alternatively, developers may find ways to adapt to the regulation. Furthermore, revenues may not be significantly affected if advertisers are interested in reaching iOS users even in the absence of detailed third-party data. Therefore, it is not immediately evident whether ATT will affect app availability.

Determining the impacts ATT may have on end-users is also not straightforward. If ATT were to decrease the availability of new apps or the continued investment of developers with existing apps, consumers may suffer harm from the reduction of variety and quality. However, not all apps provide the same surplus to users, so a reduction in variety might not affect consumer surplus in a meaningful way if consumers find adequate substitutes. 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 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 highly 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 would be more costly to consumers than a corresponding increase in exit. However, if developers did 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 effects 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.

While we cannot directly measure demand and consumer surplus, we rely on the number of customer ratings and average ratings received by apps as proxies of demand and quality. Thus, to study whether ATT has affected the quality or valuation consumers have of apps, we examine whether ATT has affected the number of ratings received by apps over time and the average rating they receive.

4 Data

We use data from AppMonsta, which collects data on the universe of apps in the Apple App Store and the Google Play Store. This data source has previously been used to study mobile app marketplaces by Ershov (2022, 2023), Janssen et al. (2022), Kircher and Foerderer (2024), Leyden (2024), and Deng et al. (2023), among others. For each app, we observe the app's name, genre (i.e., product category), developer, entry and exit dates, the date of any product updates, and the number of reviews and average ratings it received over time.

For a smaller number of apps, AppMonsta also provides information on the software developer kits (SDKs) each app uses.[4] 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 like Google and Facebook provide

[4]This data was sold under the product name MightySignal.

Table 1: Summary Statistics

	App Store		Google Play Store		
	Mean	SD	Mean	SD	N
Log Entry Count	6.21	1.42	7.96	1.08	894
Log Exit Count	6.66	1.33	8.03	1.05	894
Update	0.09	0.29	0.07	0.26	87,036,182
Log New Ratings	0.15	0.71	0.22	0.72	78,518,418
Avg. New Rating	4.07	1.25	3.89	1.30	9,694,838
Monetization	1.14	2.00	1.23	1.55	8,448,364
Ad Mediation	0.16	0.88	0.26	1.47	8,448,364
Authentication	0.47	0.75	1.14	0.90	8,448,364
Payments	0.12	0.46	0.58	0.74	8,448,364

SDKs to allow developers to easily include advertising in their apps. To collect data on the SDKs used by developers, AppMonsta periodically downloaded a number of apps from both ecosystems and analyzed the SDKs included in them. The SDKs are classified based on their functionality. We focus on four categories of SDKs whose usage is likely to be affected by the introduction of ATT: monetization, ad mediation, payments, and authentication.

We restrict our analysis to "free" apps, i.e., those that set an upfront price of $0 as they are more likely to be ad-supported and thus affected by ATT. 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. During our sample period, entry and exit were somewhat higher on Google Play than on the App Store. The platforms are similar in terms of updating frequency and customer ratings. Regarding SDK usage by developers, we see that Monetization SDKs are by far the most used, followed by Authentication SDKs.

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 been affected by different events over time, the most notable of them being the implementation of GDPR, which corresponds 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 extraneous

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

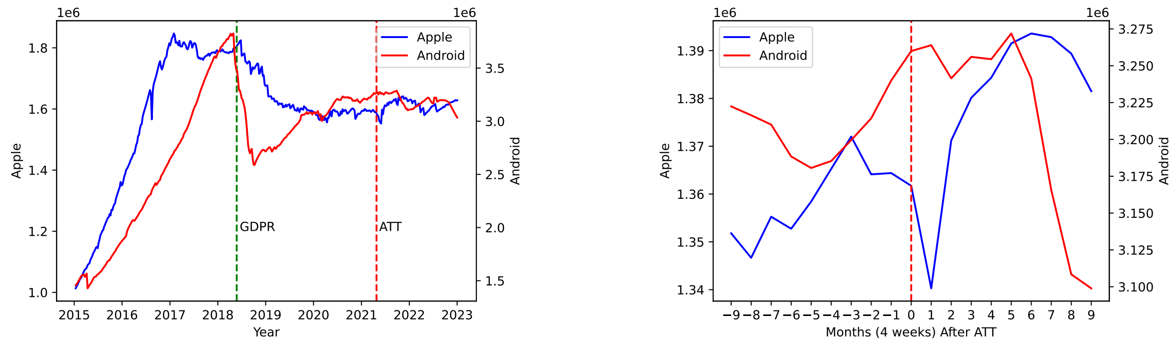

(a) Active Apps in Each Ecosystem Over Time

(b) Active Apps in Each Ecosystem From 6 Months Before ATT to 6 Months After ATT

confounding effects, we focus our analysis on the 18-month period around the implementation of ATT. Figure 2b shows the number of active apps in each ecosystem from 9 months before to 9 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 soon followed by a recovery in the number of active apps.

5 Empirical Analysis

We use a difference-in-differences (DiD) framework to study the effect of ATT on outcomes related to the state of Apple's App Store ecosystem, including the entry, exit, and update frequency of apps, the impact of the framework on consumer valuations of apps, such as the number of new ratings they receive and changes to their average ratings, and how it influenced the monetization functionalities included by developers in their apps through the use of SDKs. As we explained in Section 3, although Apple had announced they would be implementing the new framework in a future operating system update, developers only had one week's notice that the change was taking place with the release of iOS 14.5. Thus, examining the evolution of the Apple App Store, which was affected by ATT, versus the Google Play Store, which was not 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.

5.1 Market Dynamics: Entry, Exit, and Product Updates

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.[5] 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

$$Y_{p,g,t} = \alpha_1 PostATT_t \times Apple_p + \delta_{p,g} + \mu_t + \epsilon_{p,g,t}, \tag{1}$$

where p indicates a platform, g indicates a genre (i.e., product category), and t indicates a month.[6] 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, $\delta_{p,g}$ and μ_t are platform-genre and month fixed effects.

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

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

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.

We present the results of estimating Equations (1) and (2) on entry, exit, and product updates in Table 2. In column (1) of Table 2, we see that entry declined slightly on the App Store relative to the Google Play Store following the introduction of the ATT policy. A possible dynamic consistent with these results is that some developers' abilities to collect

[5]The analysis presented here considers only those apps that charge an upfront price of $0 throughout the entire sample period because ATT is most relevant to the business model of these apps. In Appendix B, we find that the results are not sensitive to the inclusion of paid apps.

[6]Months are defined as four-week intervals relative to the release of ATT.

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

	Log Entry Count (1)	Log Exit Count (2)	Update (3)
After ATT x Apple	-0.1634***	0.0651*	-0.0041***
	(0.0611)	(0.0333)	(0.0001)
Platform-Genre FE	✓	✓	
Period FE	✓	✓	✓
App FE			✓
Cohen's D	-0.3587	0.2732	-0.0186
Dependent variable mean	7.0268	7.2847	0.07768
R^2	0.91637	0.97163	0.37103
Observations	817	817	87,036,182

Observation in columns (1) and (2) are at the platform-genre-month level. Observations in column (3) are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

revenue through targeted advertising and other consumer tracking technologies may have been hampered by the new policy. However, panel (a) of Figure 3 shows that the finding is entirely driven by an initial decline in entry over the first two months of the policy, after which entry rates return to normal, and the number of active apps in the Apple ecosystem soon surpasses the pre-ATT levels (Figure Figure 2b).[7] In column (2) of Table 2, we find a positive coefficient on exit behavior on the App Store following ATT-but the effect is noisy ($p < 0.1$). Two common tools for interpreting the size of coefficients—and thus the magnitude of the effects—are Cohen's d and equivalence testing Schuirmann (1987). For all regressions, we present Cohen's d results in the tables and equivalence testing results in the Appendix C. Cohen's d for the Exit regression is just above 0.2 (the threshold for an effect to be considered "small"). The equivalence test classifies this effect as a "trivial difference"

[7]It is worth noting that some of the estimated coefficients in the pre-ATT periods are statistically significant. To the extent that there may be minor violations of the parallel trends assumption that persist into post-ATT periods, those would only further *weaken* the evidence that ATT negatively affected developer participation on the platform (see Rambachan and Roth (2023) for a discussion of these issues). Given our assertion that these effects are either short-lived or not economically significant, we view our estimates as conservative. That is, the negative impacts of ATT might be weaker than the already small estimates we present.

(following Tryon and Lewis (2008)), as the coefficients are statistically different from zero based on the t-test but statistically equivalent according to the equivalence test.

Figure 3: Entry and Exit Over Time

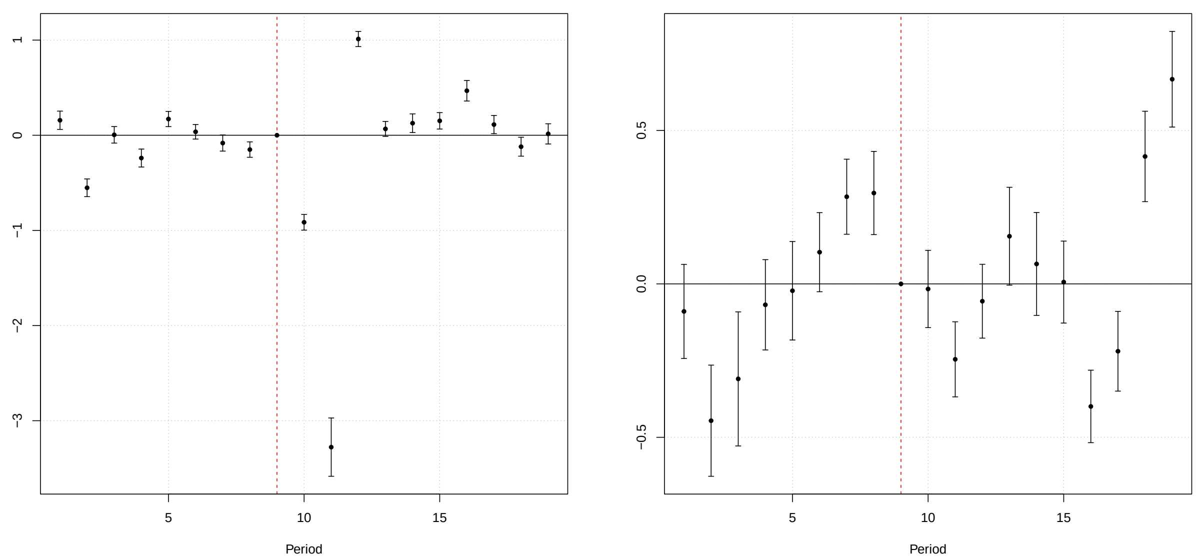

(a) Entry

(b) Exit

While some modest exit does occur on the platform, as evidenced both in Table 1 and column (2) of Table 2, this may understate the true impact of ATT on developer exit because 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 is just a one-time \$25 fee for the Google Play Store regardless of how many apps the developer has on the platform). Given this, looking at the frequency of product updates 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 (2024) 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 (5.3%) in response to ATT.[8] To get a better sense of the effect size for all models, we also calculate Cohen's d effect sizes (Cohen, 2013), which are also presented in Table 2. The Cohen's d of the update coefficient is well below the small threshold at 0.0186, and the equivalence test classifies this effect as a trivial difference (see Appendix C). This suggests that the change in updating frequency in the Apple ecosystem relative to the Android ecosystem, following ATT, was *practically* zero, even though statistically significantly different from zero.

While updating declines by a small amount on average, we find that this effect varies in both magnitude and sign across genre, as documented in column (1) of Table 3 where we re-estimate Equation (2) while allowing β_1 to vary by genre. Most notably, for games apps, which is one of the most sizable genres in terms of the number of apps and which are typically monetized through advertising, we observe that update frequency increases after ATT. This may be driven by developers adjusting their monetization strategies.[9]

5.2 Developers' Use of Software Developer Kits

Next, we consider whether ATT had any effect on how developers construct and monetize their apps, as this can indicate whether developers adjusted their business model in response to the initiative. 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 (used to let users log in using credentials from different platforms), and Payment SDKs (used for securely processing card payments).[10]

[8] During each month, 7.7% of apps get an update. A 5.3% reduction on the 7.7% of apps that get updated per month corresponds to roughly 6,000 apps reducing their updates out of approximately 1.5 million apps in the store.

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

[10] Within 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 3: Impact of ATT by Genre

	Update (1)	Log # New Ratings (2)	No Rating (3)	Avg. Rating (4)
After ATT x Apple × Genre = books	0.0020***	-0.0187***	0.0111***	-0.0220***
	(0.0005)	(0.0006)	(0.0005)	(0.0076)
After ATT x Apple × Genre = business	-0.0077***	-0.0057***	0.0074***	-0.0276***
	(0.0003)	(0.0002)	(0.0002)	(0.0089)
After ATT x Apple × Genre = education	0.0009***	-0.0149***	0.0114***	-0.0199***
	(0.0003)	(0.0003)	(0.0002)	(0.0060)
After ATT x Apple × Genre = entertainment	0.0016***	-0.0222***	0.0130***	-0.0337***
	(0.0004)	(0.0005)	(0.0004)	(0.0072)
After ATT x Apple × Genre = finance	-0.0036***	0.0035***	0.0049***	0.0471***
	(0.0007)	(0.0008)	(0.0004)	(0.0056)
After ATT x Apple × Genre = food&drink	-0.0101***	0.0073***	-0.0047***	0.0031
	(0.0004)	(0.0004)	(0.0003)	(0.0098)
After ATT x Apple × Genre = games	0.0025***	-0.0530***	0.0201***	-0.0091***
	(0.0002)	(0.0003)	(0.0002)	(0.0030)
After ATT x Apple × Genre = graphic&design	-0.0716***	-0.0637***	-0.0006	-0.0078
	(0.0055)	(0.0087)	(0.0035)	(0.0187)
After ATT x Apple × Genre = health&fitness	-0.0069***	-0.0211***	0.0117***	-0.0160**
	(0.0005)	(0.0005)	(0.0003)	(0.0066)
After ATT x Apple × Genre = lifestyle	0.0032***	-0.0105***	0.0096***	-0.0100
	(0.0004)	(0.0003)	(0.0002)	(0.0068)
After ATT x Apple × Genre = maps&navigation	0.0001	0.0078***	0.0020***	-0.0332**
	(0.0010)	(0.0009)	(0.0006)	(0.0169)
After ATT x Apple × Genre = medical	-0.0077***	-0.0087***	0.0079***	-0.0123
	(0.0007)	(0.0006)	(0.0004)	(0.0113)
After ATT x Apple × Genre = music	0.0034***	-0.0203***	0.0145***	-0.0061
	(0.0006)	(0.0007)	(0.0006)	(0.0092)
After ATT x Apple × Genre = news&magazine	-0.0004	-0.0274***	0.0112***	0.0262**
	(0.0007)	(0.0007)	(0.0005)	(0.0109)
After ATT x Apple × Genre = photo&video	-0.0032***	-0.0483***	0.0210***	0.0020
	(0.0007)	(0.0010)	(0.0006)	(0.0079)
After ATT x Apple × Genre = productivity	-0.0098***	-0.0105***	0.0084***	-0.0055
	(0.0003)	(0.0003)	(0.0002)	(0.0052)
After ATT x Apple × Genre = shopping	-0.0262***	-0.0346***	0.0176***	-0.0343***
	(0.0006)	(0.0006)	(0.0003)	(0.0083)
After ATT x Apple × Genre = social	-0.0222***	-0.0237***	0.0148***	-0.0848***
	(0.0008)	(0.0008)	(0.0005)	(0.0114)
After ATT x Apple × Genre = sports	-0.0020***	-0.0091***	0.0087***	-0.0438***
	(0.0007)	(0.0008)	(0.0005)	(0.0117)
After ATT x Apple × Genre = travel	0.0016***	0.0133***	0.0003	-0.0344***
	(0.0005)	(0.0005)	(0.0003)	(0.0109)
After ATT x Apple × Genre = weather	-0.0302***	0.0767***	-0.0060***	0.0790***
	(0.0021)	(0.0042)	(0.0017)	(0.0193)
App FE	✓	✓	✓	✓
Period FE	✓	✓	✓	✓
Dependent variable mean	0.07768	0.19973	0.87653	3.9250
R^2	0.37108	0.80223	0.58457	0.43476
Observations	87,036,182	78,518,418	78,518,418	9,694,838

Observations are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Table 4: 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 4. 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 share 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 (of merchandise or other offline goods), we expect to see an increase in the use of Payments SDKs.

In Table 6, we present the results of estimating Equation (2) for each of these four SDK categories. Monetization and Ad Mediation SDKs decline slightly in usage, while there's a slight 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 5). These results suggest an effort by some 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 a small increase in the use of Authentication SDKs. Ex ante, it is not 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. 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 some 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.

Table 5: Average SDK Usage by Genre

	App Store				Google Play Store				# Apps
	Monetization	Ad Mediation	Authentication	Payments	Monetization	Ad Mediation	Authentication	Payments	
Books	1.08	0.13	0.46	0.07	1.23	0.13	0.81	0.29	19457
Business	0.79	0.07	0.39	0.11	0.79	0.15	1.22	0.67	36089
Education	0.92	0.10	0.44	0.10	0.93	0.12	1.09	0.57	43397
Entertainment	1.20	0.16	0.47	0.10	1.28	0.24	1.02	0.49	21887
Finance	0.70	0.08	0.39	0.07	0.73	0.05	1.32	0.57	27953
Food & Drink	0.85	0.06	0.59	0.29	1.02	0.04	1.30	1.07	15638
Games	2.72	0.60	0.52	0.06	2.81	1.06	1.07	0.37	50492
Graphic & Design	1.28	0.05	0.54	0.06	1.25	0.16	0.72	0.32	2374
Health & Fitness	0.97	0.11	0.50	0.12	1.10	0.10	1.26	0.72	20546
Lifestyle	0.94	0.09	0.51	0.14	0.99	0.28	1.29	0.74	37166
Maps & Navigation	1.00	0.11	0.46	0.15	0.93	0.10	1.30	0.68	7866
Medical	0.80	0.07	0.36	0.13	0.77	0.09	1.07	0.59	9508
Music	1.17	0.15	0.45	0.08	1.38	0.29	0.88	0.42	20364
News & Magazine	1.12	0.10	0.41	0.07	1.32	0.22	1.28	0.79	14197
Photo & Video	1.23	0.18	0.44	0.12	1.31	0.28	0.81	0.48	6963
Productivity	0.95	0.11	0.38	0.10	1.03	0.18	0.95	0.44	57489
Shopping	0.90	0.09	0.59	0.29	0.96	0.06	1.48	0.94	18452
Social	1.18	0.18	0.58	0.11	1.19	0.32	1.32	0.66	10239
Sports	1.01	0.10	0.59	0.09	1.40	0.44	1.40	0.71	13943
Travel	1.00	0.08	0.53	0.23	1.01	0.12	1.32	0.70	17394
Weather	1.28	0.09	0.34	0.06	2.17	0.35	1.11	0.50	3485

Table 6: Impact of ATT on SDK Usage

	Monetization (1)	Ad Mediation (2)	Authentication (3)	Payments (4)
After ATT x Apple	-0.0258***	-0.0020**	0.0052***	0.0062***
	(0.0017)	(0.0008)	(0.0007)	(0.0004)
App FE	✓	✓	✓	✓
Period FE	✓	✓	✓	✓
Cohen's D	-0.0460	-0.0038	0.0171	0.0269
Dependent variable mean	1.2140	0.24741	1.0353	0.51105
R^2	0.88746	0.86648	0.89495	0.90325
Observations	8,448,364	8,448,364	8,448,364	8,448,364

Observations are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

While we find statistically significant effects, the coefficient estimates Cohen's d indicates that all effect sizes are below the small threshold and are classifiable as trivial differences under equivalence testing (Appendix C). It is thus important 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 to identify apps more likely to be affected by the framework. In Table 5, we present the frequency of each SDK type by genre. In Table 7, we present the results of estimating Equation (2) while allowing the effect to vary by genre. We see that (similar to our earlier discussion of Table 3) responses vary in both magnitude and 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 less of both. The results in Table 7 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

Table 7: Impact of ATT on SDK Usage by Genre

	Monetization (1)	Ad Mediation (2)	Authentication (3)	Payments (4)
After ATT x Apple × Genre = books	-0.0180*	0.0269***	0.0245***	0.0338***
	(0.0092)	(0.0038)	(0.0039)	(0.0019)
After ATT x Apple × Genre = business	0.1688***	0.0441***	0.0488***	0.0093***
	(0.0039)	(0.0017)	(0.0018)	(0.0011)
After ATT x Apple × Genre = education	0.1022***	0.0328***	0.0365***	0.0195***
	(0.0045)	(0.0019)	(0.0019)	(0.0011)
After ATT x Apple × Genre = entertainment	-0.0795***	0.0010	-0.0310***	0.0126***
	(0.0090)	(0.0038)	(0.0032)	(0.0015)
After ATT x Apple × Genre = finance	0.2730***	0.0581***	0.0435***	0.0294***
	(0.0041)	(0.0018)	(0.0019)	(0.0010)
After ATT x Apple × Genre = food&drink	0.0669***	0.0261***	-0.0784***	-0.0640***
	(0.0064)	(0.0025)	(0.0039)	(0.0037)
After ATT x Apple × Genre = games	-0.9633***	-0.2749***	-0.0458***	0.0466***
	(0.0089)	(0.0045)	(0.0017)	(0.0008)
After ATT x Apple × Genre = graphic&design	-0.0084	0.0253***	0.0438***	0.0166***
	(0.0396)	(0.0090)	(0.0143)	(0.0056)
After ATT x Apple × Genre = health&fitness	0.0741***	0.0291***	-0.0254***	0.0084***
	(0.0055)	(0.0027)	(0.0027)	(0.0015)
After ATT x Apple × Genre = lifestyle	0.0802***	0.0316***	-0.0043**	-0.0128***
	(0.0045)	(0.0020)	(0.0022)	(0.0014)
After ATT x Apple × Genre = maps&navigation	-0.0344**	-0.0115*	-0.0018	-0.0324***
	(0.0140)	(0.0064)	(0.0063)	(0.0045)
After ATT x Apple × Genre = medical	0.1561***	0.0268***	0.0571***	-0.0006
	(0.0071)	(0.0025)	(0.0032)	(0.0022)
After ATT x Apple × Genre = music	-0.1109***	-0.0142***	0.0102**	0.0350***
	(0.0111)	(0.0042)	(0.0041)	(0.0021)
After ATT x Apple × Genre = news&magazine	-0.0131**	0.0206***	0.0212***	0.0339***
	(0.0066)	(0.0025)	(0.0033)	(0.0013)
After ATT x Apple × Genre = photo&video	-0.0499***	0.0280***	0.0056	-0.0016
	(0.0113)	(0.0054)	(0.0043)	(0.0027)
After ATT x Apple × Genre = productivity	0.0843***	0.0311***	0.0641***	0.0237***
	(0.0041)	(0.0018)	(0.0018)	(0.0010)
After ATT x Apple × Genre = shopping	0.1782***	0.0428***	-0.1008***	-0.1240***
	(0.0062)	(0.0023)	(0.0032)	(0.0037)
After ATT x Apple × Genre = social	0.0188**	0.0373***	-0.0426***	0.0081***
	(0.0094)	(0.0042)	(0.0036)	(0.0020)
After ATT x Apple × Genre = sports	0.0020	0.0145***	-0.0230***	0.0064***
	(0.0073)	(0.0032)	(0.0043)	(0.0021)
After ATT x Apple × Genre = travel	0.1123***	0.0413***	-0.0160***	-0.0616***
	(0.0069)	(0.0026)	(0.0032)	(0.0028)
After ATT x Apple × Genre = weather	-0.1816***	-0.0013	0.0926***	0.0297***
	(0.0135)	(0.0038)	(0.0053)	(0.0020)
App FE	✓	✓	✓	✓
Period FE	✓	✓	✓	✓
Dependent variable mean	1.2140	0.24741	1.0353	0.51105
R^2	0.88917	0.86666	0.89505	0.90336
Observations	8,448,364	8,448,364	8,448,364	8,448,364

Observations are at the app-month level. Months are defined as four-week intervals relative to the release of ATT. Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Table 8: Effect of ATT on the Number of Ratings and Average Rating

	Log # New Ratings (1)	No Rating (2)	Avg. Rating (3)
After ATT x Apple	-0.0176***	0.0105***	-0.0101***
	(0.0002)	(0.0001)	(0.0019)
App FE	✓	✓	✓
Period FE	✓	✓	✓
Cohen's D	-0.0525	0.0476	-0.0092
Dependent variable mean	0.19973	0.87653	3.9250
R^2	0.80218	0.58455	0.43475
Observations	78,518,418	78,518,418	9,694,838

Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

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.

5.3 Consumer Valuations

We next turn to analyzing whether users' valuations of apps changed after ATT. Besides affecting incentives to update existing apps, ATT may reduce the effort developers devote to new apps, thus resulting in lower quality. While we cannot measure demand and consumer surplus directly, we can use ratings as a proxy. This is informative of consumer welfare effects, as ultimately, what matters is whether consumers can still find and enjoy apps they want and need.

Theoretically, the relationship between ATT and consumers' demand and valuation of apps is difficult to anticipate ex-ante. 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 in users' valuations and, thus, in demand and ratings. Additionally, if consumers value Apple's privacy framework, additional users may join the ecosystem, which in turn would result in more downloads for Apps. On the other side, if ATT leads developers to stop updating their products, invest less in the development of new apps, or increase the number of ads (and/or quality of ads decreases) displayed to users to compensate for lower per-impression revenue, then the quality of apps may be affected.

We find only very modest but statistically significant effects in Table 8. We observe a 1.8% decline in the number of ratings received per month and an approximately 1% increase in the number of apps that receive no reviews in a given month. While there is a decline in the average rating received by apps, it is not economically meaningful, being a 0.01 decline on a scale of 1 to 5 stars.[11]

What Table 9 cannot tell us is whether all apps are equally affected by ATT. In particular, it may be the case that apps introduced after ATT took the framework into consideration from conception, while incumbent apps, which were originally designed with a less restrictive privacy policy in mind, may have been more negatively impacted by the change. To better understand this, we next study whether the number of ratings and average ratings received by apps that enter before vs. after ATT differ. To do so, we consider how this policy affected app ratings in the first month an app was available for sale. Specifically, we compute the number of ratings each app received during its first full month listed, the average rating they received in that same period, and whether they received no ratings during the first month. Note that since each app is only observed for its first month of listing, we cannot use app fixed effects. Instead, we use platform-genre fixed effects, as in Equation (1).

We present the results of estimating this augmented version of Equation (2) with this sample in Table 9. The results show that apps released *after* ATT received more ratings in the first month than those released before (roughly a 6.4% increase), and the probability that a new app attracts no ratings during its first month is lower after ATT—although, again, Cohen's d and the equivalence testing suggest that the magnitude of these effects is small.

We see that ATT appears to affect incumbent and new-entrant apps differently. While the introduction of ATT leads to reductions, however small, in the number of new reviews and average reviews, new entrants, who may have factored ATT into their original business model and thus their decision to enter, seem more capable of quick adaptation. Specifically,

[11] Both Cohen's d and equivalence testing support the view that the magnitude of these effects is small.

Table 9: Effect of ATT on the Number of Ratings and Average Rating in the First Full Month of Sale

	No Rating (1)	Log # New Ratings (2)	Avg. Rating (3)
After ATT x Apple	-0.0277***	0.0644***	0.0087
	(0.0009)	(0.0027)	(0.0112)
Platform-Genre FE	✓	✓	✓
Period FE	✓	✓	✓
Cohen's D	-0.0789	0.0615	0.0096
Dependent variable mean	0.84456	0.39457	4.1814
R^2	0.05862	0.05374	0.03544
Observations	2,112,133	2,112,133	328,301

The results in this table use a subsample of the first full month each app was available for sale. Observations are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

we find that consumers' valuations of new products go up after the introduction of ATT. This indicates that the results in Table 8 may be primarily driven by incumbent apps rather than by new entrants. One possible mechanism for this is the effect of ATT on incumbent product updating and, therefore, incumbent quality.

6 Discussion and Limitations

The very limited effects of the ATT framework on the Apple App ecosystem contrast with the large effects found by prior studies on other privacy interventions. Additionally, the nuanced nature of the effects we find calls for exploring potential mechanisms behind them.

In their study of the effects of GDPR on the Google Play Store, Janssen et al. (2022) find that the regulation caused a third of apps to exit and reduced entry by half. This starkly contrasts with the small effects we see. To understand this difference, it is worthwhile to consider how each intervention was handled by the platforms. In response to GDPR, Google established requirements that app developers must follow and removed the apps that did not comply with these requirements. In fact, in 2017, Google started removing apps that did not comply with their data protection and privacy rules (Kollnig and Binns, 2022), and beginning in 2018, apps that didn't meet the minimum criteria of compatibility with the latest versions of Android (Mayya and Viswanathan, 2023). Thus, in the case of Google and GDPR, if developers didn't take explicit actions, their apps could be removed from the market.

Since Apple has complete control over its platform and the design of its policy, it was able to immediately bring all apps operating on its latest operating system into compliance by restricting, from the platform's side, access to the IDFA. That is, while developers could take action to optimize their products to operate under ATT, inaction would lead to a loss of access to the identifier rather than regulatory non-compliance and removal from the store. Thus, all that would happen to an app that had not been updated is that it may receive lower advertising revenues to the extent that targeted advertising earnings differ from untargeted earnings. This, coupled with the low cost of keeping many existing apps in the market, can explain why we do not observe a significant effect of ATT on exit.

Although the effects we find are very small, we do find some evidence of a statistically significant decline in update frequency and in the monthly number of ratings received and average rating. To understand the relationship between these two patterns, we explore how ratings change for apps that update more or less frequently after ATT. We define three groups of apps: *Higher* corresponds to apps with weakly more updates in the nine months after ATT than in the nine months before. 17.4% of apps in the sample fall in this category. *Lower* corresponds to apps with fewer updates in the nine months after ATT than in the nine months before. 21.5% of apps belong to this group. *Zero* corresponds to apps with no updates in the sample period. 61.1% of apps in the sample belong to this group. Given these classifications, we then estimate a variation of our primary model (Equation (2)) in the following way:

$$Y_{i,t} = \beta_1 Post\ ATT_t \times Apple_i \times UpdateFreqChange_i + \omega_i + \mu_t + \epsilon_{i,t}. \tag{3}$$

In this equation, $UpdateFreqChange$ corresponds to a categorical variable capturing

Table 10: Effect of ATT on the Number of Ratings and Average Rating in Relation to Change in Update Frequency

	Log # New Ratings (1)	No Rating (2)	Avg. Rating (3)
After ATT x Apple × Zero	-0.0168***	0.0130***	-0.0334***
	(0.0001)	(0.0001)	(0.0032)
After ATT x Apple × (Weakly) Higher	0.0034***	-0.0012***	0.0133***
	(0.0004)	(0.0002)	(0.0030)
After ATT x Apple × Lower	-0.0319***	0.0114***	-0.0086***
	(0.0003)	(0.0002)	(0.0026)
App FE	✓	✓	✓
Period FE	✓	✓	✓
Cohen's D			
Dependent variable mean	0.20093	0.87548	3.9233
R^2	0.80129	0.58330	0.43389
Observations	77,477,703	77,477,703	9,647,402

The results in this table use a subsample of the first full month each app was available for sale. Observations are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

the three update frequency groups explained above. This model is equivalent to estimating Equation (2) while allowing for the difference-in-difference estimate to vary by update frequency groups. The results presented in Table 10 suggest the negative effect on ratings that we observe in Section 5.3 is driven by apps that (i) never update during our study period, or (ii) reduce their update frequency following the implementation of ATT. Indeed, as we show in column (3), those apps that (weakly) increase their update frequency following ATT see a small increase in their average rating. We see corresponding results for the number of new reviews and the likelihood of an app receiving no reviews in a month. These findings are consistent with prior literature that has suggested a link between update frequency and ratings (Leyden, 2023; Mayya and Viswanathan, 2023).

An interesting pattern observed in Section 5.3 is that incumbents seem to receive fewer and almost indistinguishably lower ratings after ATT, while apps that enter after ATT perform better in terms of ratings than apps introduced before ATT. This pattern could be related to developers' incentives to invest in existing vs. new apps. According to reports by analytics firms, mobile apps are characterized as having a short half-life, with apps seeing most of the users they will ever get in their first six months (Thompson, 2015). Additionally, user retention rates of apps are extremely low, with only 6% of users continuing to use the app 30 days after installation on average (Statista, 2023). Considering the short life span of apps, it is reasonable that in the presence of new requirements, developers may shift their attention to introducing new apps. It also implies that for the long-term health of the app ecosystem it is more important that the new framework does not affect the incentives developers face for introducing new apps, or the engagement and valuation of users with newly introduced apps, than the effects it may have on incumbent apps. Considering that according to our estimations, ATT did not affect entry in the long term and had a positive effect on new app downloads, it is reasonable to conclude the new framework will not negatively impact users or the ecosystem.

Our study is not without limitations. First, while our data allows us to conduct a comprehensive analysis of the effects of ATT on the App Store ecosystem, we are unable to assess whether and to what extent ATT hurt individual developers. Indeed, the reduction in the quality of targeting advertising may have had a significant adverse effect on developers' profitability. For example, in a related context, Aridor et al. (2024) report that e-commerce firms that use targeted advertising saw as much as a 37% reduction in revenue following the implementation of ATT. However, to the extent that an effect is present on that margin, we can conclude that it did not result in a meaningful reduction in platform participation nor the quality of apps on the App Store.

Additionally, while ATT can be considered a fairly exogenous shock due to the short notice developers had about the precise date it would become effective, it is a change that had been announced for several months at the time it was introduced. It is thus possible that developers may have anticipated the changes and by the time the framework was introduced its effects had already been internalized. Additionally, while there does not seem to be any other major changes in the Apple or Android ecosystems during the period we study, there are some changes that may influence our results. Finally, our analysis has centered around free apps, and it is possible that paid apps respond to different dynamics or become more relevant after ATT. To address this concern, we re-estimate the models in the paper, including paid apps, and do not find meaningful differences.[12] Additionally, prior research has indicated that while some developers switch from free to paid business models following ATT, the magnitude of the change is negligible (Kesler, 2022).

7 Conclusions

A primary 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 App Tracking Transparency (ATT) framework influenced the entry, exit, and update dynamics in the Apple App Store ecosystem, the quality of the products in the ecosystem, and the monetization 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 Apple App Store with apps in the Google Play Store before and after the implementation of ATT. As Apple's new privacy 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 meaningful negative long-run effect on the availability of apps. We 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; but the magnitude of that effect, while statistically significant, is very limited. 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

[12] See Appendix B.

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 slightly 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 slightly decreased since ATT.

We also examined how developers adapted their use of SDKs after ATT. Our results show a slight decrease in the use of SDKs that rely on sharing 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 a slight 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 did not significantly affect the availability of apps in the Apple ecosystem, nor did it seem to harm the app ecosystem as a whole. While we find minimal negative effects of ATT on consumers' valuations of apps, this impact is nuanced. It seems driven by existing apps that reduce their update frequency. In contrast, we find an increase in consumers' valuation of apps that continue to update and apps introduced after ATT. Overall our results suggest that developers adapted to ATT by implementing some targeted responses, and that the overall health of the platform was unaffected by these changes.

References

Alomar, N. and Egelman, S. (2022). Developers say the darnedest things: Privacy compliance processes followed by developers of child-directed apps.

Apple Developer (2021). Upcoming apptrackingtransparency requirements. Accessed: 2024-06-17.

Aridor, G., Che, Y.-K., Hollenbeck, B., McCarthy, D., and Kaiser, M. (2024). Evaluating the impact of privacy regulation on e-commerce firms: Evidence from apple's app tracking transparency. Available at SSRN: https://ssrn.com/abstract=4698374 or http://dx.doi.org/10.2139/ssrn.4698374.

Aridor, G., Che, Y.-K., and Salz, T. (2023). The effect of privacy regulation on the data industry: empirical evidence from GDPR. *The RAND Journal of Economics*, 54(4):695–730. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1756-2171.12455.

Axon, S. (2021). 96 https://arstechnica.com/gadgets/2021/05/96-of-us-users-opt-out-of-app-tracking-in-ios-14-5-analytics-find. [Accessed 29-03-2024].

Bian, B., Pagel, M., Tang, H., and Raval, D. (2024). Consumer surveillance and financial fraud. *NBER Working Paper*, (31692).

Bygrave, L. A. (2017). Data protection by design and by default : Deciphering the eu's legislative requirements. *Oslo Law Review*, 4(2):105–120.

Cohen, J. (2013). *Statistical power analysis for the behavioral sciences*. Routledge.

Congiu, R., Sabatino, L., and Sapi, G. (2022). The Impact of Privacy Regulation on Web Traffic: Evidence From the GDPR.

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.

Deisenroth, D., Manjeer, U., Sohail, Z., Tadelis, S., and Wernerfelt, N. (2024). Digital advertising and market structure: Implications for privacy regulation. (32726).

Deng, Y., Lambrecht, A., and Liu, Y. (2023). Spillover effects and freemium strategy in the mobile app market. *Management Science*, 69(9):5018–5041.

Downes, L. (2018). GDPR and the End of the Internet's Grand Bargain. *Harvard Business Review*. Section: Government policy and regulation.

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.

Ershov, D. (2022). Competing with superstars in the mobile app market. *Working Paper*.

Ershov, D. (2023). Variety-based congestion in online markets: Evidence from mobile apps. *American Economic Journal: Microeconomics*.

Federal Trade Commission (2022). Ftc seek comments on trade regulation rule on commercial surveillance and data security, r111004. Docket ID: FTC-2022-0053.

Godinho de Matos, M. and Adjerid, I. (2022). Consumer Consent and Firm Targeting After GDPR: The Case of a Large Telecom Provider. *Management Science*, 68(5):3330–3378. Publisher: INFORMS.

Goldberg, S. G., Johnson, G. A., and Shriver, S. K. (2024). Regulating Privacy Online: An Economic Evaluation of the GDPR. *American Economic Journal: Economic Policy*, 16(1):325–358.

Hoppner, T. and Westerhoff, P. (2021). Privacy by Default, Abuse by Design: EU Competition Concerns About Apple's New App Tracking Policy.

Janssen, R., Kesler, R., Kummer, M. E., and Waldfogel, J. (2022). GDPR and the lost generation of innovative apps. *NBER Working Paper*.

Jin, G. Z., Liu, Z., and Wagman, L. (2024). The GDPR and SDK Usage In Android Mobile Apps.

Johnson, G., Lin, T., Cooper, J. C., and Zhong, L. (2023a). COPPAcalypse? The Youtube Settlement's Impact on Kids Content.

Johnson, G. A., Shriver, S. K., and Goldberg, S. G. (2023b). Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR. *Management Science*, 69(10):5695–5721. Publisher: INFORMS.

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.

Kircher, T. and Foerderer, J. (2023). Does Privacy Undermine Content Provision and Consumption? Evidence from Educational YouTube Channels.

Kircher, T. and Foerderer, J. (2024). Ban targeted advertising? an empirical investigation of the consequences for app development. *Management Science*, 70(2):1070–1092.

Kollnig, K. and Binns, R. (2022). The cost of the gdpr for apps? nearly impossible to study without platform data.

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., Binns, R., Van Kleek, M., and Shadbolt, N. (2022a). Are iPhones Really Better for Privacy? Comparative Study of iOS and Android Apps. *Proceedings on Privacy Enhancing Technologies*, 2022(2):6–24. arXiv:2109.13722 [cs].

Kollnig, K., Shuba, A., Kleek, M. V., Binns, R., and Shadbolt, N. (2022b). Goodbye tracking? impact of iOS app tracking transparency and privacy labels. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM.

Kraft, L., Skiera, B., and Koschella, T. (2023). Economic Impact of Opt-in versus Opt-out Requirements for Personal Data Usage: The Case of Apple's App Tracking Transparency (ATT).

Lakens, D. (2017). Equivalence tests: A practical primer for t tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, 8(4):355–362. PMID: 28736600.

Le Meur, G. (2023). Game over: Examining the impacto f privacy regulations on mobile games. Accepted: 2023-08-08T19:18:10Z.

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). There's an app (update) for that: Product updating under digitization. *Working Paper*.

Leyden, B. T. (2024). 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*.

Lukic, K., Miller, K. M., and Skiera, B. (2023). The Impact of the General Data Protection Regulation (GDPR) on Online Tracking.

Lukovitz, K. (2022). Privacy Update: ATT IDFA Opt-In Rate At 25 https://www.mediapost.com/publications/article/373613/privacy-update-att-idfa-opt-in-rate-at-25-overal.html. [Accessed 29-03-2024].

Mayya, R. and Viswanathan, S. (2023). Delaying Informed Consent: An Empirical Investigation of Mobile Apps' Upgrade Decisions.

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.

Moon, M. (2020). Facebook runs more newspaper ads attacking ios 14 privacy changes. *Engadget*.

Peukert, C., Bechtold, S., Batikas, M., and Kretschmer, T. (2022). Regulatory Spillovers and Data Governance: Evidence from the GDPR. *Marketing Science*, 41(4):746–768.

PwC and IAB (2023). Internet advertising revenue report, fy 2022.

Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*, 90(5):2555–2591.

Ribera, A. (2022). Trading off the orchard for an apple: the ios 14.5 privacy update. *Journal of European Competition Law & Practice*, 13.

Schmitt, J., Miller, K. M., and Skiera, B. (2021). The Impact of Privacy Laws on Online User Behavior.

Schuirmann, D. J. (1987). A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *Journal of Pharmacokinetics and Biopharmaceutics*, 15(6):657–680.

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.

Statista (2023). App retention rate by category. Accessed: 2024-06-17.

Thompson, D. (2015). Mobile apps have a short half-life: use falls sharply after the first six months. *Vox*. Accessed: 2024-06-17.

Tryon, W. W. and Lewis, C. (2008). An inferential confidence interval method of establishing statistical equivalence that corrects tryon's (2001) reduction factor. *Psychological Methods*, 13(3):272–277.

Warberg, L., Lefrere, V., Cheyre, C., and Acquisti, A. (2023). Trends in Privacy Dialog Design after the GDPR: The Impact of Industry and Government Actions. In *Proceedings of the 22nd Workshop on Privacy in the Electronic Society*, WPES '23, pages 107–121, New York, NY, USA. Association for Computing Machinery.

Woodward, E. (2023). Data friction and infrastructural platform power : an analysis of Apple's iOS 14 privacy updates.

Appendix A Facebook Newspaper Advertisement

Figure 4: Facebook campaign against Apple's ATT

Appendix B Paid App Sample

In the main body of our paper, we analyze the sample of "free" apps, or those that set an upfront price of $0 on the App Store and Google Play Store. We focus our attention on these apps because they are the ones that are primarily affected by the implementation of ATT. Moreover, existing work shows that consumers are very reticent to pay for apps, so free and paid apps may reasonably be viewed as competing in separate markets Leyden (2023).

In this appendix, we reproduce our primary results using a sample of paid apps, or those that charge a non-zero price for the duration of our sample. We are unable to reproduce our SDK analysis within this sample because the data provider only provides SDK information for a subset of free apps on the platform.[13]

Table 11: Paid App Sample: Effect of ATT on the Number of Ratings and Average Rating

	Log Entry Count (1)	Log Exit Count (2)	Update (3)	Log # New Ratings (4)	Avg. Rating (5)	No Rating (6)	Price (7)
After ATT x Apple	-0.1624**	-0.1366**	-0.0013***	-0.0193***	-0.0189*	0.0055***	-0.0340***
	(0.0631)	(0.0663)	(0.0003)	(0.0005)	(0.0103)	(0.0004)	(0.0091)
Platform-Genre FE	✓	✓					
Period FE	✓	✓	✓	✓	✓	✓	✓
App FE			✓	✓	✓	✓	✓
Cohen's D	-0.3571	-0.2866	-0.0088	-0.0883	-0.0188	0.0329	-0.0073
Dependent variable mean	3.1180	3.9422	0.03041	0.07960	4.0816	0.94054	5.8344
R^2	0.89509	0.88719	0.29945	0.73312	0.53302	0.54094	0.96588
Observations	817	817	4,175,632	3,952,291	235,013	3,952,291	4,175,632

This table presents results using the set of apps that set an upfront price above $0 for the duration of the study period. Observations in columns (1) and (2) are at the platform-genre-month level. Observations in subsequent columns are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
Robust standard errors in parenthesis.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

We reproduce Table 2 and Table 8 in the combined Table 11. As in the free sample, we find evidence of a reduction in entry, however in this case we also find a decline in exit. Paid app Updates fall to a smaller degree compared to free app updates, although the estimates are close in magnitude. And as in the free app sample, we find evidence of a reduction in the number of new ratings, an increase in the likelihood of earning no new ratings in a given month, and finally a decline in average new ratings.

Note that we are unable to analyze the effect of ATT on SDK adoption in the paid app sample, as our data provider does not offer SDK data for paid apps. Identifying SDKs in apps requires downloading and processing the app files, and so building a substantial dataset of paid app SDKs would have been prohibitively expensive.

Finally, in Table 12, we reproduce the "First Month" analysis of Table 9. The results are comparable across samples.

[13] Analyzing the presence of SDKs in apps requires downloading each app for analysis. Thus, presumably, doing so for a large sample of paid apps is prohibitively expensive.

Table 12: Paid App Sample: Effect of ATT on the Number of Ratings and Average Rating in the First Full Month of Sale

	No Rating (1)	Log # New Ratings (2)	Avg. Rating (3)
After ATT x Apple	-0.0019	0.0088	-0.0504
	(0.0043)	(0.0104)	(0.0800)
Platform-Genre FE	✓	✓	✓
Period FE	✓	✓	✓
Cohen's D	-0.0084	0.0158	-0.0559
Dependent variable mean	0.94350	0.12194	4.3177
R^2	0.02609	0.02380	0.06176
Observations	50,797	50,797	2,870

Observations are at the app-month level. Months are defined as four-week intervals relative to the release of ATT.
This table presents results using the set of apps that set an upfront price above $0 for the duration of the study period. Additionally, 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.
*** $p < 0.01$, * $p < .05$, * $p < 0.1$

Appendix C Equivalence Testing

In order to better understand the magnitude of our estimated effects, we conduct a series of equivalence tests to determine whether our estimates are meaningfully different from zero, which we distinguish from the typical statistical significance test. Specifically, we follow Schuirmann (1987) to implement for each estimated coefficient two one-sided tests (TOST).

We implement these tests in the following way. First, we select upper- and lower-bound criteria. For each outcome variable (e.g., *update*), we follow Lakens (2017) and Lefrere et al. (2022) and construct these bounds as $\pm 0.3 * SD$, where SD is the standard deviation of the outcome (pooled across the entire sample). Then, given an estimated coefficient of β, we conduct two independent t-tests under the following null hypotheses:

- $H0 : \beta < -0.3 * SD$
- $H0 : \beta > 0.3 * SD$

That is, we test whether our estimated effect of ATT is *more* extreme than each of our two bounds. Under this approach, if we reject both null hypotheses, then we conclude that the effect is *practically* zero. Following Lakens (2017), we refer to such a case as "Statistically Equivalent." We note that this equivalence test result is a separate matter from concluding that β is not statistically significantly different from zero.

In Table 13, we present the results of conducting this test for each of the outcomes in Table 2, Table 6, and Table 8. We follow Lakens (2017) in constructing the taxonomy for this table. We report on both equivalence and statistical significance. "Statistically Equivalent" refers to cases where the above two null hypotheses are rejected, and "Statistically Different" refers to the traditional test of statistical significance (which is also reported in the tables in the main text through the use of stars on the regression coefficients). Both sets of tests are conducted at the 5% level.

Tryon and Lewis (2008) provide guidance on interpreting the output of this table. They characterize cases that are "Statistically Equivalent and Statistically Different" as representing a "Trivial Difference." That is, while the estimated coefficient is found to be statistically significantly different from zero, the estimated effect is not meaningfully large given the benchmark threshold of $0.3 * SD$ of the relevant outcome.[14]

[14] Tryon and Lewis (2008) find that the approach to equivalence testing developed in their paper is equivalent to that of Schuirmann (1987).

Outcome	$\alpha = 0.05$
Update	Statistically Equivalent and Statistically Different
Log # New Ratings	Statistically Equivalent and Statistically Different
Avg. Rating	Statistically Equivalent and Statistically Different
No Rating	Statistically Equivalent and Statistically Different
Monetization	Statistically Equivalent and Statistically Different
Ad Mediation	Statistically Equivalent and Statistically Different
Authentication	Statistically Equivalent and Statistically Different
Payments	Statistically Equivalent and Statistically Different
Log Entry Count	Statistically Equivalent and Statistically Different
Log Exit Count	Statistically Equivalent and Not Different

Table 13: Equivalence Test