

Existence, Antecedents and Consequences of Non-Compliance in Mobile App Markets

Reinhold Kesler,¹ Bernd Skiera,² Lennart Kraft,³ and Tim Koschella⁴

This Version: February 28, 2025

First Version: June 20, 2024

Abstract

Digital platforms, now ubiquitous intermediaries in the modern economy, claim to uphold governance rules to ensure a level playing field for their participants. However, there is limited research exploring whether digital platforms fulfill this claim. Furthermore, the antecedents and consequences of any non-compliance remain largely unexamined. This paper addresses this research gap by examining non-compliance in the mobile app market. The empirical study compares the disclosed with the actual behavior concerning device ID transfer for advertising purposes of 852 apps available on Apple and Google platforms across 19 countries. The findings reveal that about 40% of the apps do not comply. Compliance is more prevalent among apps catering to Apple (EU) users than Google (non-European) users. Notably, older apps demonstrate greater compliance. However, popularity and reputation do not explain compliance, while app categories and connections to certain supply-side platforms do. Intriguingly, non-compliant apps earn at least 10% more advertising revenue than they would if being compliant, thus gaining a significant economic edge.

¹ Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf, Universitätsstrasse 1, 40225 Düsseldorf, Germany, email: kesler@dice.hhu.de.

² Goethe University Frankfurt am Main, Theodor-W.-Adorno-Platz 4, 60629 Frankfurt, Germany, email: skiera@wiwi.uni-frankfurt.de. Bernd Skiera is also a Professorial Research Fellow at Deakin Business School, 221 Burwood Highway, Burwood, VIC 3125, Australia.

³ DZ BANK AG, Platz der Republik, 60325 Frankfurt am Main, Germany. Goethe University Frankfurt am Main, Theodor-W.-Adorno-Platz 4, 60629 Frankfurt, Germany, email: lennart.kraft@wiwi.uni-frankfurt.de.

⁴ Real-time Technologies GmbH (Kayzen), Ackerstr. 29, 10115 Berlin, email: tim.koschella@kayzen.io.

Acknowledgement of financial support:

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 833714).

1. Introduction

Digital platforms have revolutionized global commerce by enabling seamless interactions between multiple parties. These platforms should operate on the principle of equal opportunity, governed by rules and enforcement mechanisms that aim to ensure a level playing field. Despite this principle, the reality often diverges because competitive strategies may lead parties to seek unfair advantages such as concealing unpopular behavior like tracking their users. Moreover, platforms might exercise preferential treatment towards certain parties, exploit their market power for their own benefit or choose not to enforce regulations because it is costly (Reimers and Waldfogel, 2023).

However, an unequal playing field on digital platforms brings grave consequences: platforms risk losing reputation and legal compliance, businesses face competitive disadvantages, consumers may not get the most suitable products or overpay, and regulators grapple with unfair competition and consumer harm. These outcomes necessitate effective regulation and enforcement to preserve fairness in the digital economy as foreseen by recent legislation like the European Digital Markets Act or Digital Services Act.

This paper aims to better understand non-compliance on platforms by addressing three research questions: (i) Does non-compliance exist, (ii) what are the determinants of non-compliance, and (iii) what are the economic consequences resulting from non-compliance? Our setting is the market for mobile applications, and the compliance requirement in question is transparent personal data processing, mandating that apps disclose any transfer of personal data. We study the nature of divergence between what apps claim to do with users' data and what they actually do. For this, we compare the disclosed and actual data practices of 852 apps across two platforms (Apple and Google) used in 19 countries and connected to various supply-side platforms in November 2023 and April 2024.

The study posits three hypotheses: first, there are differences in compliance by country as, for instance, stricter regulations like General Data Protection Regulation (GDPR) increase compliance; second, platforms that claim to be more concerned about the privacy of their users, such as Apple, have higher compliance rates; third, businesses that stand to lose economically from non-compliance are more likely to comply, with the implication that more popular or reputable apps, or those with more experience or resources, are more compliant.

In order to study what apps claim regarding their data usage, we first collect apps' disclosed data practices as displayed on Apple ("Privacy Label") and Google ("Data Safety"). These labels require app developers to disclose their data practices publicly in a standardized format, similar to nutrition labels. We focus our analysis on the disclosure of the transfer of the device ID for advertising purposes. This ID is important for advertising because it enables firms to profile users by collecting information about their behavior over time. These profiles help to target users with specific ads, as is evident in retargeting (Lambrecht and Tucker, 2013).

To compare this claimed behavior with their actual behavior, we retrieve data from a large demand-side platform (DSP) for programmatic mobile ads (with >1bn unique users worldwide and >100bn daily ad requests). For all apps in the sample, which make up more than two-thirds of all ad impressions of the DSP, the data encompasses the share of ad impressions displayed to those users for which the apps transfer personal data (here, their device ID) and those users without, along with average prices associated with these shares across the supply-side platforms (SSP) to which the app connects.

We combine the two datasets and define non-compliance by an app disclosing no transfer of users' device ID through the corresponding label while actually transferring it on the ad exchange. We enrich this data with app characteristics, which we web-crawl from each app's landing page on each platform. Thus, we have measures for popularity, reputation, experience, and resources through the number of ratings, the average rating, the age and the number of

installations. The empirical strategy mainly involves explaining non-compliance by country-, platform-, and app-level characteristics while leveraging the same app's availability on two platforms (i.e., Apple and Google).

The share of apps that disclose the transfer of the device ID for advertising purposes is higher on Apple's App Store, but tracking the device ID is much more prevalent on Google's Play Store. Tracking, however, should only occur if the app discloses doing so. Surprisingly, about 40% of apps across both platforms in our sample do not comply with their disclosed data usage, driven by a higher non-compliance rate on Google than Apple.

There is little difference between the apps' countries of origin, though apps are more compliant when targeting European users, on average, by 4 percentage points. App characteristics like popularity and reputation do not determine non-compliance, whereas older and larger apps tend to be more compliant. Finally, app categories and the connected supply-side platforms of an app explain sizable differences in non-compliance.

Based on the price premium of ads targeted toward users tracked with a device ID, we compute the economic advantage of non-compliant apps. Specifically, we compare their realized prices to a counterfactual where they had not transferred the device ID for advertising purposes and stayed compliant. The resulting economic advantage amounts to, on average, 10% more revenue on Apple's and 70% on Google's app platform.

In conclusion, non-compliance is prevalent on digital platforms, and non-compliant apps reap significant economic gains. Nevertheless, stricter regulation seems to mitigate non-compliance. These results highlight the importance of combating non-compliance. Accordingly, platforms should enforce compliance, businesses should track competitor adherence, consumers should trust carefully, and regulators should ensure these standards are rigorously upheld.

2. Industry Background

The landscape of our investigation is the dynamic and expansive market for mobile applications, which achieved a staggering global revenue of 533 billion US dollars in 2023, encompassing over 4 million apps across the two leading platforms, Apple and Google.¹ This vibrant ecosystem derives one-third of its revenues from sources like in-app purchases and direct app sales; two-thirds emanate from advertising revenues, predominantly generated through real-time auction processes for each ad view.

Navigating the intricacies of the mobile advertising ecosystem unveils a complex network of interactions. At its core, app developers use Supply-Side Platforms (SSP) that run the auctions to monetize the app's digital real estate by offering ad inventory; advertisers employ Demand-Side Platforms (DSP) to acquire ad inventory tailored to their campaigns. Information about the user often improves such tailoring so that more information increases the prices of ads (Johnson et al., 2020, but also our analysis below). Thus, apps usually have a strong economic incentive to collect and transfer user data, making our analysis of the enforcement to preserve user privacy particularly attractive.

Specifically, we focus on the rule that mandates transparency about the processing of personal data. It is at the heart of privacy legislation such as the GDPR or California Consumer Privacy Act (CCPA), and platforms ought to implement mechanisms to warrant it. Therefore, developers of apps are required to disclose their data practices in a standardized format on Apple (known as "Privacy Label") and Google (referred to as "Data Safety"), analogous to the way nutrition information is presented (see Appendix A.1). This data also includes the transfer of the user's device ID, i.e., the Identifier for Advertisers (IDFA) on Apple and Google Advertising Identifier (GAID), often also referred to as the Android Advertising ID (AAID).

¹ See <https://www.data.ai/en/go/state-of-mobile-2024>, last accessed on 27th February, 2025.

The transfer of this device ID helps to improve third-party advertising. For example, retargeting would not be feasible otherwise.

The basic idea of our study is to measure the discrepancy between what firms claim in their disclosure of data practices and what firms actually do with user data as measured by the availability of device IDs. Accordingly, we define non-compliance as firms disclosing not to transfer a device ID for advertising purposes while actually transferring it.

3. Related Literature

Our study contributes to three strands of literature. First, we contribute to research on the value of user data, both from the perspective of a firm and a user, as it may explain the benefit of tracking users and a firm's decision not to comply. Specifically, research has shown that the ability to identify and track users leads to increased revenues for publishers through targeted advertisements. These studies exploit privacy changes that, for instance, involved less tracking, e.g., by disabling cookies (Goldfarb and Tucker, 2011; Johnson et al., 2020), introducing GDPR (Wang et al., 2024), banning the usage of personal data (Sun et al., 2023) or providing users with more granular choices to prevent tracking (Kraft et al., 2024).

Our results confirm that more information about users increases the price of displaying ads. In addition, they outline that the resulting price differences seem to be high enough to encourage non-compliant behavior by non-disclosing the transfer of (privacy-sensitive) information for advertising purposes. Reasons for the non-disclosure could be that users may negatively react, as suggested by the high opt-out rates from tracking (Kraft et al., 2024) and the decreased demand when intrusive data practices are revealed (Bian et al., 2023).

Second, our study relates to research studying firms' non-compliance in a wide range of settings involving, for instance, regulations on taxes (Slemrod, 2019), labor markets (Ashenfelter and Smith, 1979), or the environment (Blundell et al., 2020). We contribute by

looking into the increasingly regulated digital sphere, where we consider our measurement of non-compliance to be particularly valid.

Third, our study is part of research studying non-compliance in the context of apps’ data practices, which mainly belongs to computer science. Figure 1 shows an overview of related literature, where the majority of these studies find non-compliance. However, they do not elaborate on the reasons for its existence or the economic consequences. Our contribution is to address the latter with cross-platform and cross-country evidence. Moreover, we consider our context as a more deliberate non-compliance on the part of an app developer than when considering privacy policies or all of an app’s network traffic (which would require isolating the traffic of interest). The reason is that transferring a device ID for advertising is commercially consequential.

Figure 1: Overview of Previous Studies on Apps’ Non-Compliance

Study	Comparing Privacy Labels with? actual data?	Apple	Google	Country	3p-tracker	2p-tracker	Extent of Non-Compliance	Determinants of Non-Compliance	Economic Consequences	
Koch et al. (2022)	Network traffic	✓	✓	✗	Germany	✓	✗	16% of apps	types/purposes	✗
Kollnig et al. (2022)	Network traffic	✓	✓	✗	United Kingdom	✓	✓	80% of apps	types/purposes	✗
Xiao et al. (2022)	Network traffic	✓	✓	✗	United States	✓	✗	67% of apps	types/purposes and categories	✗
Jain et al. (2023)	Privacy policy	✗	✓	✗	United States	✓	✗	88 % of apps	types and popularity of apps	✗
Ali et al. (2023)	Privacy policy	✗	✓	✗	United States	✓	✗	56 % of apps	types/purposes and case studies	✗
Khandelwal et al. (2023)	Privacy policy	✗	✓	✓	United States	✓	✗	40 % of apps	types/purposes, across & survey	✗
This Study	Demand-side platform (DSP)	✓	✓	✓	19 countries	✓	✓	... % of apps	country-, platform- and firm-level characteristics	✓

Notes: The selection follows a systematic search of literature covering combinations of “privacy labels”, “apps”, and “non-compliance”.

4. Conceptual Framework

We apply rational choice theory (Becker, 1968) to study app developer’s non-compliance and its drivers and consequences. Accordingly, a trade-off in an expected utility framework is made between profits when being and not being compliant. Losses from non-compliance can arise due to fines, where both the detection probability and the size of the fines play a role.

Additionally, losses can come from decreases in revenue because users negatively react to a possible revelation of non-compliance with a certain probability. Gains from non-compliance are increasing revenues as advertisers pay a higher ad price for users with a device ID. The difference in resulting profits determines the decision not to comply, thereby taking possible losses and gains of non-compliance along with different context factors into account. A particular aspect will be whether app developers make a uniform or differentiated decision in different contexts for both the firm's claimed behavior (i.e., its disclosure) and actual behavior regarding the device ID.

At the country level, higher fines (particularly regulations, e.g., GDPR) and a higher detection likelihood lead to higher expected losses, thus leading to more compliance. For instance, European data protection authorities imposed 665 fines until February 27, 2025, because of "non-compliance with general data processing principles" (according to enforcementtracker.com). At the platform level, those with more transparency (e.g., App Tracking Transparency on Apple) increase the detection likelihood and yield higher compliance because of these expected losses. Additionally, Apple has more to lose if it turns out that apps on Apple are not compliant. At the firm level, similarly, apps with more revenues to lose due to negative user reactions, be it because of their popularity or reputation being at stake, are more likely to be compliant. In addition, apps with more experience and resources face lower compliance cost and thus, their compliance increases.

Accordingly, we have three main hypotheses for the determinants: first, there are differences in compliance by country as, for instance, stricter regulations like GDPR lead to higher compliance than other or even no regulation; second, platforms providing more transparency about the usage of personal data, such as Apple, see higher compliance rates; third, businesses that stand to lose economically from non-compliance are more likely to comply, with the implication that more popular or reputable apps, or those with more experience and resources,

are more compliant. Apps trade off these factors against economic advantages or gains from non-compliance. Therefore, we derive the economic consequences of being non-compliant (see details in Section 6.3).

5. Description of Data

5.1. Sampling

To create our sample, we depart from apps offering their inventory space on a demand-side platform. Then, we rank these apps of each of the two platforms by the number of auctions, reflecting the number of ad impressions, on November 15, 2023, giving us the Top 1,000 apps on each platform. To identify the same apps across platforms (here referred to as twins) to leverage cross-platform differences, we take these 2,000 apps and look each up on Apptopia. This mobile analytics platform shows for a specific app on a specific platform if an identical app is also available on the other platform. Based on this information, we infer 852 twin apps, i.e., an identical app active on both platforms.

We examine these apps in 19 countries, i.e., Australia, Brazil, Canada, France, Germany, India, Indonesia, Italy, Japan, Mexico, Netherlands, Poland, Romania, Russian Federation, Spain, Switzerland, Thailand, United Kingdom and United States. These are important countries for the DSP and exhibit considerable heterogeneity in their regulatory environment regarding data protection, e.g., with seven member states of the European Union and subject to the GDPR. However, there is a further layer for each app-platform-country combination, as we can observe for each app to which supply-side platforms it connects – which is, on average, seven. Hence, our unit of observation is an app on a platform catering to users of a country and connected via an SSP.

In the following, we first describe the two datasets that enable non-compliance measurement by an app disclosing no transfer of users’ device ID through the corresponding label (Section 5.2) while actually transferring it on the ad exchange (Section 5.3). We then turn toward the third dataset (Section 5.4), which comprises app characteristics necessary for measuring the third main hypothesis, and conclude with descriptive statistics on all measures (Section 5.5). We provide additional details on the data in Appendix A.2.

For our primary analyses, we use all the datasets collected in November 2023, including the monthly data from the DSP. For selective analyses, we repeat the data collection in April 2024.

5.2. Disclosure Data

In order to study what apps claim regarding their data usage, we first collect apps’ disclosed data practices as displayed on Apple (“Privacy Label”) and Google (“Data Safety”). These labels require app developers to disclose their data practices publicly in a standardized format, similar to nutrition labels. We focus our analysis on the disclosure of the transfer of the device ID for advertising purposes. We ensure country-specific versions by pre-specifying URLs (e.g., for the USA, “&hl=en&gl=US”) and crawling from the corresponding regions through VPNs.

5.3. Behavioral Data

To compare claimed behavior in the disclosed data practices with what apps actually do with users’ data, we retrieve data from a large demand-side platform for programmatic mobile ads (catering more than 100bn daily ad requests to more than 1bn users). For all apps in the sample, the data encompasses the share of ad impressions displayed to trackable and non-trackable users, i.e., ad impressions with and without a user’s device ID (for advertising).² Furthermore,

² For apps on Apple, we additionally see whether the ad impressions come along with the identifier for vendors (IDFV), which is explained in more detail in the respective robustness check in Section 6.2.2.

the corresponding average (and median) prices are available, also distinguished by ad format. Importantly, we observe the share of ad impressions displayed to trackable and non-trackable users for every supply-side platform to which an app connects, which is another source of variation. Finally, the DSP classifies (twin) apps into specific categories (“smart app categories”), which we use rather than the ones provided by two platforms that do not necessarily overlap across the two. Accordingly, Figure 8 of the Appendix shows that 76% of apps belong to games (Casino, Casual, Core, Hypercasual); the remaining categories are Entertainment, News and Weather, Social, Tools and Utility, and Other. The large share of games among ad-based apps aligns with statistics of advertisement software development kits.³

5.4. App Characteristics Data

We enrich this data with app characteristics, which we web-crawl from the country-specific app’s landing page on each platform. While this data validates that almost all apps include advertisements (and offer free versions), they also provide measures for popularity, reputation, experience and size through the number of ratings, the average rating, the age, and the number of installations. The app’s age is inferred from Apptopia and validated by information provided in the meta-data of the app’s landing page (if available). We indicate an app as large if it surpasses 10 million installations on Google because only Google reveals this information. Figure 9 of the Appendix shows the distribution for the number of installations across apps, showing that our sample is normally distributed around 10 million installations with very large and small apps. On Google’s Play Store, app developers can disclose their physical address. We use it to infer the origin of the app developer. As the information is not uniformly formatted (and not mandatory), we use the information, if available, as input in Google Maps to extract the country. This process yields a country for more than 81% of apps, of which 74% belong to

³ See <https://42matters.com/sdk-analysis/top-ad-networks-sdks>, last accessed on 27th February, 2025.

the Top 10 countries (US, CY, SG, HK, FR, UK, IL, DE, AE, TR), as displayed in Figure 10 of the Appendix, while the remaining 26% distribute across 45 countries.

5.5. Descriptive Statistics

Table 1 provides summary statistics of all the before-mentioned variables distinguished by the platform. The apps of our sample make up more than two-thirds of all auctions. The number of auctions follows a right-skewed distribution. While Section 6.1 describes the components of non-compliance, Table 1 shows that about one-third (one-fourth) of users (developers) reside in the European Union.

The information on ratings suggests our apps are very popular and reputable because the number and grades surpass industry averages. For the latter, less than 30 percent of apps are rated at all, and the mean rating is 4.1, as reported by AppBrain for Google. Our apps have a similar age across the two platforms, with an average of more than 5 years, and two-thirds of observations belong to large apps according to our definition of more than 10 million installations on Google. Lastly, almost all these apps have in-app payments (in addition to advertising) on both platforms.

Table 1: Summary Statistics of 852 Apps on both platforms (Apple and Google) and up to 19 countries

	Apple			Google		
	Mean	SD	N	Mean	SD	N
Number of Auctions (in millions)	929.75	3,194.94	836	1,235.41	3,958.51	807
Share of Trackable Traffic	0.31	0.13	836	0.78	0.11	807
Share of Actual Device ID Transfer	0.71	0.17	836	0.89	0.10	807
Share of Disclosed Device ID Transfer	0.82	0.38	836	0.67	0.47	807
Share of Non-Compliance	0.11	0.27	836	0.29	0.42	807
Share of EU Users	0.30	0.10	836	0.32	0.10	807
Share of EU Developers	0.23	0.42	836	0.25	0.43	807
Number of Ratings (in thousands)	41.89	98.49	836	1,147.56	3,129.21	807
Average Rating	4.49	0.33	835	4.28	0.32	807
App Age (in Years)	5.81	3.68	835	5.37	2.92	804
Large Dummy (10M+ Installations)	0.66	0.47	791	0.66	0.47	807
Share of In-App Prices	0.90	0.30	836	0.90	0.30	807
N	836			807		

Notes: The number of apps for each platform is lower than 852 as not every app of our twin combination is available on the respective platform or demand-side platform (DSP).

6. Results


6.1. Existence of Non-Compliance

We derive non-compliance by comparing the firm's claimed behavior (i.e., its disclosure) with its actual behavior regarding the device ID. Table 1 reveals that apps on Apple are more likely to disclose the transfer of the device ID and that tracking of the device ID is more prevalent on Google's Play Store. However, Table 1 does not reveal whether an app's decision to disclose and transfer is uniform in all settings or differs across settings, e.g., countries. Therefore, we utilize the variety of observations for each app developer in the following to track both the claimed and actual behavior in different contexts. The number of observations of each app on both platforms amounts to up to 19 countries, multiplied by the number of supply-side platforms to which the app connects.

Starting with the actual transfer (defined as the trackable traffic being above 5%) in the left panel of Figure 2, we find that (almost) all apps transfer the device ID at least once on both platforms. In Figure 11 of the Appendix, we define three categories: whether the transfer is made nowhere ($< 10\%$), somewhere (10-90%), or everywhere ($> 90\%$) across the different observations for each app. On Apple, the vast majority of apps only transfer "somewhere"; on Android, this share is only about 40 percent, with the remainder being "everywhere". Hence, apps differentiate with respect to their decision to transfer the device ID, which may hint at an intentional choice.

For the disclosure behavior, the right panel of Figure 2 suggests that many apps disclose the device ID transfer at least once, more often on Apple, however. Interestingly, for more than 30 percent of apps, there is a disclosure on one platform, not the other (table not reported). Looking at the distribution of disclosure within an app in Figure 12 of the Appendix, we find that apps disclose the transfer either everywhere or not at all, i.e., apps do not selectively disclose across countries.


Figure 2: Distribution of Transfer & Disclosure of Device ID (at least once) by Platform



Notes: The analysis is done at the app level. Each app on both platforms has observations for up to 19 countries multiplied by its connected supply-side platforms. For each app, we compute whether there is at least one observation across all observations of an app denoting a transfer (disclosure) on the right (left) panel. We receive the share of such apps translating to the x-axis. The y-axis distinguishes the platform.

While the above analyses focus on whether there is a transfer, we also look at how often the device ID is transferred, as it can explain the decision to non-comply. For this, Figure 3 displays the share of trackable traffic (i.e., ad impressions with device ID) for both platforms, which complements Table 1. Distinguishing whether the app disclosed the transfer of the device ID (left panel) or not (right panel) shows distinctive differences across platforms.

Figure 3: Distribution of the Share of Trackable Traffic per App (by Disclosure and Platform)



Notes: The left panel shows the share of trackable traffic for apps that disclosed the transfer of the device ID, while the right panel shows it for apps that disclosed not to transfer. Hence, we divide our total number of observations (on the level of app-platform-country-SSP) into those with and without a disclosed transfer.

The share of trackable traffic should be zero for apps that disclose not to transfer the device ID. The share of trackable traffic for apps that disclose the transfer of the device ID is often lower than 100% because many users do not allow apps to transfer it.

Reading example: The average value of the share of ad impressions with a device ID on Google – for observations with a disclosure of the transfer (left panel) – is 0.78. This value corresponds to a share of trackable traffic of 78% and means that 78% of advertising auctions consist of users for which the advertiser receives the users’ device ID. It represents compliant behavior because the app disclosed the transfer of the device ID.


While there is a bunching of apps on Apple around 0 for the trackable traffic if they do not disclose such practice, on Google, the pattern is qualitatively similar to the distribution with disclosed device ID transfers. In case of non-disclosure, we consider an app non-compliant if the trackable traffic is above a threshold of 5%. Figure 3 shows that this threshold is not crucial.⁴ Table 1 reveals that apps have, on average, 19.9% of their observations across both platforms not in compliance with their disclosed data usage. Specifically, Apple's App Store has an average of 11.3% of apps being non-compliant, whereas it amounts to 28.8% on Google's Play Store.

The summary statistics at the observation level (see Table 3 of the Appendix) show similar averages. If we look at how many apps are non-compliant at least once, this share increases to almost 40% for an app (composed of twins) and to 17% and 33% for Apple and Google, respectively (table not reported). Studying similarly in Figure 3, we find substantial differentiation in whether apps decide to non-comply nowhere, somewhere, and everywhere; so, apps are selectively non-compliant across their observations.

We shed further light on this differentiation in Figure 13 of the Appendix. There, for each app, we compare the non-compliance rate (i.e., the share of observations corresponding to non-compliance) on both platforms by the user location. While the evidence suggests more non-compliance on Google, we also find an increase in non-compliance for both platforms when the user resides in non-European countries.

⁴ In fact, the results remain similar with thresholds at 0 (which is the value that should occur if apps do not transfer the device ID) and 10 percent (which is a value that allows for a rather high number of accidental transfers, e.g., because of technical glitches).

Figure 4: Differentiation of Non-Compliance by Platform



Notes: Non-Comp. refers to Non-Compliance. The analysis is done at the app level. Each app on both platforms has observations for up to 19 countries multiplied by its connected supply-side platforms. For each app, we compute how many observations denote non-compliance. For example, an app is available in 18 countries and connects to 6 supply-side platforms in each country. If we observed non-compliance in 3 countries on all of its connected supply-side platforms, then the share of non-compliant observations would be $(3 \times 6) / (18 \times 6) = 16.67\%$.

As a result, every app belongs to one of the three categories: non-compliance nowhere ($< 10\%$), somewhere ($10\text{-}90\%$), or everywhere ($> 90\%$). Results on the right (left) panel correspond to the Google (Apple) sample.

6.2. Antecedents of Non-Compliance

In this subsection, we use regression analysis to analyze the antecedents systematically.

6.2.1. Baseline Estimations

We aim to explain non-compliance by country-, platform-, and app-level characteristics and leverage the availability of the same app on two platforms and different supply-side platforms. Therefore, we estimate the following linear probability model:

$$Y_{ijcs} = \beta_0 + \beta_1 EU\ User_c + \beta_2 EU\ Dev_i + \beta_3 Google_j + SSP_s + X_{ijc} + \varepsilon_{ijcs} \quad (1)$$

Accordingly, Y corresponds to a dummy variable indicating non-compliance (Yes = 1, No = 0). The unit of observation is an app (i), on a platform (j), targeting users in a specific country (c) and connected to different supply-side platforms (s). Information on the location of the app user ($EU\ User$) and app developer ($EU\ Dev.$) is dichotomized by distinguishing EU vs. non-EU. Dummy variables indicate the platform ($Google = 1, Apple = 0$) and the supply side platform (equal to 1 for the respective exchange, 0 otherwise). X captures app-level characteristics, which comprise measures of the third hypothesis along with control variables such as the presence of in-app payments and the category.

Table 2: Antecedents of Non-Compliance (Linear Probability Model)


	Non-Compliance Dummy {Yes = 1, No = 0}					
	(1)	(2)	(3)	(4)	(5)	(6)
EU User	-0.04*** (0.00)	-0.04*** (0.00)	-0.04*** (0.00)	-0.04*** (0.00)	-0.04*** (0.00)	-0.04*** (0.00)
EU Developer	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)
Google		0.16*** (0.02)	0.17*** (0.02)	0.15*** (0.02)	0.15*** (0.02)	0.15*** (0.02)
Log. (# Ratings)			-0.00 (0.00)			
Average Rating				-0.01 (0.02)		
App Age (in Years)					-0.02*** (0.00)	
Large Dummy (10M+)						-0.04+ (0.02)
In-App Price Dummy			0.02 (0.03)	0.02 (0.03)	0.02 (0.03)	0.02 (0.03)
Category FE	Yes	Yes	Yes	Yes	Yes	Yes
SSP FE	Yes	Yes	Yes	Yes	Yes	Yes
Mean Dep. Var.	0.16	0.16	0.16	0.16	0.16	0.16
R ²	0.04	0.09	0.09	0.09	0.10	0.09
N (Apps)	852	852	851	851	852	807
N (Total)	200552	200552	196903	196903	198224	194342

+ p<0.10, * p<0.05, ** p<0.01, *** p<0.001. SSP refers to Supply-Side Platforms; Dep. Var. to Dependent Variable. FE denotes Fixed Effects.

Standard errors are clustered at the twin level. Reference categories are non-EU and Apple.

Table 2 displays the regression results, which show some differences between countries. While apps are more compliant when targeting European users (by 4 percentage points), being a developer in the EU is not statistically significant. More importantly, apps are considerably more often non-compliant on Google than on Apple. However, app characteristics like popularity (measured by the number of ratings) and reputation (measured by the average rating) do not seem to determine non-compliance. In contrast, more experienced (measured by older apps) and apps with more resources (measured by larger apps) tend to be more compliant. Finally, the coefficients plot displayed in Figure 5 shows category-specific differences in non-compliance. Specifically, Casino Games are most likely to be non-compliant. For the SSPs, Figure 3 suggests that non-compliance is more likely among the largest (in terms of the number of auctions) and includes market leaders like Google Adx and Vungle.

Figure 5: Coefficients of Baseline Regression for Categories and Supply-Side Platforms (SSPs)



Notes: The plots show coefficients of the previously indicated control variables of the baseline regression explaining non-compliance by categories (left) and SSPs (right). Reference categories are casual games and BidSwitch, respectively. The plot contains the estimate and the 95 % confidence interval.

The results remain qualitatively unchanged when we include all app-level characteristics in one specification. While the app-level measures may correlate and confirm the hypothesized signs, the varying statistical significance may suggest differences in the underlying drivers of non-compliance.

The empirical evidence supports hypotheses 1 and 2, while the results on the firm-level characteristics are less conclusive and point to experience being a factor for non-compliance.

The results with alternative model specifications remain unchanged, as shown in Table 4 of the Appendix, e.g., exploiting within-twin and within-app variation. Dropping Casino Games from our sample, our results remain qualitatively unchanged (table not reported).

6.2.2. Alternative Identifier

The IDFA is at the heart of Apple’s App Tracking Transparency (ATT), salient to users (through prompts), and is more important to Apple. Thus, Apple might examine compliance more for this identifier but less for another (less prominent) identifier, namely the Identifier For Vendors (IDFV). Therefore, we repeat our analysis for the IDFV. It is an identifier that uses the same ID for a user in apps belonging to the same firm. Our results indicate that non-compliance for the IDFV is, on average, 10.4 (8.5) percent on the app level (observation level), comparable to the values for IDFA, as shown before.

Moreover, we repeat our baseline model from Table 5 of the Appendix. We only replace the IDFA with the IDFAV when it comes to the identifier to compute non-compliance. The Google dummy variable is missing as the IDFAV is unavailable for that platform. The corresponding regression results in Table 5 of the Appendix suggest that the coefficients for European users remain negative and statistically significant for a less obvious identifier on Apple. All app characteristics of hypothesis 3 are statistically insignificant, whereas the patterns for categories and SSPs remain substantively similar.

6.2.3. Variation over Time

In this final robustness check, we aim to examine the consistency of app developers' behavior over time by comparing their actions in November 2023 and April 2024. Specifically, we examine whether, how, and where (e.g., platform or country) behavioral changes occur during this period, considering potential shocks such as impacts from the European regulations enacted at the beginning of 2024. We restrict the following analysis to apps observed twice within the same platform, country, and supply-side platform to ensure an appropriate comparison. 850 out of 852 apps fulfill this restriction, yielding 150,000 observations.

We depart from the finding that in the majority of observations, the decision to comply or non-comply persists (95 percent). The remaining observations suggest a slightly growing non-compliance. Focusing on apps changing their decision to comply or non-comply, we find that compliance is achieved by stopping the transfer rather than disclosure and, similarly, non-compliance by starting to transfer. In most cases, the latter happens for apps that were already non-compliant in other instances. We repeat our baseline regression in Table 6 of the Appendix by pooling the two cross-sections and considering app fixed effects. While our key results remain, we also find an increasing divergence between apps catering to EU vs. non-EU users. Specifically, non-compliance rates decrease over time for EU users while it slightly increases for non-EU users.


6.3. Consequences of Non-Compliance

We also examine the economic advantage of not sticking to the disclosed data practice. Therefore, we build upon the idea that advertising prices are higher when targeted toward users trackable through device IDs (see Kraft et al., 2024). These higher prices give non-compliant apps an economic advantage by enabling them to increase the advertising revenue of their trackable traffic. To quantify this advantage, we calculate the “realized price” by the weighted average of trackable and non-trackable traffic prices. The weights are the shares of trackable and non-trackable traffic. We then compare this realized price to a “counterfactual” or “legitimate price,” derived from the hypothetical scenario where all traffic is non-trackable and priced accordingly. The economic advantage is then the percentage difference between the realized price and this counterfactual price, indicating the respective revenue increase and, thus, quantifying the financial benefit gained through non-compliance.

We have price data on the observation level of an app on a platform in a country, thus aggregating all SSPs compared to our previous analyses.⁵ As a result, we have the prices for each app-platform-country combination by ad format (banner, interstitial, rewarded) and the availability of the device ID. We follow Kraft et al. (2024) to back out the price premium associated with trackable traffic compared to non-trackable traffic. So, we regress (average) prices of ads on the availability of the device ID, country, and platform, along with the ad format as a control variable and app fixed effects. We also weight the regression by the number of auctions used to calculate the ad prices. The regression gives us the price difference by the presence of the device ID per country and platform.

⁵ Table 7 of the Appendix shows that the baseline regression results remain qualitatively unchanged.

Figure 6: Price Premiums (Derived From Device ID Coefficient) Plots Across Countries and Platforms



Notes: The plots show coefficients of the device ID dummy variable of the following regression by countries for Apple (left) and Google (right): $\log(\text{average price})$ weighted by the number of auctions and explained by the availability of the device ID, country, and platform, along with the ad format as a control variable and app fixed effects. The plot contains the estimate and the 95 % confidence interval. The impact of the device ID's availability on the ad's price is $(\exp(x)-1)$, where x represents the value of the coefficient of the device ID.

For example, the device ID coefficient is 0.555 for the United States and the Apple platform. Thus, $\exp(0.555)-1=74\%$. So, the price for trackable traffic in the US for Apple devices is, on average, 74% higher than for non-trackable traffic in the US for Apple devices.

Country codes: AU stands for Australia, BR for Brazil, CA for Canada, FR for France, DE for Germany, IN for India, ID for Indonesia, IT for Italy, JP for Japan, MX for Mexico, NL for Netherlands, PL for Poland, RO for Romania, RU for Russian Federation, ES for Spain, CH for Switzerland, TH for Thailand, UK for United Kingdom, US for United States.

A simple example shall illustrate the derivation of the economic advantages. Suppose a non-compliant Apple app in the United States has 60 % of trackable traffic. The price premium for Apple devices in the US is 74% (see Figure 6). Thus, we can combine this information to compute the economic advantage as follows:


$$Avg. Price_{Non-Compliant} = (60\% * 1.74 USD) + (40\% * 1.00 USD) = 1.444 USD$$

$$Avg. Price_{Compliant} = (0\% * 1.74 USD) + (100\% * 1.00 USD) = 1.00 USD$$

$$Economic\ advantage = \frac{Avg. Price_{Non-Compliant}}{Avg. Price_{Compliant}} - 1 = 44.4 \%$$

We repeat this calculation for both platforms and all apps in our sample. On average, non-compliant apps generate 10% and 70% higher prices and, thus, more advertising revenue on Apple's and Google's app platforms. Figure 7 provides the corresponding distribution.

Figure 7: Distribution of Economic Advantage Across 852 Apps and both Platforms (Apple and Google)



Note: The figure is based on all non-compliant observations (departing from an app active on two platforms across 19 countries). The mean of 11.06 in the Apple sample means that, on average, the non-compliant apps realize 11.06% higher advertising revenue because they transferred the device ID despite the apps telling users it would not do so (our counterfactual). The economic advantage occurs because advertisers pay higher ad prices if they know the users’ device ID.

7. Summary, Conclusion, and Implications

In summary, this paper reveals a significant prevalence of non-compliance in the market for mobile applications regarding the transparency of personal data usage. Comparing apps’ disclosed versus actual data practices in the context of advertising, we find a notable size and discrepancy in non-compliance rates across Apple’s and Google’s platforms. Tracking and non-compliance are more prevalent on Google’s Play Store than on Apple’s App Store. Additionally, the lower non-compliance rates observed on Apple extend beyond the prominent Identifier for Advertisers (IDFA) to a less prominent device identifier. Our analysis further reveals that apps differentiate with respect to their decision to transfer the device ID across the countries of the users and the supply-side platforms to which the app connects.

In contrast, apps choose to disclose the transfer of such data uniformly across jurisdictions. As a result, we show some variation of compliance across different regions, where stricter

regulations, such as the GDPR, appear to encourage better adherence. Non-compliance remains relatively stable over our observation period. We only observe a slight decrease when catering to European users. Interestingly, app characteristics, apart from age, play a minor role in the decision not to comply. In contrast, the category and the supply-side platform to which the app connects explain sizeable differences in non-compliance. Most importantly, non-compliant apps gain a significant economic advantage of at least 10 percent through higher advertising revenue for their (non-disclosed) trackable traffic.

Our study makes several key contributions to understanding non-compliance in mobile app data practices. First, we provide direct evidence of discrepancies between disclosed and actual data-sharing behaviors across Apple's and Google's ecosystems, highlighting persistent transparency challenges despite newly introduced instruments like privacy labels. Our granular analysis shows that apps vary in transferring device identifiers across countries and supply-side platforms, indicating systematic and potentially deliberate non-compliance. Second, our measure of non-compliance is more direct than previous research and more consequential in our context as it translates into a clear monetary advantage for non-compliant apps. Third, we demonstrate that stricter privacy regulations, such as the GDPR, are related to better compliance, though the development over time is only modest. Fourth, our main results are robust across various analyses, remaining stable even when accounting for potential measurement errors. Finally, we identify the app's age, category, and supply-side platform affiliation as determinants of non-compliance, offering insights for policymakers seeking more effective enforcement.

Despite the principles of equal opportunity on a level playing field that digital platforms purport to uphold, our paper shows a considerable divergence. This divergence also suggests that the benefits of non-compliance outweigh possible costs. Hence, the findings underscore the need for increasing the costs of non-compliance through rigorous enforcement by platforms as well

as more stringent regulatory oversight, ensuring that standards are strictly applied and violations appropriately penalized. This enforcement would lead to fairness and safeguarding the interests of all parties involved, which also involves monitoring the fair play of their competitors and cautious trust by consumers. A comprehensive approach is vital for maintaining the integrity of digital platforms and their users' trust, thereby ensuring a truly level playing field in the rapidly evolving digital landscape and an effective implementation of regulations in place worldwide.

While our study provides insights into non-compliance in mobile app data practices, several avenues for future research emerge. First, future work may examine the long-term impact of evolving privacy regulations, assessing whether and by how much stricter enforcement leads to sustained improvements. Second, our finding that certain supply-side platforms align with higher non-compliance highlights the need to explore the ad ecosystem's value chain. Future research could investigate the role of intermediaries in enabling or discouraging non-compliant practices. Third, while our study focuses on device identifier transfers, future work could look at other sensitive and valuable data types, such as location data, to determine whether similar non-compliance patterns persist. Finally, given that non-compliant apps gain a measurable financial advantage, future research should explore the broader market distortions caused by non-compliance, including long-run market outcomes of compliant vs. non-compliant firms.

References

- Ali, M. M., Balash, D. G., Kanich, C., & Aviv, A. J. (2023). Honesty is the Best Policy: On the Accuracy of Apple Privacy Labels Compared to Apps' Privacy Policies. arXiv:2306.17063v1 [cs.CR].
- Ashenfelter, O., & Smith, R. S. (1979). Compliance With the Minimum Wage Law. *Journal of Political Economy*, 87(2), 333-350.
- Bian, B., Ma, X., & Tang, H. (2023). The Supply and Demand for Data Privacy: Evidence from Mobile Apps. Available at SSRN 3987541.
- Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, 76(2), 169-217.
- Blundell, W., Gowrisankaran, G., & Langer, A. (2020). Escalation of Scrutiny: The Gains From Dynamic Enforcement of Environmental Regulations. *American Economic Review*, 110(8), 2558-2585.
- Goldfarb, Avi and Catherine Tucker. (2011). Privacy Regulation and Online Advertising, *Management Science*, 57(1), 57–71.
- Jain, A., Rodriguez, D., del Alamo, J. M., & Sadeh, N. (2023). ATLAS: Automatically Detecting Discrepancies Between Privacy Policies and Privacy Labels. arXiv:2306.09247.
- Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer Privacy Choice in Online Advertising: Who Opt's Out and at What Cost to Industry?. *Marketing Science*, 39(1), 33-51.
- Khandelwal, R., Nayak, A., Chung, P., & Fawaz, K. (2023). The Overview of Privacy Labels and their Compatibility with Privacy Policies. arXiv:2303.08213v2 [cs.CY].
- Koch, S., Wessels, M., Altpeter, B., Olvermann, M., & Johns, M. (2022). Keeping Privacy Labels Honest. *Proceedings on Privacy Enhancing Technologies*, 2022(4), 486-506.
- Kollnig, K., Shuba, A., Van Kleek, M., Binns, R., & Shadbolt, N. (2022). Goodbye Tracking? Impact of iOS App Tracking Transparency and Privacy Labels. *ACM FAccT '22*, ACM, New York, NY, USA (pp. 508-520).
- Kraft, L., Bleier, A., Skiera, B., & Koschella, T. (2024). Granular Control and Privacy Decisions: Evidence from Apple's App Tracking Transparency (ATT). Available at SSRN 4598472.
- Lambrecht, A., & Tucker, C. (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research*, 50(5), 561-576.
- Reimers, I., & Waldfogel, J. (2023). A Framework for Detection, Measurement, and Welfare Analysis of Platform Bias. NBER Working Paper 31766.

Slemrod, J. (2019). Tax Compliance and Enforcement. *Journal of Economic Literature*, 57(4), 904-954.

Sun, T., Yuan, Z., Li, C., Zhang, K., & Xu, J. (2023). The Value of Personal Data in Internet Commerce: A High-Stakes Field Experiment on Data Regulation Policy. *Management Science*, 70(4), 2645–2660.

Wang, P., Jiang, L., & Yang, J. (2024). The Early Impact of GDPR Compliance on Display Advertising: The Case of an Ad Publisher. *Journal of Marketing Research*, 61(1), 70-91.

Xiao, Y., Li, Z., Qin, Y., Bai, X., Guan, J., Liao, X., & Xing, L. (2022). Lalaine: Measuring and Characterizing Non-Compliance of Apple Privacy Labels at Scale. arXiv:2206.06274v2 [cs.CR].

Appendix

A.1. Details on Privacy Disclosures

Process for Developers: Both Apple and Google rely on self-reporting by app developers. Their disclosures rest upon standardized privacy questionnaires and forms submitted through Apple’s App Store Connect and Google’s Play Console. They provide guidance on what to disclose, though app developers remain ultimately responsible for the accuracy of their disclosures. Moreover, app developers must update their privacy disclosures if their data practices change. As it is an honor system, platforms do not automatically verify privacy labels at the time of app submission, but they may do so later manually.

Enforcement and Non-Compliance: Given the honor system, enforcement is largely reactive – driven by external complaints and audits, e.g., from users, researchers, or media coverage. If data practices are misleading or false, potential sanctions include rejection of app updates or removal of the app, requirement to correct disclosures, developer account suspension or bans for repeated violations. Neither platform discloses specific methods for verifying compliance, though there are anecdotal reports of app removals on both.

A.2 Details on Data

Representativeness: We argue for better generalizability due to a cross-platform and cross-regional setup and wide coverage across developer regions and installations. While the sample comprises ad-based apps, advertising is the primary revenue source for most apps. We also ensure that no specific categories drive our results through subsample analyses.

Measurement Error: Non-compliance, our main measure, is based on claimed and actual behavior, and accurately measuring both is essential. We consider our disclosure variable conservative, as we treat any mention of a device ID in the privacy labels as a disclosure. Additionally, we provide robustness checks for different thresholds to determine when actual device ID transfers should not be considered a technical limitation, artifact, or inconsistency.

A.3. Additional Descriptive Statistics

Figures 8-10 provide more information on the underlying sample by showing distributions for app categories, installation numbers and developer locations corresponding to Sections 5.3. and 5.4. Table 3 provides summary statistics on a different aggregation level than Table 1 in Section 5.5. Finally, relating to Section 6.1., Figures 11-12 show graphs on the differentiation of app developers regarding disclosure and transfer, our two components to compute non-compliance, whereas Figure 13 delves deeper into characterizing non-compliance by platform and country of the user.

Figure 8: Distribution of the Apps' Categories




Figure 9: Distribution of the Apps' Number of Installations




Figure 10: Distribution of the Apps' Developer Location across the Top 10 Countries

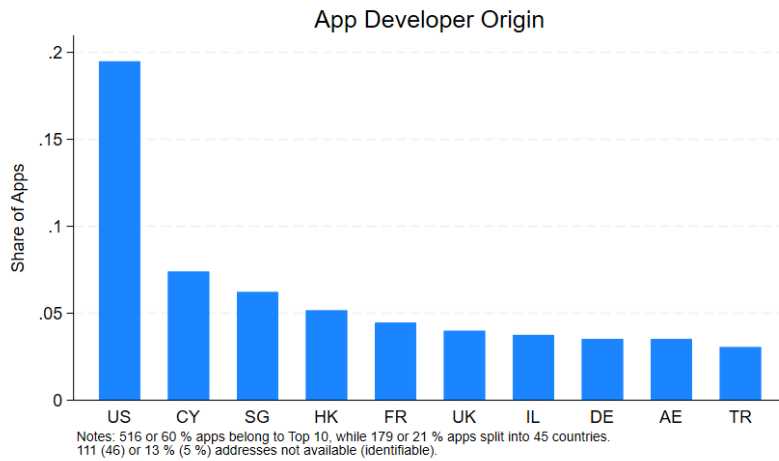



Table 3: Summary Statistics (Observation Level)


	Apple			Google		
	Mean	SD	N	Mean	SD	N
Number of Auctions (in millions)	8.18	109.23	94,978	9.44	203.72	105,574
Share of Trackable Traffic	0.31	0.31	94,978	0.78	0.33	105,574
Share of Actual Device ID Transfer	0.74	0.44	94,978	0.89	0.31	105,574
Share of Disclosed Device ID Transfer	0.88	0.32	94,978	0.73	0.44	105,574
Share of Non-Compliance	0.08	0.28	94,978	0.24	0.42	105,574
Share of EU Users	0.33	0.47	94,978	0.34	0.47	105,574
Share of EU Developers	0.28	0.45	94,978	0.29	0.46	105,574
Number of Ratings (in thousands)	38.63	165.51	93,074	1,273.89	3,455.84	104,447
Average Rating	4.53	0.31	92,456	4.30	0.37	104,447
App Age (in Years)	5.35	3.58	94,950	4.91	2.82	105,265
Large Dummy (10M+ Installations)	0.72	0.45	90,702	0.74	0.44	105,574
Share of In-App Prices	0.94	0.24	93,074	0.93	0.25	105,486
N	94,978			105,574		

Figure 11: Differentiation of Actual Transfer by Platform



Notes: The analysis is done at the app level. Each app on both platforms has observations for up to 19 countries multiplied by its connected supply-side platforms. For each app, we compute how many observations denote a transfer. For example, an app is available in 18 countries and connects to 6 supply-side platforms in each country. If we observed transfers in 3 countries on all of its connected supply-side platforms, then the share of observations with transfers would be $(3 \times 6) / (18 \times 6) = 16.67\%$. As a result, every app falls into the three categories: nowhere ($< 10\%$), somewhere (10-90%), or everywhere ($> 90\%$). Results on the right (left) panel correspond to the Google (Apple) sample.

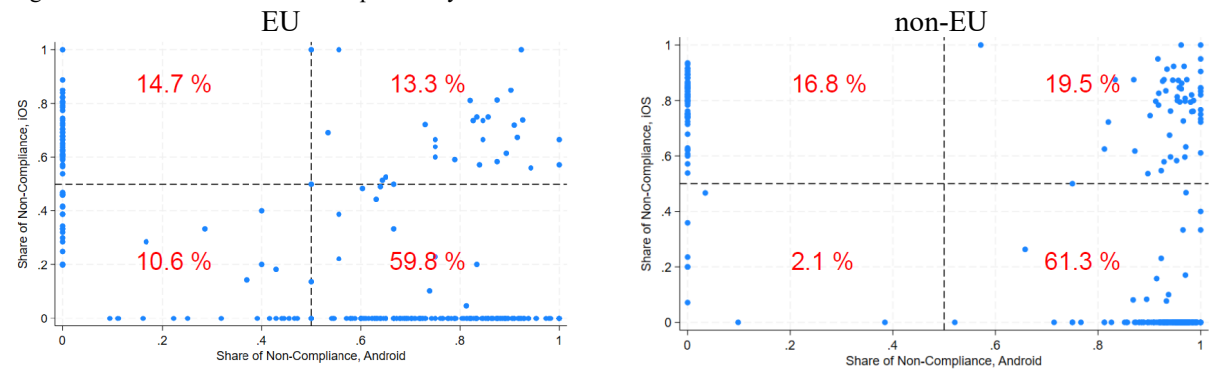
Figure 12: Differentiation of Disclosure by Platform



Notes: The analysis is done at the app level. Each app on both platforms has observations for up to 19 countries multiplied by its connected supply-side platforms. For each app, we compute how many observations denote disclosure. For example, an app is available in 18 countries and connects to 6 supply-side platforms in each country. If we observed disclosures in 3 countries on all of its connected supply-side platforms, then the share of observations with disclosures would be $(3 \times 6) / (18 \times 6) = 16.67\%$.

As a result, every app falls into the three categories: nowhere ($< 10\%$), somewhere ($10-90\%$), or everywhere ($> 90\%$). Results on the right (left) panel correspond to the Google (Apple) sample.

Figure 13: Differentiation Non-Compliance by Platform and User Location



Notes: The analysis is done at the app level. For each app, we compute how many observations denote non-compliance. As a result, every app has a share of non-compliant observations for each platform, with the x-axis denoting Google and the y-axis denoting Apple. Results on the right (left) panel correspond to the non-EU (EU) sample.

A.4. Additional Regression Analyses

Table 4 provides robustness checks of the baseline results, as shown in Section 6.2.1, by exploiting within-variation. Table 5 shows the regression results corresponding to Section 6.2.2. on using an alternative, less important, identifier for the platform of Apple. In contrast, Table 6 displays the regression results when studying variation over time with a two-period panel. Finally, Table 7 shows that our baseline regression results remain similar for a different aggregation level of the dataset, which we used for the calculations of Section 6.3.

Table 4: Robustness of Baseline Regression (exploiting within-variation)

	Non-Compliance Dummy {Yes = 1, No = 0}			
	(1)	(2)	(3)	(4)
EU User	-0.04*** (0.00)	-0.03*** (0.00)	-0.04*** (0.00)	-0.03*** (0.00)
Google	0.14*** (0.02)	0.14*** (0.02)		
Average Rating	0.02 (0.02)	0.02 (0.02)	0.00 (0.00)	0.00 (0.00)
Twin FE	Yes	Yes	No	No
App FE	No	No	Yes	Yes
SSPs FE	Yes	Yes	Yes	Yes
Mean Dep. Var.	0.16	0.15	0.16	0.15
R ²	0.51	0.48	0.84	0.86
N (Apps)	851	790	1636	1566
N (Total)	196903	152866	196897	152852

+ p<0.10, * p<0.05, ** p<0.01, *** p<0.001; SSP denotes Supply-Side Platform; FE denotes Fixed Effects.

Standard errors are clustered at the twin level. Reference categories are non-EU and Apple.

Columns 2 and 4 restrict the sample to having twin apps available in each platform, country, and SSP.

“Twin” denotes the developer of the two apps enabling variation within platforms, while “app” corresponds to the respective platform version.

Table 5: Regression Results for Alternative Identifier (IDFV)

	Non-Compliance Dummy {Yes = 1, No = 0}				
	(1)	(2)	(3)	(4)	(5)
EU User	-0.01*** (0.00)	-0.01 (0.00)	-0.01** (0.00)	-0.01** (0.00)	-0.01** (0.00)
EU Developer	-0.01 (0.02)	-0.00 (0.02)	-0.00 (0.02)	-0.00 (0.02)	0.00 (0.02)
Log. (# Ratings)		0.00 (0.00)			
Average Rating			-0.01 (0.02)		
App Age (in Years)				-0.00 (0.00)	
Large Dummy					-0.02 (0.02)
In-App Price Dummy		0.02 (0.03)	0.02 (0.03)	0.03 (0.03)	0.04 (0.03)
Categories	Yes	Yes	Yes	Yes	Yes
SSPs	Yes	Yes	Yes	Yes	Yes
Mean Dep. Var.	0.08	0.08	0.08	0.08	0.08
R ²	0.06	0.05	0.05	0.06	0.06
N (Apps)	836	835	835	835	791
N (Total)	94,978	92,456	92,456	93,046	88,856

+ p<0.10, * p<0.05, ** p<0.01, *** p<0.001; SSP denotes Supply-Side Platform; FE denotes Fixed Effects. Standard errors are clustered at the twin level. Reference categories are non-EU. Apple-only sample.

Table 6: Regression Results for Two-Period Panel

	Non-Compliance Dummy {Yes = 1, No = 0}			
	(1)	(2)	(3)	(4)
Apr '24 Dummy	0.00 (0.00)	0.01* (0.00)	0.00 (0.00)	0.01* (0.00)
EU User	-0.05*** (0.00)	-0.04*** (0.00)	-0.04*** (0.00)	-0.03*** (0.00)
Apr '24 Dummy x EU User		-0.03*** (0.00)		-0.03*** (0.00)
EU Developer	-0.02 (0.02)	-0.02 (0.02)		
Google	0.15*** (0.02)	0.15*** (0.02)	0.15*** (0.02)	0.15*** (0.02)
App Age (in Years)	-0.02*** (0.00)	-0.02*** (0.00)		
Average Rating			0.00 (0.02)	0.00 (0.02)
Categories FE	Yes	Yes	No	No
App FE	No	No	Yes	Yes
SSPs FE	Yes	Yes	Yes	Yes
Mean Dep. Var.	0.17	0.17	0.17	0.17
R ²	0.10	0.10	0.50	0.50
N (Apps)	850	850	849	849
N (Total)	300,172	300,172	297,179	297,179

+ p<0.10, * p<0.05, ** p<0.01, *** p<0.001; SSP denotes Supply-Side Platform; FE denotes Fixed Effects. Standard errors are clustered at the twin level. Reference categories are Nov '23, non-EU and Apple.

Table 7: Antecedents of Non-Compliance (Linear Probability Model, Aggregated over Supply-Side Platforms (SSPs))

	Non-Compliance Dummy {Yes = 1, No = 0}					
	(1)	(2)	(3)	(4)	(5)	(6)
EU User	-0.02*** (0.00)	-0.02*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
EU Developer	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)	-0.03 (0.02)
Google		0.17*** (0.02)	0.18*** (0.03)	0.16*** (0.02)	0.16*** (0.02)	0.16*** (0.02)
Log. (# Ratings)			-0.00 (0.00)			
Average Rating				-0.02 (0.02)		
App Age (in Years)					-0.02*** (0.00)	
Large Dummy (10M+)						-0.04+ (0.02)
In-App Price Dummy			0.00 (0.03)	0.00 (0.03)	0.01 (0.03)	0.01 (0.03)
Category FE	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.05	0.10	0.09	0.09	0.11	0.10
N (Apps)	852	852	851	851	852	807
N (Total)	28,049	28,049	26,783	26,783	27,364	26,720

+ p<0.10, * p<0.05, ** p<0.01, *** p<0.001; FE denotes Fixed Effects.

Standard errors are clustered at the twin level. Reference categories are non-EU and Apple.