Goodbye Tracking? Impact of iOS App Tracking Transparency and Privacy Labels

Konrad Kollnig
konrad.kollnig@cs.ox.ac.uk
Department of Computer Science,
University of Oxford
Oxford, United Kingdom

Anastasia Shuba
ashuba22@gmail.com
Independent Researcher
USA

Max Van Kleek
max.van.kleek@cs.ox.ac.uk
Department of Computer Science,
University of Oxford
Oxford, United Kingdom

Reuben Binns
reuben.binns@cs.ox.ac.uk
Department of Computer Science,
University of Oxford
Oxford, United Kingdom

Nigel Shadbolt
nigel.shadbolt@cs.ox.ac.uk
Department of Computer Science,
University of Oxford
Oxford, United Kingdom

ABSTRACT

Tracking is a highly privacy-invasive data collection practice that has been ubiquitous in mobile apps for many years due to its role in supporting advertising-based revenue models. In response, Apple introduced two significant changes with iOS 14: App Tracking Transparency (ATT), a mandatory opt-in system for enabling tracking on iOS, and Privacy Nutrition Labels, which disclose what kinds of data each app processes. So far, the impact of these changes on individual privacy and control has not been well understood. This paper addresses this gap by analysing two versions of 1,759 iOS apps from the UK App Store: one version from before iOS 14 and one that has been updated to comply with the new rules.

We find that Apple's new policies, as promised, prevent the collection of the Identifier for Advertisers (IDFA), an identifier for cross-app tracking. Smaller data brokers that engage in invasive data practices will now face higher challenges in tracking users – a positive development for privacy. However, the number of tracking libraries has – on average – roughly stayed the same in the studied apps. Many apps still collect device information that can be used to track users at a group level (*cohort tracking*) or identify individuals probabilistically (*fingerprinting*). We find real-world evidence of apps computing and agreeing on a fingerprinting-derived identifier through the use of server-side code, thereby violating Apple's policies. We find that Apple itself engages in some forms of tracking and exempts invasive data practices like first-party tracking and credit scoring from its new tracking rules. We also find that the new Privacy Nutrition Labels are sometimes inaccurate and misleading, especially in less popular apps.

Overall, our observations suggest that, while Apple's changes make tracking individual users more difficult, they motivate a countermovement, and reinforce existing market power of gatekeeper companies with access to large troves of first-party data. Making the privacy properties of apps transparent through large-scale analysis remains a difficult target for independent researchers, and a key obstacle to meaningful, accountable and verifiable privacy protections.

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06…$15.00
https://doi.org/10.1145/3531146.3533116

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; *Economics of security and privacy*; • **Networks** → Mobile and wireless security.

KEYWORDS

mobile apps, Apple, iOS, data protection, privacy, platform policies, gatekeeper power, App Tracking Transparency, Privacy Nutrition Labels

ACM Reference Format:
Konrad Kollnig, Anastasia Shuba, Max Van Kleek, Reuben Binns, and Nigel Shadbolt. 2022. Goodbye Tracking? Impact of iOS App Tracking Transparency and Privacy Labels. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3531146.3533116

1 INTRODUCTION

Tracking, the large-scale collection of data about user behaviour, is commonplace across both mobile app ecosystems, Android and iOS. While some see tracking as a 'necessary evil' to making apps available at lower prices by showing users personalised advertising and selling their data to third parties [13, 39], tracking can have highly disproportionate effects on the lives of individuals and society as a whole [49, 56]. As a countermeasure, Apple introduced the *Apple Tracking Transparency* (ATT) framework – alongside mandatory *Privacy Nutrition Labels* [29, 30] – with iOS 14, see Figure 1.

The emergence of more robust privacy measures in everyday technology is partly motivated by new data protection and privacy laws around the globe, particularly the General Data Protection Regulation (GDPR) in the EU and UK since May 2018 [34]. Among other aspects, the GDPR protects any data that can be related to individuals ('personal data'), and requires a legal basis for any processing of such personal data. This requirement has the effect that app tracking, which usually classifies as 'high-risk' data processing,

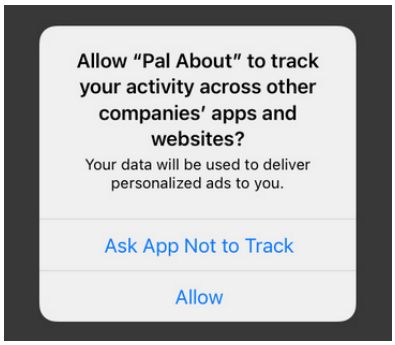

(a) Apple Tracking Transparency (ATT).

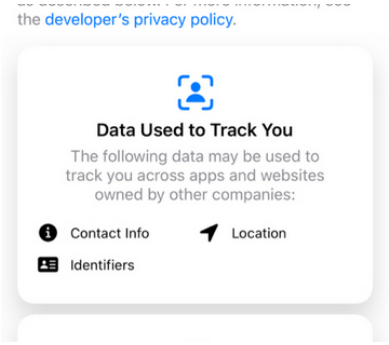

(b) Privacy Nutrition Label.

Figure 1: Overview of Apple's new privacy measures, introduced with iOS 14 [3].

needs prior user consent [33, 43]. Additionally, the 2009 ePrivacy Directive, which regulates data processing in electronic systems in the EU and UK, also requires consent to tracking [11, 33]. Despite these legal requirements, a large proportion of apps engaging in tracking have in the past been observed not to seek the required prior user consent [33, 43, 48].

Starting with iOS 14.5 in April 2021, iOS apps must now ask users for explicit permission before tracking them, see Figure 1a. If an iOS user asks an app not to track, this has the direct effect that this app cannot access the Identifier for Advertisers (IDFA) anymore. The IDFA is a random, unique identifier provided by the operating system to apps for tracking users across multiple sessions of a single app and across apps. Additionally, apps are obliged to stop certain tracking practices under the Apple's App Store policies (more in Section 5). Preliminary data suggests that the vast majority of users (between 60% and 95%) choose to refuse tracking when asked for it under the new system [4, 21, 27].

While potentially good for user privacy, the ATT has been reported to have vastly increased Apple's share of advertising on iOS – as part of its Apple Search Ads on the App Store – and to have decreased the efficacy of ads from competing companies. An important reason for this, as argued by Eric Seufert and others, is that Apple's own tracking technologies may not fall under Apple's definition of tracking [40]. It has also been reported that many marketing companies have shifted advertising budgets from iOS to Android [18]. The Financial Times estimated that the loss for leading tech companies from the new policy would be around $10bn [20], but also reported that companies deemed the 'effect of Apple's privacy changes was less than feared' [16]. Apple's privacy changes may prompt a rise in paid apps and in-app purchases [31], and thereby particularly affect those individuals who are already worse off financially.

In addition to the changes relating to the ATT, app developers must now self-declare what types of data they collect from users, and for what purposes – called *Privacy Nutrition Labels* [29, 30], see Figure 1. As such, these labels aim to make it easier for end-users to understand the data practices of apps, instead of having to study lengthy privacy policies, which few users do [38]. There is, however, a risk that many users may just ignore the new (and potentially oversimplified) privacy labels (as they commonly do with privacy policies [38]), gain a false sense of security, or not understand the consequences for their privacy (which tends to be highly individual [44]), and that developers may not honestly self-declare their actual data practices [60]. Despite these concerns, the labels have the potential to shift developers' existing data practices towards being more privacy-preserving, through increased transparency and end-user awareness.

Based on the above observations, this paper analyses the following research questions:

(1) What impact have the ATT and Privacy Nutrition Labels had – thus far – on tracking, particularly on the extent and quality of tracking?
(2) To what extent do apps disclose their tracking practices in their Privacy Nutrition Labels?
(3) What implications do the ATT and Privacy Nutrition Labels have for the power relations between the actors in the digital advertising system, including mobile OS providers, digital advertisers, app developers and marketers?

To analyse these questions, this paper analyses privacy in 1,759 iOS apps, for each of which we downloaded two versions: one from before Apple's new rules and one that has been updated since. We use a combination of app code and network analysis to gain rich insights into the data practices of the studied apps.

The remainder of this paper is structured as follows. We first review related work in Section 2. Next, we introduce our app download and analysis methodology in Section 3. We turn to the results from our app code and network analysis in Section 4. We discuss our findings in Section 5 and the limitations of our study in Section 5.1. We conclude the paper and outline direction for future work in Section 6. Code and data to replicate our results are available at https://www.platformcontrol.org/.

2 BACKGROUND

2.1 Related work

Previous research extensively studied privacy in mobile apps. Two main methods have emerged in the academic literature: dynamic and static analysis.

Dynamic analysis observes the run-time behaviour of an app, to gather evidence of sensitive data leaving the device. Early research focused on OS instrumentation, i.e. modifying Android [14] or iOS [1]. With growing complexity of mobile operating systems, recent work has shifted to analysing network traffic [24, 46, 48, 50, 51, 51, 52, 56]. This comes with certain limitations. One problem is limited scalability, since every app is executed individually. Another issue is that not all privacy-relevant parts of apps may be invoked during analysis, potentially leading to incomplete results.

Static analysis dissects apps without execution. Usually, apps are decompiled, and the obtained program code is analysed [12, 26]. The key benefit of static analysis is that it can analyse apps quickly, allowing it to scale to millions of apps [6, 9, 34, 57, 59]. However, static analysis can involve substantial computational effort and – unlike dynamic analysis – does not allow the observation of real data flows because apps are never actually run. Programming techniques, such as the use of code obfuscation and native code, can pose further obstacles. This is especially true for iOS apps, which are often harder to analyse and decompile – compared to Android – and are encrypted by default [6, 35, 62]. While this iOS encryption might legitimately protect *paid* apps against piracy, Apple also encrypts all free apps downloaded from the App Store. By contrast, Google only encrypts paid apps (not free ones) when downloaded from its Play Store. The encryption of iOS apps by Apple – even of free ones – is problematic for research efforts because it drives researchers into legal grey areas of copyright law [35]. Partly because of these difficulties, our recent work [35] was the first large-scale app privacy analysis study on iOS apps since 2013 [1]. We avoided legal problems relating to copyright law by conducting part of the analysis on-device through using the popular app instrumentation tool `Frida` [22].

In this paper, we follow the methodology of our previous paper, which used a combination of both dynamic and static analysis, so as to compare the privacy practices of the studied apps before and after the introduction of Apple's new privacy rules. We discuss our methodology for this paper in more detail in Section 3.

2.2 Regulation of App Platforms

The centrality of app platforms – i.e. Apple's iOS and Google's Android ecosystem – makes them a target for effective privacy regulation, however such regulation is limited [54, 63]. The US Federal Trade Commission (FTC) established some baseline rules for app stores in 2013. They strongly recommended to app platforms to require just-in-time consent for sensitive data access, to seek privacy policies from app developers, and to implement system-wide opt-out mechanism from data collection [15]. Despite not being law, Google and Apple followed many of the recommendations, and have not seen further public recommendations from the FTC since.

In the EU and UK, there exists no targeted regulation of app stores. The Regulation on platform-to-business relations (P2BR) contains general provisions for online intermediaries, including app stores, but does little to enact better privacy protections [63]. Data protection laws, such as the GDPR and the ePrivacy Directive, arguably place the primary responsibility for data protection with the app developers, not usually with app platform providers – although this is subject to ongoing debate; this lack of data protection obligations within the entire software development process – not just deployment – has been widely criticised [8, 28].

While no targeted regulation exists, app platforms face increasing scrutiny by courts and regulators. In the case *Epic Games v Apple* running since 2020, a US District Court judge largely found no monopolistic behaviour of Apple, but did identify some anticompetitive conduct in Apple's business practices. The judge ordered Apple to allow app developers to inform app users of alternative payment methods. Both Apple and Epic Games have appealed the ruling. In the EU, following a complaint of Spotify against Apple from 2019, the European Commission identified multiple anticompetitive aspects about Apple's ecosystem in a preliminary ruling – the case is, however, still ongoing. In January 2022, the Dutch competition authority demanded changes from Apple to its App Store policies; Apple has to date not fulfilled the demands of the regulators in their entirety, and has instead chosen to pay a weekly penalty of €5 million up to a maximum of €50 million [5].

The challenges in keeping up with regulation of platforms have spurred a recent countermovement by lawmakers. In South Korea, parliament amended the Telecommunication Business Act to force app stores to allow alternative payment methods and reduce commissions [47]. In response, Apple lowered the share it takes from App Store revenues of small developers (making less than $1 million per year) from 30% to 15%. In the US, Congress is debating a new Open App Markets Act that aims to address common competition concerns around app stores and passed the Senate Judiciary Committee with a strong a 20—2 bipartisan vote in February 2022. In the EU, lawmakers are seeking to enact two new pieces of legislation that aim to improve the regulation of digital markets, the Digital Markets Act and the Digital Services Act. Any new legal requirement for app platforms will likely have implications worldwide, due to the nature of digital ecosystems.

In sum, there currently exist few specific legal obligations for app platforms. Instead, they are encouraged to self-regulate their conduct. The following analysis shall shine a light on how the recent policy changes by Apple, a highly prominent example of this self-regulation, have affected the actual privacy practices of mobile apps.

3 METHODOLOGY

In this section, we describe our analysis methodology (depicted in Figure 2), which follows the one that we previously used for a comparative analysis of iOS and Android apps' privacy practices [35]. Code and data to replicate our results are available at https://www.platformcontrol.org/. We therefore keep our description of the methodology short and refer the reader to the original paper for details.

3.1 App Selection and Download

This section details our process for selecting and downloading apps from the Apple App Store (step 1 in Figure 2). For the selection of apps, we revisited the same 12,000 iOS apps as in our previous study [35]. These apps were selected by first generating a large list of apps available on the Apple App Store between December

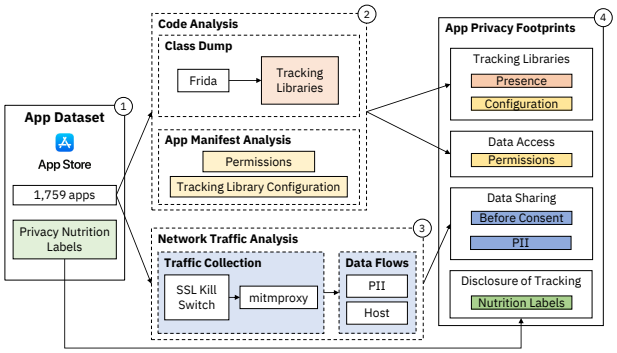

Figure 2: Overview of our analysis methodology (Section 3): First, (1) we select and download 1,759 apps from before the introduction of the ATT, and 1,759 from after. We also collect apps' Privacy Nutrition Labels. Next, we perform (2) Code Analysis to examine permissions and tracking libraries usage; and (3) Network Traffic Analysis to analyse tracking domains contacted at the first app start and the sharing of personal data. The results of this analysis (Section 4) are detailed App Privacy Footprints (4) of the downloaded apps.

2019 and February 2020. We then downloaded a random subset ($n = 12,000$) of those apps that were last updated since 2018 so as to focus on apps currently in use. For this work, we re-downloaded those apps that were updated to comply with Apple's ATT and privacy label rules, in October 2021. This resulted in a dataset of 1,759 *pairs* of apps, one from before iOS 14 and one from after. This number of apps is comparatively small because many apps had not yet been updated since the new rules, while some other apps had been removed from the store (2,713 out of 12,000 apps were not available on the App Store anymore). We additionally scraped the Privacy Nutrition Labels for the newly downloaded apps.

3.2 Code Analysis

To identify the presence of tracking libraries (step 2 in Figure 2), we extracted the names of all classes loaded by each app using the tool `Frida` [22] and checked them against a list of known tracker class names from our previous paper [35]. We also examined the app manifest (every iOS app must provide such a file) to determine how certain tracking libraries are configured – many tracking libraries allow developers to restrict data collection using settings in the manifest file, e.g. to disable the collection of unique identifiers or the automatic SDK initialisation at the first app start. This can help set up tracking libraries in a legally compliant manner. For example, 'Data minimisation' is one of the key principles of GDPR (Article 5.1 (c)), and user opt-in is required prior to app tracking in the EU and UK [33]. We analysed the privacy settings provided by some of the most prominent tracking libraries: Google AdMob, Facebook, and Google Firebase.

Beyond analysing tracking in apps, we also obtained a list of permissions that apps can request. Permissions form an important part of the security model of iOS as they protect sensitive information on the device, such as apps' access to the camera or address book. As such, permissions are different to the new privacy labels, which do not affect the runtime behaviour of apps. We extracted apps' permissions by automatically inspecting the manifest file.

3.3 Network Analysis

To analyse apps's network traffic (step 3 in Figure 2), we executed every app on a real device – one iPhone SE 1st Gen with iOS 14.2, and one with iOS 14.8 – for 30 seconds without user interaction. We captured network traffic using the tool `mitmdump`. We disabled certificate validation using `SSL Kill Switch 2`, after gaining system-level access on both iPhones (known as 'jailbreak'). On the iPhone with iOS 14.2, we did not opt-out from ad personalisation from the system settings, thereby assuming user opt-in to use the IDFA (reflecting the assumption that many users, who would reject tracking, do not do so because the option is in the less prominent settings on the OS [35]). On the iPhone with iOS 14.8, we asked all apps not to track from the system settings. Although in Android privacy research real user behaviour is simulated via various automation tools [7, 25, 45, 46, 48, 50, 55], Apple's restrictions on debugging and instrumentation have hindered the development of such tools for iOS. Tracking libraries are usually initialised at the first app start and without user consent [33, 35, 42, 48], and they can thus be detected without user interaction in the network traffic, as done in our analysis.

4 RESULTS

In this section, we present our findings from analysing two versions – one from before and one from after the release of iOS 14 and the ATT – of 1,759 iOS apps (step 4 in Figure 2). We analysed 199.6 GB of downloaded apps, extracted 3.2 GB in information about classes in apps' code, and collected 3.9 GB of data in apps' network traffic. Installing and instrumentation failed for 74 iOS apps; we have excluded these apps from our subsequent analysis and focus on the remaining 1,685 apps.

First, we focus on the tracking libraries found in the code analysis (Section 4.1) and whether or not they were configured for data minimisation (Section 4.1.1). Next, in Section 4.2, we analyse apps' access to the IDFA (which is now protected by the ATT) and also their permissions. Following up, in Section 4.3, we report on the data sharing of apps before consent is provided, with a particular focus on whether apps that are instructed not to track actually do so in practice. Lastly, in Section 4.4, we check whether and to what extent apps disclose their tracking practices in their Privacy Nutrition Labels.

4.1 Tracking Libraries

Apps from both before the ATT and after widely used tracking libraries (see Figure 3a). The median number of tracking libraries included in an app was 3 in both datasets. The mean before was 3.7, the mean after was 3.6. 4.75% of apps from before ATT contained more than 10 tracking libraries, compared to 4.75% after. 86.39% contained at least one before ATT, and 87.52% after.

The most prominent libraries have not changed since the introduction of ATT. The top one was the SKAdNetwork library (in 78.4% of apps before, and 81.8% after). While part of Apple's

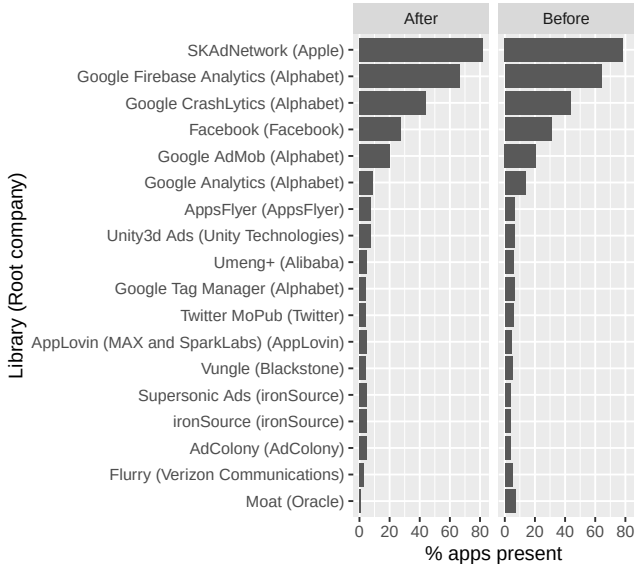

	Median	Mean	Q1	Q3	Count > 10	None
Before	3	3.7	2	5	4.75%	13.61%
After	3	3.6	2	4	4.75%	12.48%

(a) Top tracking libraries in app code.

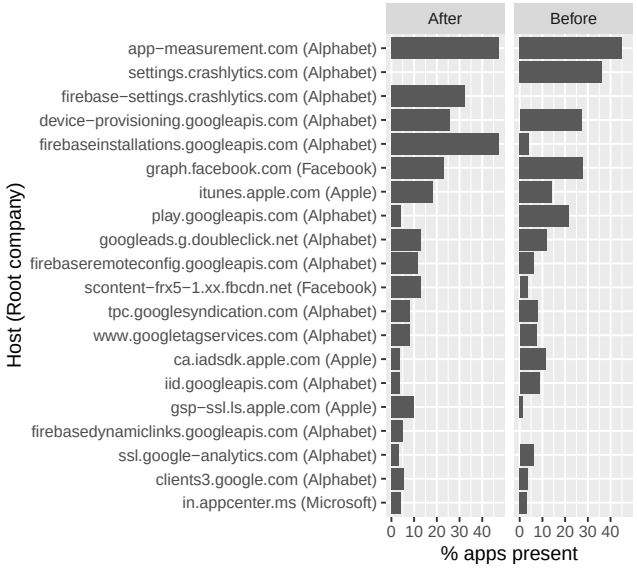

	Median	Mean	Q1	Q3	Count > 10	None
Before	3	4.0	1	6	4.75%	13.43%
After	4	4.7	2	7	7.19%	10.70%

(b) Top tracking hosts contacted at the first app start.

Figure 3: Third-party libraries (integrated in apps, but not necessarily activated) and contacted tracking domains of apps, as well as the companies owning them (in brackets). Shown are the top 15 tracking libraries and domains for before and after the new privacy changes under iOS 14.

privacy-preserving advertising attribution system, this library discloses information about what ads a user clicked on to Apple, from which Apple could (theoretically) build user profiles for its own advertising system. Following up with Apple about this potential issue (by one of the authors exercising the GDPR's *right to be informed* under Article 13), they did not deny the fact that this data might be used for advertising, but assured us that any targeted ads would only be served to segments of users (of at least 5,000 individuals with similar interests). Google Firebase Analytics ranked second (64.3% of apps from before ATT, and 67.0% after), and Google Crashlytics third (43.6% before, 44.4% after).

Overall, Apple's privacy measures seem not to have affected the integration of tracker libraries into *existing* apps.

4.1.1 Configuration for Data Minimisation. Among the apps that used Google AdMob, 2.9% of apps from before and 4.5% from after chose to delay data collection. Choosing to delay data collection can be helpful for app developers, to seek consent before enabling tracking and to fulfil legal obligations. Among the apps using the Facebook SDK, there was an increase in those which delayed the sending of app events (6.7% before, and 12.5% after); an increase in those which delayed the SDK initialisation (1.0% before ATT, 2.2% after), and an increase in those which disabled the collection of the IDFA (5.0% before, 8.6% after). Among apps using Google Firebase, 0.6% permanently deactivated analytics before ATT and 0.8% after, 0.0% disabled the collection of the IDFA before and 0.6% after, and 0.6% delayed the Firebase data collection before ATT and 1.0% after.

Overall, we found that only a small fraction of apps made use of data-minimising SDK settings in their manifest files. One reason for this observation might be that some developers are not aware of these settings because tracking companies tend to have an interest in less privacy-preserving defaults regarding data collection [33, 39]. This fraction has subtly increased since the introduction of the ATT.

4.2 Data Access and Permissions

Most prevalent permissions. Figure 4 shows the most prevalent permissions before and after the introduction of the ATT. On average, there was an increase in permission use (4.3 permissions before, 4.7 after – excluding the new *Tracking* permission). `CameraUsage` (for camera access) was the most common permission (62.6% before ATT, 66.9% after), closely followed by `PhotoLibraryUsage` (65.8% before ATT, 66.9% after), and `LocationWhenInUseUsage` (53.8% before ATT, 58.0% after).

Tracking permission and access to IDFA. As part of ATT, apps that want to access the IDFA or conduct tracking must declare the `TrackingUsage` permission in their manifest. 24.7% of apps from our dataset chose to declare this permission, and might ask users for tracking. At the same time, the share of apps that contain the `AdSupport` library, necessary to access the IDFA in the app code, stayed unchanged at 50.8% of apps. This means that 50.8% of

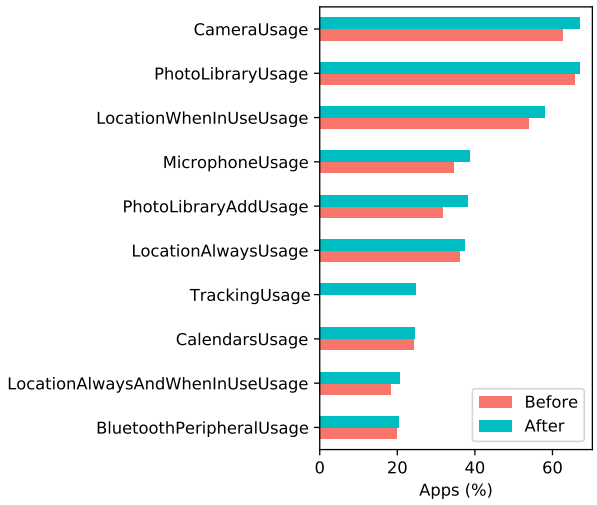

Figure 4: Top 10 permissions that apps can request.

apps from after the ATT could access the IDFA on earlier versions of iOS than 14.5, but only 24.7% can on iOS 14.5 or higher.

Tracking permission and integration of tracking SDKs. The share of apps that both contained a tracking library and could request tracking varied somewhat between the used tracking library. 69.3% of the 350 apps that integrated Google AdMob declared the `TrackingUsage` permission; 78.7% of the 110 apps that integrated Unity3d Ads; 50.0% of the 116 apps that integrated Moat; and 77.3% of the 54 apps that integrated Inmobi. Whether the app is from before or after the ATT, the vast majority of apps (between 97 and 100%) that integrated any of these tracking libraries also integrated the `AdSupport` library, and could therefore access the IDFA if running on iOS versions before 14.5.

4.3 Data Sharing

4.3.1 Before Consent. This section analyses how many tracking domains apps contacted before any user interaction has taken place; the next Section 4.3.2 then analyses what data was shared with trackers. Since tracking libraries usually start sending data right at the first app start [33, 35, 42, 48], this approach provides additional evidence as to the nature of tracking in apps – and without consent. Our results are shown in Figure 3b.

The average number of tracking domains contacted was somewhat higher for apps from after the introduction of the ATT (4.0 before, 4.7 after). The most popular domains were related to Google's analytics services: `firebaseinstallations.googleapis.com` (4.1% of apps before the ATT, 47.4% after) and `app-measurement.com` (45.2% before, 47.2% after). Since both endpoints are related to Google Firebase, the large increase in `firebaseinstallations.googleapis.com` prevalence likely reflects internal restructuring of Firebase following Google's acquisitions of other advertising and analytics companies. For example, Google acquired the crash reporting software Crashlytics from Twitter in January 2017, which is clearly reflected in our data. Google deprecated the old API endpoint (`settings.crashlytics.com` and changed it to `firebase-settings.crashlytics.com`) from November 2020. This had the direct effect that all Crashlytics users must now also use Google Firebase. The domain `settings.crashlytics.com` was contacted by 36.4% for apps from before the ATT, and `firebase-settings.crashlytics.com` by 32.3% after the ATT. While this might point to a small difference in the adoption of Google Crashlytics, the exact same number of apps (734, 43.6%) integrated the Crashlytics library into their code, before and after the ATT. Similarly, the exact same number of apps integrate the Facebook SDK (523, 31.1%); the share of apps that contacted the associated API endpoint `graph.facebook.com` at the first start fell from 27.7% to 23.1%. The Google Admob SDK, too, was integrated in the same number of apps (350, 20.8%), and did not see a decline in apps that contact the associated API endpoint `googleads.g.doubleclick.net` (12.1% before, 12.9% after).

Information	Example	Before	After
iPhone Name	MyPhone	2.5%	4.2%
iPhone Model	iPhone8,4|iPhone SE	60.2%	74.5%
Carrier	Three	20.2%	20.2%
Locale	en_GB|en-gb	85.7%	90.1%
CPU Architecture	ARM64|16777228	13.7%	16.1%
Board Config	N69uAP	3.1%	4.5%
OS Version	14.8|18H17	79.9%	86.9%
Timezone	Europe/London	3.9%	3.4%

Figure 5: Proportion of *all* apps that shared device information. This information can potentially be used for fingerprinting or cohort tracking.

Overall, data sharing with tracker companies before any user interaction remains common, even after the introduction of the ATT. This is in potential violation with applicable data protection and privacy laws in the EU and UK, which require prior consent [33].

4.3.2 Exposure of Personal Data. We found that 26.0% of apps from before the ATT shared the IDFA over the Internet, but none from after the ATT. In this sense, the ATT effectively prevents apps from accessing the IDFA. Despite Apple's promises, closer inspection of the network traffic showed that both Apple and other third parties are still able to engage in user tracking.

We found that iPhones continued to share a range of information with third-parties, that can potentially be used for device fingerprinting or cohort tracking, see Table 5. Only *timezone* saw a subtle decrease in the number of apps that shared this information. It is not clear why apps need to access or share some of this information, e.g. the carrier name (shared by 20.2% of apps) or the iPhone name (shared by 3–4% of apps). Meanwhile, some types of information, particularly the iPhone name, might allow the identification of individuals, especially when combined with other information.

In our analysis, we found 9 apps that were able to generate a mutual user identifier that can be used for cross-app tracking, through the use of server-side code. These 9 apps used an 'AAID' (potentially leaning on the term Android Advertising Identifier) implemented and generated by Umeng, a subsidiary of the Chinese tech company Alibaba. The flow to obtain an AAID is visualised in Figures 8a and 8b in the Appendix. As expected, the IDFA is

Domain	Company	Apps	User ID	Locale	Model	OS Version
firebaseinstallations.googleapis.com	Google	47.4%	✓	✓		
app-measurement.com	Google	47.2%	✓	✓		
firebase-settings.crashlytics.com	Google	32.3%	✓	✓	✓	✓
device-provisioning.googleapis.com	Google	25.8%	✓	✓	✓	✓
graph.facebook.com	Facebook	23.1%	✓	✓	✓	✓
itunes.apple.com	Apple	18.3%	✓	✓	✓	✓
fbcdn.net	Facebook	13.0%		✓		
googleads.g.doubleclick.net	Google	12.9%	✓	✓	✓	✓
firebaseremoteconfig.googleapis.com	Google	11.8%	✓	✓		
gsp-ssl.ls.apple.com	Apple	9.9%	✓	✓	✓	✓
tpc.googlesyndication.com	Google	8.3%		✓		✓
www.googletagservices.com	Google	8.1%		✓		✓
clients3.google.com	Google	5.3%		✓		
firebasedynamiclinks.googleapis.com	Google	5.2%	✓	✓		✓
in.appcenter.ms	Microsoft	4.3%	✓	✓	✓	✓
play.googleapis.com	Google	4.2%	✓	✓	✓	✓
skadsdk.appsflyer.com	AppsFlyer	4.0%	✓	✓		
gsp64-ssl.ls.apple.com	Apple	3.9%		✓	✓	✓
api.onesignal.com	OneSignal	3.7%		✓		
ca.iadsdk.apple.com	Apple	3.7%	✓	✓	✓	✓

Table 1: 20 most common tracking domains after ATT: sharing of user identifiers with third-parties, alongside device information. Empty cells mean that we did not observe the sharing of a certain type of information, although this might still take place.

only zeros because we used the opt-out provided by iOS 14.8; we observed, however, that the IDFV (ID for Vendors), a non-resettable, app-specific identifier was shared over the Internet, see Figure 8a. The sharing of device information for purposes of fingerprinting would be in violation of the Apple's policies, which do not allow developers to 'derive data from a device for the purpose of uniquely identifying it' [3]. Other experts and researchers have also voiced concerns that tracking might continue [19, 37, 41, 61].

We reported our observations to Apple on 17 November 2021, who promised to investigate the problem. We conducted a follow-up investigation on 1 February 2022, and re-downloaded and analysed a range of iOS apps. Some of the apps still continued to retrieve a unique identifier from the URL https://aaid.umeng.com/api/postZdata. Other apps now contacted the URL https://utoken.umeng.com/api/postZdata/v2, and applied additional encryption (rather than just HTTPS) to the requests and responses. This encrypted data had roughly the same size as before (~750 bytes for the request, ~350 bytes for the response) and the same mimetype (`application/json` for the request, `application/json;charset=UTF-8` for the response). The issue seems thus to be present still, but has now been hidden away from the public through the use of encryption. We have tried to reproduce these experiments for a few apps on iOS 15 and higher, but did not observe the same behaviour; there currently exists no public jailbreak for these iOS versions, and similar investigations as ours are therefore not (yet) possible on these iOS versions. There is a possibility that the issue has been fixed on iOS 15 or higher, or that we did not pick up the same behaviour in our small-scale testing (about 10 apps instead of more than 1000). However, Apple did not provide further details to us.

Analysing the top 20 most commonly contacted domains, we could confirm that installation-specific identifiers (see column 'User ID') are commonly collected alongside further device-specific information, see Table 1. While these installation-specific identifiers are usually randomly generated at the first app start, large tracking companies can likely still use these identifiers to build profiles of an app user's journey across apps, using their server-side code to link different identifiers together (e.g. through the user's IP address, other device information, and first-party data). Companies also receive information about a user's locale (i.e. the display language), the device model, and the OS version. Such information can be used to disambiguate different users connecting from the same IP address (e.g. households sharing the same Wi-Fi router) – and even across different IP addresses through the use of additional, first-party data that large tracking companies hold.

Table 1 does not include all the different kinds of information that we observed being sent to tracking domains because the kinds of information varied between companies. For example, Google assigned an `android_id` to an iOS app upon first contact with the company that was then used for all subsequent communication with Google's API endpoints. This identifier differed between apps, and did not seem to be used for cross-app tracking on-device (it might be on Google's servers). When contacting the domain `googleads.g.doubleclick.net`, Google collected the current system volume and the status of the silencing button. As already described above, `ca.iadsdk.apple.com` collected a `purchaseTimestamp`, that can be used to identify the user, and is not accessible for other app developers. The domain `gsp64-ssl.ls.apple.com`, belonging to Apple's location services, even collected the IP address and port that we used for proxying the network traffic through `mitmdump` as part of our analysis. We did not observe any other domains that had access to this information, underlining Apple's privileged data access. Crucially, for many of the observed transmissions between apps and servers, we could not even determine what data was sent, due to use of encryption [37] and closed-source communication protocols.

System-Level Tracking by Apple. We found that iPhones exchanged a range of unique user identifiers directly with Apple, see Figure 9 in the Appendix. We observed that network requests, which included various unique user identifiers and other personal

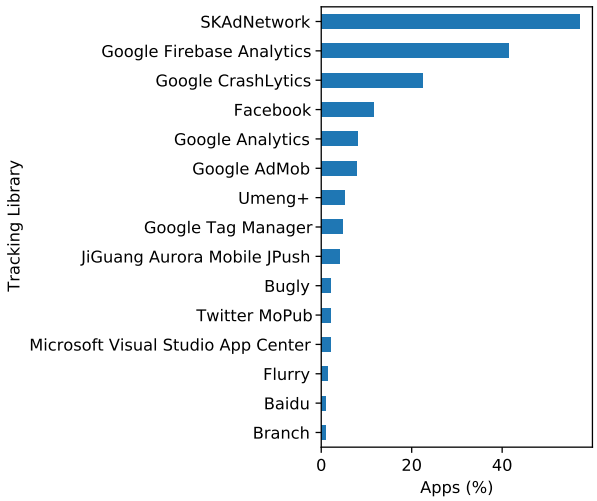

Figure 6: Top tracking libraries in apps that claim in their Privacy Nutrition Labels not to collect any data.

data, were issued following the interaction with apps and connected to Apple's App Store and advertising technologies. While this does not allow user-level apps to gain access to these user identifiers, Apple itself can use these identifiers to enrich its own advertising services. Indeed, Apple claims in its privacy policy that it may use users' interactions with its advertising platform and with the App Store to group users into segments (of at least 5,000 individuals), and show adverts to these groups [2]. Specifically, we found that the App Store collected the UDID, the serial number of the device, the DSID (an identifier linked to a user's Apple account), and a `purchaseTimestamp`. All of these identifiers can be used by Apple to single out individual users. Crucially, the UDID has been inaccessible to app developers other than Apple since 2013 [53], but Apple continues to have access to this identifier. Moreover, Apple collects the serial number, which cannot be changed and is linked to a user's iPhone. This might be unexpected for some users. These findings are in-line with previous reports that both Google and Apple collect detailed information about their users as part of regular device usage [36].

4.4 Disclosure of Tracking in Privacy Nutrition Labels

We now consider whether and to what extent apps (from after the introduction of iOS 14) disclose their tracking activities in their Privacy Nutrition Labels.

Among the studied apps, 22.2% claimed that they would not collect any data from the user. This was often not true: as shown in Figure 6, 80.2% of these apps actually contained at least one tracker library (compared to 93.1% for apps that did disclose some data sharing), and 68.6% sent data to at least one known tracking domain right at the first app start (compared to 91.4%). On average, apps that claimed not to collect data contained 1.8 tracking libraries (compared to 4.3), and contacted 2.5 tracking companies (compared to 4.2). Among the 22.2% of apps claiming not to collect data, only 3 were in the App Store charts. As noticed above (see Table 1), tracking libraries usually create a unique user identifier. Among the apps that used the SKAdNetwork, 42.0% disclosed their access to a 'User ID', 42.2% of apps using Google Firebase Analytics, 48.2% of apps using Google Crashlytics, and 53.2% of apps using the Facebook SDK. 63.2% of apps using Google Firebase Analytics disclosed that they collected any data about 'Product Interaction' or 'Other Usage Data', and about 70% of apps using the Facebook SDK, Google Analytics, or Google Tag Manager. Additionally, apps can disclose their use of 'Advertising Data': 27.5% of apps with the SKAdNetwork did so, 66.0% of apps with Google AdMob, 80.9% of apps with Unity3d Ads, and 45.4% apps with AppsFlyer.

All of this points to notable discrepancies between apps' disclosed and actual data practices. App developers might be able to address this, but are often not fully aware of all the data that is collected through third-party tracking software [13, 39]. Conversely, Apple itself might be able to reduce this discrepancy through increased use of automated code analysis, in particular applied to third-party tracking software.

5 DISCUSSION

Tracking continues, and reinforces the power of gatekeepers and opacity of the mobile data ecosystem. Our findings suggest that tracking companies, especially larger ones with access to large troves of first-party data, can still track users behind the scenes. They can do this through a range of methods, including using IP addresses to link installation-specific IDs across apps and through the sign-in functionality provided by individual apps (e.g. Google or Facebook sign-in, or email address). Especially in combination with further user and device characteristics, which our data confirmed are still widely collected by tracking companies, it would be possible to analyse user behaviour across apps and websites (i.e. fingerprinting and cohort tracking). A direct result of the ATT could therefore be that existing power imbalances in the digital tracking ecosystem get reinforced.

We even found a real-world example of Umeng, a subsidiary of the Chinese tech company Alibaba, using their server-side code to provide apps with a fingerprinting-derived cross-app identifier, see Figure 8 in the Appendix. The use of fingerprinting is in violation of Apple's policies [3], and raises questions around the extent to which Apple can enforce its policies against server-side code. ATT might ultimately encourage a shift of tracking technologies behind the scenes, so that they are outside of Apple's reach. In other words, Apple's new rules might lead to even less transparency around tracking than we currently have, including for academic researchers.

Privacy Nutrition Labels can be inaccurate and misleading, and have so far not changed data practices. Our results suggest that there is a discrepancy between apps' disclosed (in their Privacy Nutrition Labels) and actual data practices. We observed that many (mostly less popular) apps gave incomplete information or falsely declared not to collect any data at all. These observations are not necessarily to blame on app developers, who often have no idea of how third-party libraries handle users' personal data [13, 33, 39]. As reported in Section 4.1.1, the proportion of app developers that make use of data-minimising settings of popular

tracker libraries has roughly doubled, but these developers still remain a small minority. The Privacy Nutrition Labels have not (yet) had an impact on developers' actual practices at large, but might do so in the long run by both increasing app users' privacy expectations and making app developers rethink their privacy practices [29, 30]. As they stand, the labels can be misleading and create a false sense of security for consumers.

Are the most egregious and opaque trackers tamed now? The reduced access to permanent user identifiers through ATT could substantially improve app privacy. While in the short run, some companies might try to replace the IDFA with statistical identifiers, the reduced access to non-probabilistic cross-app identifiers might make it very hard for data brokers and other smaller tracker companies to compete. Techniques like fingerprinting and cohort tracking may end up not being competitive enough compared to more privacy-preserving, on-device solutions. We are already seeing a shift of the advertising industry towards the adoption of such solutions, driven by decisions of platform gatekeepers (e.g. Google's FloC / Topics API and Android Privacy Sandbox, Apple's ATT and Privacy Nutrition Labels) [17, 34], though more discussion is needed around the effectiveness of these privacy-protecting technologies. The net result, however, of this shift towards more privacy-preserving methods is likely going to be more concentration with the existing platform gatekeepers, as the early reports on the tripled marketing share of Apple [16], the planned overhaul of advertising technologies by Facebook/Meta and others [17], and the shifting spending patterns of advertisers suggest [18]. Advertising to iOS users – being some of the wealthiest individuals – will be an opportunity that many advertisers cannot miss out on, and so they will rely on the advertising technologies of the larger tech companies to continue targeting the right audiences with their ads.

Failure of GDPR enforcement, and power of platforms. Apple's new rules should not have a dramatic effect on the tracking of users in the EU and UK, given that existing data protection laws in these jurisdictions already ban most forms of third-party tracking without user consent [33, 43]. While there was vocal outcry over Apple's new privacy measures by advertisers, the adtech industry was aware of tightened EU and UK data protection rules since April 2016, and had plenty of time to work out a way to ensure compliance with basic provisions of the GDPR, until May 2018, including the need to seek consent from users before engaging in tracking [33]. Broad empirical evidence, from this and other pieces of research [32, 33, 35, 45, 48, 62], suggests that apps' compliance with the GDPR is somewhat limited.

At the same time, it is worrying that a few changes by a private company (Apple) seem to have changed data protection in apps more than many years of high-level discussion and efforts by regulators, policymakers and others. This highlights the relative power of these gatekeeper companies, and the failure of regulators thus far to enforce the GDPR adequately. An effective approach to increase compliance with data protection law and privacy protections in practice might be more targeted regulation of the gatekeepers of the app ecosystem; so far, there exists no targeted regulation in the US, UK and EU (see Section 2.2).

Apple's Double Standards I: Making and Enforcing App Store Policies. Our analysis shows that Apple has a competitive advantage within the iOS ecosystem in various ways. First, it both makes the rules for the App Store and interprets them in practice. This is reflected in Apple's definition of tracking, which ostensibly exempts its own advertising technology [2]: 'Tracking refers to the act of linking user or device *data collected from your app* with user or device data collected from *other companies' apps, websites, or offline properties* for *targeted advertising or advertising measurement purposes*. Tracking also refers to sharing user or device data with *data brokers*.' (emphasis added) [3] In other words, for tracking to fall under Apple's definition, it must fulfil three conditions, or be done by a data broker.

Apple's definition hinges on a distinction between first-party and third-party data collection, when this is not usually the root of privacy problems. This is why the W3C defines tracking as 'the collection of data regarding a particular user's activity across multiple distinct contexts and the retention, use, or sharing of data derived from that activity outside the context in which it occurred.' [58]. Rather than *companies*, this definition is centred around different *contexts*, as is commonly sought to be protected in privacy theory (e.g. contextual integrity [44]) and in privacy and data protection law (e.g. purpose limitation under Article 5 of the GDPR). Apple's definition of tracking might both betray the expectation of consumers who expect that tracking would stop (when first-party tracking, notably by Apple itself, continues to be allowed), and motivate other companies to consolidate and join forces leading to increased market concentration.

Apple additionally foresees a list of exempt practices [3] (see Figure 7 in the Appendix for an excerpt). These include 'fraud detection, fraud prevention, or security purposes', which might be interpreted extremely broadly by tracking companies. The exempt practices further allow tracking by a 'consumer reporting agency'. The term 'consumer reporting agency' is defined in the US Fair Credit Reporting Act (FCRA), regulating the relationship between these agencies and other 'furnishers of information' relating to consumers. By explicitly exempting credit scoring, Apple might try to avoid liability, and it might not have much choice under current US law. The exemption of credit scoring is nonetheless problematic because the use of personal data for credit scoring can have disproportionate impacts on individuals, and might be protected by other data protection and privacy laws. This might create the (false) impression for some app developers that other legal conditions do not apply, and a *false sense of security* for many consumers.

Apple's Double Standards II: Access to Data. Being the maker of the iOS ecosystem, Apple has a certain competitive advantage, by being able to collect device and user data, including hardware identifiers, that other app developers do not have access to, and use this for its own business purposes. For example, by collecting the device's serial number regularly, Apple can accurately tie the point-of-sale of its devices to activities on the device itself, and track the device lifecycle in great detail. Some of Apple's own apps, including the App Store itself, have access to this information because they are not distributed via the App Store and hence do not fall under the rules governing the App Store, including those that relate to tracking of users. These observations support the known concerns around fair competition in the App Store.

5.1 Limitations

A few limitations of our study are worth noting. First, for practical reasons, we were not able to analyse all the apps in the App Store, only a reasonably large subset of free apps in the App Store's UK region. Furthermore, for the purposes of examining the effect of ATT, we only focused on apps that already existed on the App Store before iOS 14 – newly released apps may adopt different strategies. Regarding our analysis methods, our instruments are also potentially limited in several ways. The results of our static analysis must be interpreted with care, since not all code shipped in an app will necessarily be invoked in practice. We may have overestimated tracking in certain contexts, e.g., if tracking code was included but not used. In our network analysis, we performed this off-device, meaning that all device traffic was analysed in aggregate. The risk here is that we may wrongly attribute some communications to an app that in fact was generated by some other app or subsystem on the device. To minimise this risk, we uninstalled all pre-installed apps, and ensured no apps were running in the background. We also used jailbreaking (i.e. gained full system access by exploiting a vulnerability in the iOS operating system) to circumvent certificate validation, which might make some apps alter their behaviour. In all parts of our analysis, we consider all apps equally, regardless of popularity [7] and usage time [55], both of which can impact user privacy. Likewise, we treat all tracking domains, libraries and companies equally, though they might pose different risks to users.

6 CONCLUSIONS & FUTURE WORK

Overall, we find that Apple's new policies largely live up to its promises on making tracking more difficult. Tracking libraries cannot access the IDFA anymore, and this directly impacts the business of data brokers. These data brokers can pose significant risks to individuals, since they try to amass data about individuals from a wide range of contexts and sell this information to third-parties. At the same time, apps still widely use tracking technology of large companies, and send a range of user and device characteristics over the Internet for the purposes of cohort tracking and user fingerprinting. We found real-world evidence of apps computing a mutual fingerprinting-derived identifier through the use of server-side code (see Section 4.3.2 and Figure 8 in the Appendix) – a violation of Apple's new policies [3], highlighting limits of Apple's enforcement power as a privately-owned data protection regulator [23, 54]. Indeed, Apple itself engages in some forms of user tracking (see Section 4.3.2 and Figure 9) and exempts invasive data practices like first-party tracking and credit scoring from its definition of tracking. Lastly, we found the Privacy Nutrition Labels to be sometimes incomplete and inaccurate, especially in less popular apps (Section 4.4).

Apple's privacy changes have led to positive improvements for user privacy. However, we also found various aspects that are in conflict with Apple's marketing claims and might go against users' reasonable privacy expectations, e.g. that the new opt-in tracking prompts would stop all tracking, that the new Privacy Nutrition Labels would always be correct and be verified by Apple, or that Apple would be subject to the same restrictions to data access and privacy rules as other companies. There is a risk that individuals will develop even more resignation over the use of their data online if they are provided with with misleading or ineffective privacy solutions [10, 49]. This resignation could in the long run undermine privacy efforts and adversely affect fundamental rights, such as the rights to data protection and privacy.

Despite positive developments over the recent months and years, especially through initiatives by Apple, there is still some way to go for app privacy. Violations of various aspects of data protection and privacy laws remain widespread in apps [32, 33, 35, 45, 48, 62], while enforcement of existing data protection laws against such practices stays sporadic. Apple's privacy efforts are hampered by its closed-source philosophy on iOS and the opacity around the enforcement of its App Store review policies. To strengthen iOS privacy, Apple has already started to prevent IP-based tracking by routing traffic to trackers via its own servers when using the iOS browser ('Privacy Relay'). As a direct response to our findings, Apple could consider extending the Privacy Relay to tracking within apps, thereby making the tracking of users through their IP address more difficult [41]. However, this would also further extend Apple's reach over the iOS ecosystem and potentially allow the company to track users even more accurately.

More generally, the key decision makers with regards to privacy technologies must establish robust transparency and accountability measures that allow for independent assessment of any privacy guarantees and promises. This is especially true, given the current lack of targeted regulations for app platforms like Google Play and the Apple App Store (see Section 2.2). In the case of Apple, improved transparency measures must necessarily involve the phasing out of encryption of free iOS apps by default, which currently forces independent privacy researchers into legal grey areas and severely hampers such research efforts (see Section 2.1). This is why most previous privacy research focused on Android and the last large-scale privacy study into iOS apps had been conducted in 2013 [1], until the recent release of the method used in this study [35].

We conclude that the new changes by Apple have traded more privacy for more concentration of data collection with fewer tech companies. Stricter privacy rules may encourage even less transparency around app tracking, by shifting tracking code onto the servers of dominant tracking companies. Despite the new rules, large companies, like Google/Alphabet and Facebook/Meta, are still able to track users across apps, because these companies have access to unique amounts of first-party data about users. Apple is now able to track its customers even more accurately, by taking a larger share in advertising technologies and getting unique access to user identifiers, including the device serial number. This underlines that privacy and competition problems can be highly intertwined in digital markets and need holistic study.

Future work. In this work, we only analysed apps that were already present on the App Store before iOS 14 and the ATT; it would be interesting to analyse how the ATT has impacted the privacy properties of *newly released* apps on the App Store. It would also be helpful to develop a new automation tool for iOS apps to observe apps' data practices automatically, even beyond the first app start – as studied in this paper. It would be pertinent to study user tracking by platforms in more detail, and also how the Privacy Nutrition Labels inform individuals around app privacy.

ACKNOWLEDGMENTS

We thank Jake Stein and Alexander Fanta for helpful comments and Ulrik Lyngs for help with data analysis. Konrad Kollnig was funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/R513295/1. Max Van Kleek has been supported by the PETRAS National Centre of Excellence for IoT Systems Cybersecurity, which has been funded by the UK EPSRC under grant number EP/S035362/1. Max Van Kleek, Reuben Binns, and Nigel Shadbolt have been supported by the Oxford Martin School EWADA Programme.

REFERENCES

[1] Yuvraj Agarwal and Malcolm Hall. 2013. ProtectMyPrivacy: Detecting and Mitigating Privacy Leaks on iOS Devices Using Crowdsourcing. In *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services - MobiSys '13*. ACM Press, Taipei, Taiwan, 97. https://doi.org/10.1145/2462456.2464460

[2] Apple. 2021. Apple Advertising & Privacy. https://www.apple.com/legal/privacy/data/en/apple-advertising/.

[3] Apple. 2021. User Privacy and Data Use. https://developer.apple.com/app-store/user-privacy-and-data-use/.

[4] AppsFlyer. 2021. Initial data indicates ATT opt-in rates are much higher than anticipated — at least 41%. https://www.appsflyer.com/blog/trends-insights/att-opt-in-rates-higher/.

[5] Authority for Consumers and Markets. 2022. ACM obliges Apple to adjust unreasonable conditions for its App Store. https://www.acm.nl/en/publications/acm-obliges-apple-adjust-unreasonable-conditions-its-app-store.

[6] Reuben Binns, Ulrik Lyngs, Max Van Kleek, Jun Zhao, Timothy Libert, and Nigel Shadbolt. 2018. Third Party Tracking in the Mobile Ecosystem. In *Proceedings of the 10th ACM Conference on Web Science - WebSci '18* (Amsterdam, Netherlands). ACM Press, New York, NY, USA, 23–31. https://doi.org/10.1145/3201064.3201089

[7] Reuben Binns, Jun Zhao, Max Van Kleek, and Nigel Shadbolt. 2018. Measuring Third-party Tracker Power across Web and Mobile. *ACM Transactions on Internet Technology* 18, 4 (2018), 1–22. https://doi.org/10.1145/3176246

[8] Lee A Bygrave. 2017. Data Protection by Design and by Default: Deciphering the EU's Legislative Requirements. *Oslo Law Review* 1 (2017), 105–120. https://doi.org/10.18261/issn.2387-3299-2017-02-03

[9] Kai Chen, Xueqiang Wang, Yi Chen, Peng Wang, Yeonjoon Lee, XiaoFeng Wang, Bin Ma, Aohui Wang, Yingjun Zhang, and Wei Zou. 2016. Following Devil's Footprints: Cross-Platform Analysis of Potentially Harmful Libraries on Android and iOS. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Jose, CA, 357–376. https://doi.org/10.1109/SP.2016.29

[10] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–13. https://doi.org/10.1145/3313831.3376389

[11] Datenschutzkonferenz. 2021. Orientierungshilfe Der Aufsichtsbehörden Für Anbieter von Telemedien.

[12] Manuel Egele, Christopher Kruegel, Engin Kirda, and Giovanni Vigna. 2011. PiOS: Detecting Privacy Leaks in iOS Applications. In *Proceedings of the Network and Distributed System Security Symposium (NDSS) 2011*. The Internet Society, San Diego, California, 15 pages.

[13] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2021. "Money makes the world go around": Identifying Barriers to Better Privacy in Children's Apps From Developers' Perspectives. In *Conference on Human Factors in Computing Systems (CHI '21)* (Yokohama, Japan, 2021). ACM Press, NY, USA, 1–24. https://doi.org/10.1145/3411764.3445599

[14] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. 2010. TaintDroid: An Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation (OSDI'10)*. USENIX Association, Vancouver, BC, 393–407.

[15] Federal Trade Commission. 2013. Mobile Privacy Disclosures–Building Trust Through Transparency. https://www.ftc.gov/sites/default/files/documents/reports/mobile-privacy-disclosures-building-trust-through-transparency-federal-trade-commission-staff-report/130201mobileprivacyreport.pdf.

[16] Financial Times. 2021. Alphabet and Microsoft smash estimates with $110bn revenue haul. https://www.ft.com/content/273aeecb-57a8-40f8-a2ba-8a21a635b289.

[17] Financial Times. 2021. Apple reaches quiet truce over iPhone privacy changes. https://www.ft.com/content/69396795-f6e1-4624-95d8-121e4e5d7839.

[18] Financial Times. 2021. Apple's privacy changes create windfall for its own advertising business. https://www.ft.com/content/074b881f-a931-4986-888e-2ac53e286b9d.

[19] Financial Times. 2021. China's tech giants test way around Apple's new privacy rules. https://www.ft.com/content/520ccdae-202f-45f9-a516-5cbe08361c34.

[20] Financial Times. 2021. Snap, Facebook, Twitter and YouTube lose nearly $10bn after iPhone privacy changes. https://www.ft.com/content/4c19e387-ee1a-41d8-8dd2-bc6c302ee58e.

[21] Flurry. 2021. iOS 14.5 Opt-in Rate - Daily Updates Since Launch. https://www.flurry.com/blog/ios-14-5-opt-in-rate-att-restricted-app-tracking-transparency-worldwide-us-daily-latest-update/.

[22] Frida. [n. d.]. Frida: A world-class dynamic instrumentation framework. https://frida.re.

[23] Daniel Greene and Katie Shilton. 2018. Platform privacies: Governance, collaboration, and the different meanings of "privacy" in iOS and Android development. *New Media & Society* 20, 4 (2018), 1640–1657. https://doi.org/10.1177/1461444817702397

[24] Catherine Han, Irwin Reyes, Amit Elazari, Joel Reardon, Alvaro Feal, Kenneth A. Bamberger, Serge Egelman, and Narseo Vallina-Rodriguez. 2019. Do You Get What You Pay For? Comparing The Privacy Behaviors of Free vs. Paid Apps.. In *The Workshop on Technology and Consumer Protection (ConPro '19)*. Institute of Electrical and Electronics Engineers, NY, USA, 7 pages.

[25] Catherine Han, Irwin Reyes, Álvaro Feal, Joel Reardon, Primal Wijesekera, Narseo Vallina-Rodriguez, Amit Elazari, Kenneth A. Bamberger, and Serge Egelman. 2020. The Price is (Not) Right: Comparing Privacy in Free and Paid Apps. *Proceedings on Privacy Enhancing Technologies* 2020, 3 (2020), 222–242. https://doi.org/10.2478/popets-2020-0050

[26] Jin Han, Qiang Yan, Debin Gao, Jianying Zhou, and Robert H Deng. 2013. Comparing Mobile Privacy Protection through Cross-Platform Applications. In *Proceedings 2013 Network and Distributed System Security Symposium* (San Diego, CA). Internet Society, 16.

[27] International Association of Privacy Professionals. 2021. Apple's ATT rollout presents uncertain path for adtech. https://iapp.org/news/a/apples-att-rollout-presents-uncertain-path-for-adtech/.

[28] Lina Jasmontaite, Irene Kamara, Gabriela Zanfir-Fortuna, and S Leucci. 2018. Data Protection by Design and by Default: Framing Guiding Principles into Legal Obligations in the GDPR. *European Data Protection Law Review* 4 (2018), 168–189. https://doi.org/10.21552/edpl/2018/2/7

[29] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "Nutrition Label" for Privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09* (Mountain View, California, 2009). ACM Press, 1. https://doi.org/10.1145/1572532.1572538

[30] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. 2010. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1573–1582. https://doi.org/10.1145/1753326.1753561

[31] Reinhold Kesler. 2022. The Impact of Apple's App Tracking Transparency on App Monetization. *Work in Progress* (2022), 22 pages.

[32] Konrad Kollnig. 2019. Tracking in Apps' Privacy Policies. *arXiv preprint arXiv:2111.07860* (2019), 10 pages. arXiv:2111.07860 [cs] http://arxiv.org/abs/2111.07860

[33] Konrad Kollnig, Reuben Binns, Pierre Dewitte, Max Van Kleek, Ge Wang, Daniel Omeiza, Helena Webb, and Nigel Shadbolt. 2021. A Fait Accompli? An Empirical Study into the Absence of Consent to Third-Party Tracking in Android Apps. *Proceedings of the Seventeenth Symposium on Usable Privacy and Security* (2021).

[34] Konrad Kollnig, Reuben Binns, Max Van Kleek, Ulrik Lyngs, Jun Zhao, Claudine Tinsman, and Nigel Shadbolt. 2021. Before and after GDPR: Tracking in Mobile Apps. 10, 4 (2021), 30 pages. https://doi.org/10.14763/2021.4.1611

[35] Konrad Kollnig, Anastasia Shuba, Reuben Binns, Max Van Kleek, and Nigel Shadbolt. 2022. Are iPhones Really Better for Privacy? A Comparative Study of iOS and Android Apps. *Proceedings on Privacy Enhancing Technologies* 2022, 2 (2022), 6–24. https://doi.org/10.2478/popets-2022-0033

[36] Douglas J Leith. 2021. Mobile Handset Privacy: Measuring The Data iOS and Android Send to Apple And Google. (2021), 10.

[37] Lockdown Privacy. 2021. Study: Effectiveness of Apple's App Tracking Transparency. https://blog.lockdownprivacy.com/2021/09/22/study-effectiveness-of-apples-app-tracking-transparency.html.

[38] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The Cost of Reading Privacy Policies. *I/S: A Journal of Law and Policy for the Information Society* (2008), 26.

[39] Abraham H Mhaidli, Yixin Zou, and Florian Schaub. 2019. "We Can't Live Without Them!" App Developers' Adoption of Ad Networks and Their Considerations of Consumer Risks. *Proceedings of the Fifteenth Symposium on Usable Privacy and Security* (2019), 21.

[40] Mobile Dev Memo. 2021. ATT advantages Apple's ad network. Here's how to fix that. https://mobiledevmemo.com/att-advantages-apples-ad-network-heres-how-to-fix-that/.

[41] Mobile Dev Memo. 2021. Why isn't Apple policing mobile ads fingerprinting? https://mobiledevmemo.com/why-isnt-apple-policing-mobile-ads-

fingerprinting/.
[42] Trung Tin Nguyen, Michael Backes, Ninja Marnau, and Ben Stock. 2021. Share First, Ask Later (or Never?) Studying Violations of GDPR's Explicit Consent in Android Apps. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 3667–3684. https://www.usenix.org/conference/usenixsecurity21/presentation/nguyen
[43] Trung Tin Nguyen, Michael Backes, Ninja Marnau, and Ben Stock. 2021. Share First, Ask Later (or Never?) Studying Violations of GDPR's Explicit Consent in Android Apps. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 3667–3684. https://www.usenix.org/conference/usenixsecurity21/presentation/nguyen
[44] Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Washington Law Review* 79 (2004), 39.
[45] Ehimare Okoyomon, Nikita Samarin, Primal Wijesekera, Amit Elazari, Narseo Vallina-Rodriguez, Irwin Reyes, Alvaro Feal, and Serge Egelman. 2019. On The Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies. *The Workshop on Technology and Consumer Protection (ConPro '19)* (2019), 7 pages.
[46] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. ReCon: Revealing and Controlling PII Leaks in Mobile Network Traffic. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services - MobiSys '16*. ACM Press, Singapore, Singapore, 361–374. https://doi.org/10.1145/2906388.2906392
[47] Reuters. Reteuers. S.Korea targets Apple over new app store regulation. https://www.reuters.com/technology/skorea-targets-apple-over-new-app-store-regulation-2021-10-15/.
[48] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. 2018. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies* 2018, 3 (2018), 63–83. https://doi.org/10.1515/popets-2018-0021
[49] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and Creepiness in App Space: Perceptions of Privacy and Mobile App Use. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems - CHI '14* (Toronto, Ontario, Canada). ACM Press, 2347–2356. https://doi.org/10.1145/2556288.2557421
[50] Anastasia Shuba and Athina Markopoulou. 2020. NoMoATS: Towards Automatic Detection of Mobile Tracking. *Proceedings on Privacy Enhancing Technologies* 2020, 2 (2020), 45–66. https://doi.org/10.2478/popets-2020-0017
[51] Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. 2018. NoMoAds: Effective and Efficient Cross-App Mobile Ad-Blocking. In *Proceedings on Privacy Enhancing Technologies 2018*. 125–140. https://doi.org/10.1515/popets-2018-0035
[52] Yihang Song and Urs Hengartner. 2015. PrivacyGuard: A VPN-based Platform to Detect Information Leakage on Android Devices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM '15)*. 15–26. https://doi.org/10.1145/2808117.2808120
[53] The Verge. 2013. Apple to finally stop accepting apps that use outdated UDID device identifier on May 1st. https://www.theverge.com/2013/3/21/4133288/apple-to-finally-stop-accepting-apps-that-use-outdated-udid-device-identifier-may-1st.
[54] Joris van Hoboken and R O Fathaigh. 2021. Smartphone platforms as privacy regulators. *Computer Law & Security Review* 41 (2021), 105557. https://doi.org/10.1016/j.clsr.2021.105557
[55] Max Van Kleek, Reuben Binns, Jun Zhao, Adam Slack, Sauyon Lee, Dean Ottewell, and Nigel Shadbolt. 2018. X-Ray Refine: Supporting the Exploration and Refinement of Information Exposure Resulting from Smartphone Apps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18* (Montreal QC, Canada). ACM Press, 1–13. https://doi.org/10.1145/3173574.3173967
[56] Max Van Kleek, Ilaria Liccardi, Reuben Binns, Jun Zhao, Daniel J. Weitzner, and Nigel Shadbolt. 2017. Better the Devil You Know: Exposing the Data Sharing Practices of Smartphone Apps. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. ACM Press, 5208–5220. https://doi.org/10.1145/3025453.3025556
[57] Nicolas Viennot, Edward Garcia, and Jason Nieh. 2014. A Measurement Study of Google Play. In *The 2014 ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '14)*. 221–233. https://doi.org/10.1145/2591971.2592003
[58] W3C Working Group. 2019. Tracking Compliance and Scope. https://www.w3.org/TR/tracking-compliance/#tracking.
[59] Haoyu Wang, Zhe Liu, Jingyue Liang, Narseo Vallina-Rodriguez, Yao Guo, Li Li, Juan Tapiador, Jingcun Cao, and Guoai Xu. 2018. Beyond Google Play: A Large-Scale Comparative Study of Chinese Android App Markets. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. 293–307. https://doi.org/10.1145/3278532.3278558
[60] Washington Post. 2021. I checked Apple's new privacy 'nutrition labels.' Many were false. https://www.washingtonpost.com/technology/2021/01/29/apple-privacy-nutrition-label/.
[61] Washington Post. 2021. When you 'Ask app not to track,' some iPhone apps keep snooping anyway. https://www.washingtonpost.com/technology/2021/09/23/iphone-tracking/.
[62] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N. Cameron Russell, and Norman Sadeh. 2019. MAPS: Scaling Privacy Compliance Analysis to a Million Apps. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (2019), 66–86. https://doi.org/10.2478/popets-2019-0037
[63] R. Ó Fathaigh and J. van Hoboken. 2019. European Regulation of Smartphone Ecosystems. *European Data Protection Law Review* 5, 4 (2019), 476–491. https://doi.org/10.21552/edpl/2019/4/6

APPENDIX

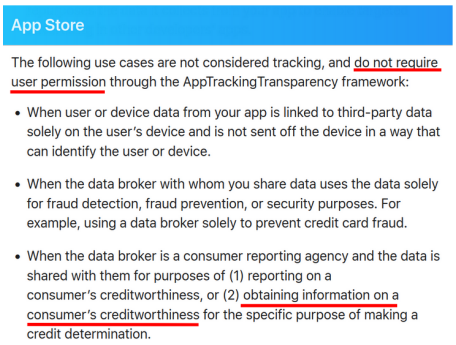

Figure 7: Apple's definition of tracking: Excerpt from Apple's exempt data practices, including credit scoring, from requiring user opt-in under ATT (emphasis added) [3]. We discuss the limitations of Apple's definition of tracking in Section 5.


```
{
    "sdk_version": "1.2.0",
    "bundle_id": "[Redacted]",
    "hw_model": "N69uAP",
    "kid": "[Redacted]",
    "total_storage": "30745123781",
    "country": "GB",
    "zdata": "[Redacted]",
    "app_version": "[Redacted]",
    "app_name": "[Redacted]",
    "sdk_type": "IOS",
    "storage": "14078912372",
    "zdata_ver": "1.1.0",
    "source_id": "umeng",
    "idfv": "7EBDAFC8-97BB-4FDB-B4D3-E2F4EA040B8C",
    "timezone": "1",
    "os_version": "14.8",
    "model": "iPhone8,4",
    "hostname": "MyPhone",
    "appkey": "[Redacted]",
    "idfa": "00000000-0000-0000-0000-
        000000000000"
}
```


(a) Request: Sending a range of device information to Umeng at https://aaid.umeng.com/api/postZdata.


```
{
    "aaid": "BAEC362C-49FC-494B-B0A7-175D990B059D",
    ...
}
```


(b) Response: Umeng returns an identifier that is shared by multiple apps, and can be used for cross-app tracking.

Figure 8: Fingerprinting in apps, even after the ATT. This is likely in violation of Apple's new policies and the expectations of many end-users (personal data redacted). We provide more results on the circumvention of the ATT in Section 4.3.2.

```
<plist version="1.0">
<dict>
        ...
        <key>dsid</key>
        <string>[Apple ID]</string>
        <key>guid</key>
        <string>[UDID]</string>
        <key>serialNumber</key>
        <string>[serial number]</string>
        ...
</dict>
</plist>
```

(a) Request of Apple App Store to https://buy.itunes.apple.com/WebObjects/MZFinance.woa/wa/renewVppReceipt?guid=[UDID].

```
{
    "attributionMetadataExistsOnDevice": false,
    "toroId": "[Redacted]",
    "purchaseTimestamp": "2021-11-01T15:15:05Z",
    "adamId": 477718890,
    "attributionDownloadType": 0,
    "developmentApp": false,
    "anonymousDemandId": "[Redacted]",
    "bundleId": "ru.kinopoisk",
    "attributionKey": "[Redacted]"
}
```

(b) Request (shortended) of Apple's advertising framework to https://ca.iadsdk.apple.com/adserver/attribution/v2.

Figure 9: Sharing of unique user identifiers with Apple (personal data redacted). We explain more about the tracking of users by Apple in Section 4.3.2.