

A Fait Accompli? An Empirical Study into the Absence of Consent to Third-Party Tracking in Android Apps

Konrad Kollnig
University of Oxford

Reuben Binns
University of Oxford

Pierre Dewitte
KU Leuven

Max Van Kleek
University of Oxford

Ge Wang
University of Oxford

Daniel Omeiza
University of Oxford

Helena Webb
University of Oxford

Nigel Shadbolt
University of Oxford

Abstract

Third-party tracking allows companies to collect users’ behavioural data and track their activity across digital devices. This can put deep insights into users’ private lives into the hands of many strangers, and often happens without users’ awareness or explicit consent. EU and UK data protection law, however, *requires consent*, both 1) to access and store information on users’ devices and 2) to legitimate the processing of personal data as part of third-party tracking, as we analyse in this paper.

This paper further investigates whether and to which extent consent is implemented in mobile apps. First, we analyse a representative sample of apps from the Google Play Store. We find that most apps engage in third-party tracking, but only 4% obtained consent before doing so, indicating potentially widespread violations of EU privacy law. Second, we examine the 13 most common third-party tracking libraries in detail. While most acknowledge that they rely on app developers to obtain consent on their behalf, they typically fail to put in place robust measures to ensure this: disclosure of consent requirements is limited; default consent implementations are lacking; and compliance guidance is difficult to find, hard to read, and poorly maintained.

1 Introduction

Third-party tracking, the deliberate collection, processing and sharing of users’ behavioural data with third-party companies, has become widespread across both mobile app ecosystems [15, 75, 76] and the web [15, 59]. The use of third-party

trackers benefits app developers in several ways, notably by providing analytics to improve user retention, and by enabling the placement of attributable personalised advertising within apps, which often translates into a vital source of revenue for them [28, 54]. However, it also makes app developers dependent on privacy-invasive data practices that involve the processing of large amounts of personal data [34, 41, 54], with little awareness from users and app developers [24, 63, 66, 76]. Data protection and privacy legislation such as the General Data Protection Regulation (GDPR) [32] in the EU and the UK, and the Children’s Online Privacy Protection Act (COPPA) [71] in the US, establish clear rules when it comes to the processing of personal data and provide additional safeguards when it comes to information relating to children. As explained in Section 3, consent is a necessary precondition for third-party tracking in mobile apps.

The implementation of consent in mobile apps has—since the end of 2020—sparked a fierce public battle between Apple and Facebook over tracking controls in iOS 14 [1, 2]. To give users more control over their privacy, Apple plans to introduce an opt-in mechanism from the use the Advertising Identifier (AdID)—similar to how apps currently request location or contacts access. Facebook, like many other mobile advertising companies, is concerned that most users will not consent to tracking if asked more clearly [3]; iOS users can already opt-out from the use of AdID, but are not explicitly asked by every app. Despite the importance of consent, there exists little empirical evidence as to whether mobile apps implement any type of consent mechanisms prior to engaging in third-party tracking.

Despite their crucial role within the software development life cycle, putting the blame of implementing consent incorrectly on app developers might be misguided. Many lack legal expertise, depend on the use of tracking software, and face limited negotiation power in the design of tracker software, which is usually developed by large, multinational companies [11, 20, 28, 42, 54]. At the same time, failure to implement appropriate consent mechanisms in software impacts individuals’ choice over data collection and their informational self-

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2021.
August 8–10, 2021, Vancouver, B.C., Canada.

determination, and may expose vulnerable groups—such as children—to disproportionate data collection. This underlines the need for robust privacy guarantees in code.

Driven by these observations, the present contribution aims to answer the following research questions:

1. Do app developers need to obtain valid user consent before engaging in third-party tracking in the EU and UK? (*consent requirements for tracking*)
2. To what extent do apps engage in third-party tracking, and obtain valid user consent before doing so? (*practices of app developers*)
3. To what extent do third-party tracking companies encourage and support app developers to obtain consent as and where required? (*practices of tracker companies*)

Contributions. In answering these questions, this paper makes three contributions. First, we clarify the role of consent in the regulatory framework applicable in the EU and the UK when it comes to the processing of personal data for third-party tracking. Second, we provide empirical evidence as to a widespread absence of consent mechanisms to legitimise third-party tracking in 1,297 apps. Third, we analyse the guidance provided by 13 commonly used tracker companies and assess whether they inform app developers about how to translate consent in code (see Figure 1 and Table 2).

Structure. The rest of this paper is structured as follows. Section 2 reviews the relevant literature surrounding the concept of consent, app privacy analysis, and existing system-wide tracking controls for Android. Section 3 discusses the role of consent for third-party tracking in the EU and UK by drawing on the guidance issued by national Data Protection Authorities (DPAs). Section 4 analyses the presence of consent for third-party tracking in 1,297 Android apps randomly sampled from the Google Play Store. Section 5 reviews the guidance offered by tracker companies to app developers. After discussing the limitations of our approach in Section 6, we turn to the discussion of our results in Section 7 and our conclusions in Section 8.

2 Background

In this section, we discuss previous relevant literature, including on the concept of consent, analysing privacy in apps, and existing system-wide tracking controls for Android. In particular, we highlight the limits of consent, and the dependence of end-users on the privacy options implemented by their apps and smartphone operating system.

2.1 Promises and Limits of Consent

Consent is a pillar of privacy and data protection law, in the US, EU, and many other jurisdictions and international

frameworks. As an approach to privacy protection, consent is associated with the regime of *notice & choice* [66]. For many data-processing activities, companies that want to process data from an individual must

1. Adequately inform the individual (*Notice*), and
2. Obtain consent from the individual (*Choice*).

These two fundamental requirements are often implemented in software through the provision of a privacy policy, accompanied by consent options for the end-user.

The limitations of the notice & choice paradigm have been explored in a range of scholarship. Regarding “*notice*”, it has been documented that most people do not read privacy policies, and that when they try to, have difficulties understanding them [61] and do not have enough time to read every such policy [53].

Regarding “*choice*”, evidence suggests that many individuals struggle with privacy decisions in practice [4, 64]. The mismatch between stated and observed privacy preferences is known as the “*privacy paradox*” [56], although this so-called “*paradox*” may be best explained by structural forces that prevent alignment between values and behaviour [67, 74]. Individuals often have no real choice but to accept certain data processing because some digital services—such as Facebook or Google—have become indispensable [18]. Even when offered a level of choice, individuals face ubiquitous tracking [15], are tricked into consent [57], and do not get an adequate compensation in exchange for their data [22]. Because of the limits to individual privacy management, various scholars argue that the regime of notice & choice does not provide *meaningful* ways for individuals to manage their privacy [12, 13, 66].


Despite such limitations, consent remains a key component of many privacy and data protection regimes. For the purpose of this present contribution, we do not assume that consent is the only or best way to address privacy and data protection issues. Rather, we aim to investigate whether, in addition to all these problems and limitations, the basic process of consent itself is even being followed where it is currently required in the context of third-party tracking in apps.

2.2 Analysing Privacy in Apps


There is a vast range of previous literature that has analysed the privacy practices of mobile apps, and third-party tracking in particular. Two main methods have emerged in the academic literature: dynamic and static analysis.

Dynamic analysis executes an app, and analyses its runtime behaviour. While early research analysed apps by modifying the operating system [7, 29], recent work has focused on analysing apps’ network traffic [43, 52, 58, 60, 62, 63, 65, 68, 72].

As for system modification, Enck et al. modified Android so that sensitive data flows through and off the smartphone



(a) Facebook links to a page that is supposed to explain how to implement consent in practice.



(b) Google AdMob links to an outdated library, creating unnecessary friction for consent implementation.



(c) Flurry links to a broken GDPR guide.

Figure 1: Many trackers provide information on what developers need to know to implement consent. These guides are often difficult to find, hard to read, and poorly maintained. 3 out of 13 common trackers linked to unmaintained or broken pages.

could be monitored easily [29]. Agarwal and Hall modified iOS so that users were asked for consent to the usage of sensitive information by apps [7], before the introduction of run-time permissions by Apple in iOS 6.

As for network analysis, Ren et al. instrumented the VPN functionality of Android, iOS, and Windows Phone to expose leaks of personal data over the Internet [62]. Conducting a manual traffic analysis of 100 Google Play and 100 iOS apps, they found regular sharing of personal data in plain text, including device identifiers (47 iOS, 52 Google Play apps), user location (26 iOS, 14 Google Play apps), and user credentials (8 iOS, 7 Google Play apps). Van Kleek et al. used dynamic analysis to expose unexpected data flows to users and design better privacy indicators for smartphones [72]. Reyes et al. used dynamic analysis to assess the compliance of children’s apps with COPPA [63], a US privacy law to protect children. Finding 73% of studied children’s apps to transmit personal data over the Internet, they argued that none of these apps obtained the required “*verifiable parental consent*” because their automated testing tool could trigger these network calls, and a child could likely do so as well. Okoyomon et al. found widespread data transmissions in apps that likely needed to be but were not disclosed in apps’ privacy policies, and raised doubts about the efficacy of the notice & choice regime [58] (as discussed in the previous section).

Dynamic analysis offers the advantage that is relatively simple to do, is largely device-independent, and can be used to monitor what data sharing takes place in practice; however, it has a number of limitations. The information gathered might be incomplete if not all code paths within the app involving potential data disclosures are run when the app is being analysed. Further, network-based dynamic analysis may wrongly attribute system-level communications to a studied app, e.g. an Android device synchronising the Google Calendar in the background, or conducting a network connectivity check with Google servers. Nonetheless, network-based dynamic analysis is a versatile and practical approach that has been used many times.

Static analysis infers the behaviour of an app without the need for execution. This process often relies on decompiling

an app and analysing the retrieved program code [27, 44]. The main advantage of static analysis is that it enables the analysis of apps at much larger scale (e.g. millions rather than hundreds) [15, 19, 73, 75]. As opposed to dynamic analysis, static analysis may require substantial computing resources and does not permit the direct observation of network traffic because apps are never run.

Egele et al. developed an iOS decompiler and analysed 1,407 iOS apps. They found that 55% included third-party tracking libraries [27]. Viennot et al. analysed more than 1 million apps from the Google Play Store, and monitored the changing characteristics of apps over time [73]. They found a widespread presence of third-party tracking libraries in apps (including Google Ads in 35.73% of apps, the Facebook SDK in 12.29%, and Google Analytics in 10.28%). Similarly, Binns et al. found in analysing nearly 1 million Google Play apps that about 90% may share data with Google, and 40% with Facebook [15].

The presence of consent to tracking in apps has received relatively little research attention; to the best of our knowledge, no large-scale studies in this area exist. With static analysis, it is difficult to detect at what stage a user might give consent, because of the varied implementations of consent in app code. However, network-based dynamic analysis makes this kind of consent analysis possible, and at a reasonable scale. We demonstrate this in Section 4.

2.3 Alternatives to In-App Consent

Before turning to the legal analysis concerning when consent for third-party tracking within individual apps may be required, it is worth considering the options users currently have to limit app tracking on Android at a system level. This is pertinent to our subsequent analysis because, if system-level controls were sufficient, the question of efficacy and compliance with individual app-level consent requirements might be redundant. The options for users fall into three categories: system settings, system modification, and system APIs.

System Settings. The Android operating system offers users certain possibilities to limit unwanted data collection.

Users can manage the types of data each app can access through *permissions*. This does not stop tracking, but blocks access to certain types of data, such as location. A problem inherent to the permission approach is that trackers share permission access with the apps they come bundled with. This means that, if a user allows location access to a maps app with integrated trackers, all these trackers have access as well. This, in turn, might give users a false sense of security and control. Google offers users the possibility to opt-out from personalised advertising. If users choose to do so, apps are encouraged to cease using the system-wide *Google Advertising Identifier* (AdID) for personalised advertising (although apps can continue to access the AdID). Unlike iOS, Android does not offer the option to opt-out from analytics tracking using the AdID, or to prevent apps from accessing this unique user identifier.

System Modification. Since the early days of Android, many developers have set out to modify its functionality and implement better privacy protections. *Custom ROMs* are modified versions of Android that replace the default operating system that comes pre-installed on Android smartphones. Popular examples are Lineage OS and GrapheneOS, which both try to reduce the dependency on Google on Android and increase user privacy. Another is TaintDroid, which monitors the flow of sensitive information through the system [29]. A popular alternative to custom ROMs is *rooting* devices by using exploits in the Android system to gain elevated access to the operating system, or by changing the bootloader of the Android system. Rooting is a necessary prerequisite for many privacy-focused apps, including AdAway [5], XPrivacy [17], and AppWarden [10]. System modification grants maximum control and flexibility regarding tracking, but requires a high level of technical expertise. It also relies on security vulnerabilities, often creating risks for (non-expert) users.

As a result, Google has recently begun to restrict attempts to modify Android by preventing custom ROMs from running apps using *Google’s Safety Net*. This is meant to protect sensitive apps (e.g. banking apps) from running on unsafe devices, but is also used by other popular apps such as Pokemon GO and Snapchat [35]. Some Internet outlets have declared the “*end for Android rooting, [and] custom ROMs*” [69].

System APIs. Another alternative to system modification is to develop apps that build on the capabilities of Android’s system APIs to detect and block network traffic related to tracking. Such is possible without the need for system modification at the cost of more advanced functionality. Popular apps in this category include AdGuard (using a local VPN on the Android device) [6] and DNS66 (changing the DNS settings of the Android device) [50]. Another is NetGuard [16], a firewall that allows users to monitor network connections through a VPN, and to block certain domains manually. All these tools block connections regardless of the actual content of the communications. These content-agnostic approaches can lead to overblocking and break apps.

Alternative tools aim for more fine-grained protection by removing sensitive information from network requests, such as device identifiers or location data [52, 68]. Unfortunately, these content-based approaches rely on breaking secured network connections, and on installing a self-signed root certificate on the user’s device. This practice was banned by Google with the introduction of Android 7 in 2016 because of the security risks it entails [37]. While these apps grant users the possibility to block tracking through system APIs, Google does not allow them on the Play Store [38]. Instead, users must sideload them onto their device from alternative sources, such as GitHub and F-Droid.

In conclusion, while there exists a wide array of options for end-users to reduce tracking, none of them can provide the granularity of consent implemented inside each individual app. Many of the existing tools require a high level of technical expertise, including root access or modifications to the operating system, and are therefore unsuitable for non-expert users. This makes many users dependent on the privacy solutions offered by apps themselves and their operating system.

3 When is Consent to Tracking required?

In this section, we analyse whether consent is a prerequisite for third-party tracking under EU and UK law, as well as its role under the Google Play Store policy. We focus on these jurisdictions as they have relatively stringent and specific rules on consent and third-party tracking. While similar rules exist in other jurisdictions (such as the COPPA in the US, which requires parental consent for tracking), recent regulatory actions and rich guidance issued by European regulators offer an ideal setting for a large-scale analysis. To our best knowledge, the question of consent requirements for third-party app tracking in the EU and UK has not been previously analysed in the academic literature.

Two main legal instruments are relevant to the issue of consent to third-party tracking on mobile apps: the GDPR and the ePrivacy Directive¹.

3.1 GDPR and the Need for a Lawful Ground

Applicable since 25 May 2018, the GDPR grants users various rights over their *personal data*. It also imposes a wide array of obligations on so-called *controllers*, paired with high fines for non-compliance. One of the cornerstones of the legislative reform is the concept of *Data Protection by Design*; this obliges controllers to implement appropriate technical and organisational measures to ensure and demonstrate compliance with all rules stemming from the Regulation throughout the entire personal data processing life cycle (Article 24(1) and 25(1) GDPR). All companies operating in the EU or UK

¹The ePrivacy Directive is currently under revision by EU lawmakers. A change to the current regulatory requirements in practice is expected no earlier than in two years due to nature of the EU legislative process.

must follow the rules set out in the GDPR.² Compliance with the GDPR is monitored and enforced by the *Data Protection Authorities* (DPAs) instituted in each EU Member State and in the UK. If a controller fails to comply, DPAs have the power to impose fines that can go up to €20 million or 4% of the company’s worldwide annual turnover, whichever is higher.

For the purpose of this paper, we assume that app developers qualify as *controllers*. In other words, that they “*determine the purposes and the means of the processing of personal data*” (Article 4(7) GDPR). While this might well be the case when the company actually processing the personal data at stake is also in charge of the development of the app, it is important to highlight that controllership does not always end up on their shoulders. This is the case, for instance, when a company outsources the development of its app to an external team of software developers working on the basis of clear-cut specifications and requirements, in which case the latter is likely to be considered as a *processor* (Article 4(8) GDPR) or a *third party* (Article 4(10) GDPR).

If app developers want to collect personal data for whatever purpose, they need to rely on one of the six *lawful grounds* listed in Article 6(1) GDPR. Only two usually apply in the context of mobile apps, namely: *consent* and *legitimate interests*³. On the one hand, and as specified in Article 4(11) GDPR, a valid *consent* is

any freely given, specific, informed and unambiguous indication of the data subject’s wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her

As recently clarified by the Court of Justice of the European Union, this bans the use of pre-ticked boxes to gather consent [23]. *Legitimate interests*, on the other hand, is a viable alternative to consent but requires a careful balancing exercise between the controller’s interests in processing the personal data and the data subjects’ interests and fundamental rights (Article 6(1)f GDPR) [30]. The other guarantees stemming from the GDPR (including transparency, security, purpose and storage limitation, and data minimisation) remain applicable regardless of the lawful ground used to legitimise the processing.

Consent for High-Risk Data Processing. While the controller’s legitimate interests could potentially be a viable option for legitimising third-party tracking on mobile apps, this processing is likely to qualify as a *high-risk data processing activity*.⁴ Features of third-party tracking that indicate such

²More specifically, all companies based in the EU and UK, as well as companies monitoring the behaviour of, or offering goods and services to, individuals located in the EU and UK, fall within the territorial scope of application of the GDPR (Article 3 GDPR).

³The remaining four lawful grounds listed in Article 6(1) GDPR being the fulfilment of a *contract*, a *legal obligation*, the data subject’s *vital interests*, and the performance of a *public task*.

⁴The Article 29 Working Party—an EU body to provide guidance on data

high-risk processing include the use of “*evaluation or scoring*”, “*systematic monitoring*”, “*data processed on a large scale*”, “*data concerning vulnerable data subjects*”, or “*innovative use or applying new technological or organisational solutions*”. Some of these features undoubtedly apply to third-party tracking, since tracking companies usually engage in large-scale data collection, at a high-frequency, across different services and devices, with limited user’s awareness.

The Information Commissioner’s Office (ICO)—the UK’s DPA—discourages the use of legitimate interest for high-risk activities and recommends that controllers instead rely on another lawful ground such as consent [45]. Similarly, after having analysed the case of tracking deployed on a webshop selling medical and cosmetic products [24], the German DPA came to the conclusion that the average website visitor could not reasonably expect tracking of their online activities to take place, especially when it operates across devices and services. In that case, it argued, the website visitor is not in a position to avoid the data collection. These are the concrete manifestations of the balancing exercise required by Article 6(1)f. All in all, the above-mentioned considerations disqualify the use of the controllers’ legitimate interests as an appropriate lawful ground to legitimise third-party tracking in mobile apps.

3.2 ePrivacy and the Need for Consent for Local Storage of and Access to Data

In addition to the GDPR, the ePrivacy Directive also applies to third-party tracking. This is a *lex specialis*, meaning that, when both the ePrivacy Directive and the GDPR apply in a given situation, the rules of the former will override the latter. This is the case for third-party tracking, since Article 5(3) of the ePrivacy Directive specifically requires consent for accessing or storing non-technically necessary data on a user’s device. It is widely accepted, and reflected in DPAs’ guidance, that most tracking activities are not technically necessary, and therefore require consent to store data on a user’s device [47]. So, if tracker software involves accessing or saving information on a user’s smartphone—as third-party trackers typically do on a regular basis—this requires prior consent. As a result, while consent was already the most reasonable option under the GDPR, it becomes the only viable one when combining both regulatory frameworks.

As stated above, the GDPR provides a range of possible lawful bases of which consent is just one; however, the ePrivacy Directive Article 5(3) more specifically requires consent for accessing or storing non-technically necessary data on a

protection law (now the European Data Protection Board)—has listed the 9 features commonly found in such high-risk activities, namely: 1) Evaluation or scoring, 2) Automated-decision making with legal or similar significant effect, 3) Systematic monitoring, 4) Sensitive data or data of a highly personal nature, 5) Data processed on a large scale, 6) Matching or combining datasets, 7) Data concerning vulnerable data subjects, 8) Innovative use or applying new technological or organisational solutions, and 9) Prevention of data subjects from exercising a right or using a service or a contract. [31]

user’s device. As a consequence, any further processing by the third party which is not covered by the initial consent interaction would usually require the third party to obtain fresh consent from the data subject.

Recent guidance and enforcement action from various DPAs have also demonstrated how the GDPR and the ePrivacy requirements apply to situations where consent is the basis for processing by one controller, and when that data is provided to another controller for further processing. Article 7(1) of the GDPR requires that, where consent is the lawful ground, the controller must be able to *demonstrate* that the data subject has consented. The ICO’s guidance states that third-party services should not only include contractual obligations with first parties to ensure valid consent is obtained, but “*may need to take further steps, such as ensuring that the consents were validly obtained*” [46]. It notes that, while the process of getting consent for third-party services “*is more complex*”, “*everyone has a part to play*” [46]. The responsibility of third parties has been further illustrated in an enforcement action by the CNIL (the French DPA), against Vectaury, a third-party tracking company [21]. This showed how the validity of consent obtained by an app developer is not “*transitive*”, i.e. does not carry over to the third party. If a first party obtains consent “*on behalf*” of a third party, according to a contract between the two, the third party is *still* under the obligation to verify that the consent is valid.

To summarise the implications of GDPR and ePrivacy in the context of third-party tracking: consent is typically required for access to and storage of data on the end-user’s device. Even if that consent is facilitated by the first party, third parties must also be able to demonstrate the validity of the consent for their processing to be lawful on that basis.

3.3 Requirements of the Google Play Store

In addition to EU and UK privacy law, Google imposes a layer of contractual obligations that apps must comply with. These policies apply worldwide—so beyond the jurisdiction of the EU and UK—and might oblige all app developers to implement adequate mechanisms to gather consent for third-party tracking. Google’s *Developer Content Policy* highlights that in-app disclosure and consent might need to be implemented when “*data collection occurs in the background of your app*” [40]. The Developer Content Policy also requires that developers abide by all applicable laws. It is unclear how strictly compliance with these policies—and in particular with all applicable laws—is verified and enforced by Google.

4 Tracking in Apps Before and After Consent

The previous section established that third-party tracking in apps requires valid user consent under the EU and UK regulatory framework. Despite these legal obligations, it yet not clear how and whether consent is realised in practice. In order

to examine the extent to which regulation around consent is implemented in practice, we conducted two studies—Study 1 (in this section) to see how consent is implemented in a representative sample of Google Play apps, and Study 2 (in the following Section 5) to examine how app developers were supported and encouraged to implement consent by the providers of tracker libraries.

4.1 Methodology

We studied a representative sample of 1,297 free Android apps from the UK Google Play Store. This sample was chosen randomly (through random sampling without replacement) from a set of 1.63 million apps found on the Google Play Store between December 2019 and May 2020. We explored the presence of apps on the app store by interfacing with Google Play’s search function, similar to previous research [73]. The selected apps were run on a Google Pixel 4 with Android 10. Each app was installed, run for 15 seconds, and then uninstalled. We did not interact with the app during this time, to record what companies the app contacts before the user can be informed about data collection, let alone give consent. During app execution, we recorded the network traffic of all tested apps with the popular NetGuard traffic analysis tool [16]. We did not include any background network traffic by other apps, such as the Google Play Services. For apps that showed full-screen popup ads, we closed such popups, and took note of the presence of display advertising. We assessed whether each contacted domain could be used for tracking and, if so, to what tracking company it belonged, using a combination of the App X-Ray [14] and Disconnect.me [25] tracker databases. 15 seconds after having installed the app, we took a screenshot for further analysis, and uninstalled it.

We inspected the screenshots for any form of display advertising, privacy notice or consent. We took note of any display advertising (such as banner and popups advertising) observed. We classified any form of information about data practices as a privacy notice, and any *affirmative* user agreement to data practices as consent. While this definition of consent is arguably less strict than what is required under EU and UK law, this was a deliberate choice to increase the objectivity of our classification, and provide an upper bound on compliance with EU and UK consent requirements. We then re-installed and ran those apps that asked for consent, granted consent, and repeated the network capture and analysis steps above, i.e. monitoring network connections for 15 seconds, followed by a screenshot, and finally, removed the app once again.

4.2 Results

Of the 1,297 apps, 96 did not show a working user interface. Some apps did not start or showed to be discontinued. Other apps did not provide a user interface at all, such as widgets and Android themes. We therefore only considered the re-

Hosts	Company	Apps
adservice.google.com	Alphabet	19.7%
tpc.googlesyndication.com	Alphabet	17.2%
lh3.googleusercontent.com	Alphabet	14.2%
android.googleapis.com	Alphabet	12.9%
csi.gstatic.com	Alphabet	11.6%
googleads.g.doubleclick.net	Alphabet	10.3%
ade.googlesyndication.com	Alphabet	9.7%
connectivitycheck.gstatic.com	Alphabet	9.5%
config.uca.cloud.unity3d.com	Unity	7.5%
ajax.googleapis.com	Alphabet	6.9%
api.uca.cloud.unity3d.com	Unity	6.8%
android.clients.google.com	Alphabet	6.7%
gstatic.com	Alphabet	5.8%
graph.facebook.com	Facebook	5.5%

Table 1: Top contacted tracker domains by 1,201 randomly sampled apps from the Google Play Store, at launch, before any interaction with the apps.

maining 1,201 apps. 909 apps (76%) were last updated after the GDPR became applicable on 25 May 2018.⁵ On average, the considered apps were released in August 2018 and last updated in December 2018. All apps were tested in August 2020, within a single 24-hour time frame.

Widespread tracker use. Apps contacted an average of 4.7 hosts each at launch, prior to any user interaction. A majority of such apps (856, 71.3%) contacted known tracker hosts. On average, apps contacted 2.9 tracker hosts each, with a standard deviation of 3.5. The top 10% of apps contacted at least 7 distinct hosts each, while the bottom 10% contacted none. Alphabet, the parent company of Google, was the most commonly contacted company (from 58.6% of apps), followed by Facebook (8.2%), Unity (8.2%), One Signal (5.6%), and Verizon (2.9%). Apps that we observed showing display ads contacted a significantly higher number of tracker hosts (on average 6.0 with ads vs 2.2 without).

Dominance of Google services. The 9 most commonly contacted domains all belong to Google; the top 2 domains are part of Google’s advertising business (adservice.google.com, linked to Google’s Consent API, and tpc.googlesyndication.com, belonging to Google’s real-time advertisement bidding service). 704 apps (58.6%) contacted at least one Google domain; the top (Google) domain was contacted by 236 apps (19.7%). Such breadth and variation is reflective of the corresponding variety of services that Google offers for Android developers, including an ad network (Google AdMob), an ad exchange (Google Ad Manager,

⁵It is worth noting, however, that both the need for a lawful ground—an obligation under Directive 95/46—and the consent requirement for access to and storing on terminal equipment—an obligation under the ePrivacy Directive—were already applicable before 25 May 2018. The latter has merely provided clarification on the conditions for consent to be valid.

formerly known as DoubleClick), and various other services. Domains by other tracker companies, such as Unity and Facebook, were contacted less frequently by apps (see Table 1).

Google’s tracking was also observed to be deeply integrated into the Android operating system. It has been known that the Google Play Services app—required to access basic Google services, including the Google Play Store—is involved in Google’s analytics services [55]. In our network analysis, this app seemed to bundle analytics traffic of other apps and send this information to Google in the background with a time delay. Without access to encrypted network traffic (as explained in Section 2.3), this makes it impossible to attribute network traffic to individual apps from our sample, when such network traffic could also be related to other system apps (some of which, such as the Google Phone app, use Google Analytics tracking themselves). As a consequence, we are likely under-reporting the number of apps that share data with Google, since we only report network traffic that could be clearly attributed.

Consent to tracking is widely absent. Only 9.9% of apps asked the user for consent. Apps that did so contacted a larger number of tracker hosts than those that did not (3.7 with consent vs 2.8 that did not). A slightly larger fraction (12.2% of all apps), informed the user to some extent about their privacy practices; apps in this category also contacted a larger number of trackers than those that did not (3.6 that informed vs 2.8 that did not). 19.1% of apps that did not ask for consent showed ads, compared to only 2.5% of apps that asked for consent. Once consent was granted, the apps contacted an average of 4.2 tracker hosts (higher than the 3.7 before granting consent, and the 2.8 for apps without any consent flows).

Consent limited to using or not using an app. Most apps that ask for consent force users into granting it. For instance, 43.7% of apps asking for consent only provided a single choice, e.g. a button entitled “Accept Policy and Use App” or obligatory check boxes with no alternative. A further 20.2% of apps allowed users to give or refuse consent, but exited immediately on refusal, thus providing a *Hobson’s choice*. Only 42 of the apps that implemented consent (comprising a mere 3.5% of all apps) gave users a genuine *choice to refuse* at least some aspect of tracking or other data practices. However, those apps had some of the highest numbers of tracker hosts, and contacted an average of 5.2 on launch. Among these apps, if consent was granted, the number of tracker hosts increased to 8.1, but, interestingly, an increase was also observed even if data tracking was opted-out (from the pre-consent 5.2 to 7.5 post-opt-out).

Consent limited to the personalisation of ads. Consent was often limited to an opt-out from personalised ads. 37 of the 42 apps that implement a genuine choice to refuse consent restrict this choice to limiting personalised advertising. We observed that 23 of these 37 apps used Google’s Consent API, a toolkit provided by Google for simplifying the consent

workflow for apps with personalised ads. None of the apps using the Consent API, however, ended up asking users to agree to further tracking activities, such as analytics. Only 4 apps provided the option to refuse analytics; all 4 of these did so in addition to providing the option to opt-out of personalised advertising. One further app in our sample requested consent to process health data.

5 Support and Guidance from Trackers

The previous section found a widespread absence of consent to third-party tracking in apps. As explained in Section 3, both first and third parties have a part to play in facilitating valid consent, and third parties need to take steps to ensure consent obtained by first parties is valid. At the same time, it has been reported that many app developers believe the responsibility of tackling risks related to ad tracking lie with the third-party companies [54], and need clear guidance regarding app privacy [11]. In this section, we assess the efforts, that providers of tracker libraries make, to encourage and support app developers in implementing a valid consent mechanism.

5.1 Methodology

Our qualitative analysis focuses on the 13 most common tracker companies on Android (according to [33]), and three types of document that each of them provides: 1) a step-by-step implementation guide, 2) a privacy policy, and 3) further publicly available documentation. While there may be other ways in which providers of tracking libraries support app developers to facilitate valid consent, we reason that these are the standard means by which such support would be provided. Step-by-step implementation guides serve as a primary resource for app developers and summarise the essential steps of implementing a tracker library in code. Since the implementation of consent must be done in code, consent implementation is one essential step for those trackers that require consent.

In assessing this documentation, we assume the perspective of an app developer who is motivated to comply with any explicit requirements mentioned by the tracker provider, and to follow their instructions as to how to do so, but lacks in-depth knowledge about how the GDPR and ePrivacy Directive apply to their use of a given third-party tracking software [42]. We also assume that app developers are likely to read documentation only so far as necessary to make the third-party library functional, often through trial-and-error [49, 51], and stop studying other resources once the tracker implementation is functional.

5.2 Results

Our results are summarised in Table 2. We detail our main findings in the following paragraphs.

Most trackers are unclear about their use of local storage. Whether a tracker accesses and/or stores information on a user’s device is essential in determining the need to implement consent, as explained in Section 3.2. As such, we would expect to find information stating whether or not access and/or storage takes place as part of the standard operation of the tracker. However, we did not find such information for 6 out of 13 trackers. For the others, this information was difficult to find. AppsFlyer rightly states in its online documentation that “*there are no cookies for mobile apps or devices*” [8]. While this is true from a technical perspective, EU and UK law do not differentiate between cookies and other information saved on a user’s device. Crucially, we did not find any tracker stating *not* to save information on a user’s device. In the absence of such a denial, app developers would run the risk of assuming they do not need to obtain consent for data accessed and/or stored by the tracker.

Most trackers expect app developers to obtain consent. Despite being unclear about their use of local storage, a closer inspection of the tracker policies and documentation found that most trackers instruct developers to request consent from EU users (11 out of 13). AppLovin is an exception, but does require consent if developers want to show personalised ads (which tend to be more lucrative than contextual ads). For AppsFlyer, we could not find any information regarding the need to ask users for consent. The need to ask for consent was sometimes difficult to find, and required a careful reading of the policies and documentation provided. Some developers are bound to overlook this, and unnecessarily compromise on the users’ right to choose over tracking.

Few trackers implement consent by default. We further inspected whether tracker libraries provide their own consent implementation. If they do, an app developer would not need to make any further modification to the app code. However, only a minority of tracker libraries (3 out of 13) integrates an implementation of user consent by default, and none of the five most common trackers do so. Unity Ads and Twitter MoPub provide consent flows that are automatically shown, without further action by the app developer. Facebook Ads only shows ads, if the app user 1) has agreed to personalised ads in their Facebook account settings, and 2) uses the Facebook app on their phone. However, Facebooks opts-in users by default to their personalised advertising, unless they disable this behaviour from their Facebook settings (checked 14 February 2021). While Google AdMob provides a consent library, this is not implemented by default. Indeed, Google AdMob expects the app developer to retrieve consent from the user, but shows personalised ads even if the developer does not implement their consent library.

Limited disclosure of consent requirements in step-by-step guides. We find that 3 out of 13 tracker libraries disclose the potential need for consent in their step-by-step implementation guides. This is despite 11 out of 13 trackers mentioning the need to implement consent in other places of their online

Tracker	Apps	Expects consent (in EU / UK)	Implements consent (by default)	Mentions consent (in implementation guide)	Discloses local data storage
Google Analytics	50%	Yes	No	No	Yes
Google AdMob	45%	Yes	No	Yes	Yes
Google Crashlytics	29%	Yes	No	No	Yes
Facebook App Events	20%	Yes	No	No	?
Google Tag Manager	19%	Yes	No	No	Yes
Facebook Ads	14%	Yes	Yes*	No	?
Flurry	9%	Yes	No	No	?
Unity Ads	8%	Yes	Yes	No	Yes
Inmobi	8%	Yes	No	Yes	?
Twitter MoPub	6%	Yes	Yes	No	Yes
AppLovin	6%	No	No	No	?
AppsFlyer	5%	?	No	Yes	?
OneSignal	4%	Yes	No	No	Yes

Table 2: Consent requirements and implementation for 13 commonly used Android trackers. App shares according to the Exodus Privacy Project [33]. The **trackers in bold** require consent, but do neither implement such by default nor mention the need to do so in their implementation guides. ?: We did not find any information. *: Facebook opts-in users by default to their personalised advertising, unless they disable this behaviour from their Facebook settings or do not use the Facebook app.

documentation. Google AdMob mentions the need to retrieve consent amongst other “*examples of actions that might be needed prior to initialization*” [39] of AdMob. Inmobi points out that developers need to “*obtain appropriate consent from the user before making ad requests to InMobi for Europe*” [48] in the Section on “*Initializing the SDK*”. AppsFlyer offers developers to “*postpone start [of the tracker library] until you receive user consent due to GDPR or CCPA requirements, etc.*” [9] in Section 3.4 on “*Delay SDK initialization*”. It is not clear from these three implementation guides what other reasons are to “*delay initialisation*” beyond legal compliance, and why this is not clarified. At least 6 out of 13 trackers require consent, but neither implement such by default nor inform app developers of the need to do so in the implementation guides. If AppLovin needs consent (despite not stating to do so, but as suggested by our legal analysis in Section 3), this figure would increase to 7 out of 13 trackers.

Compliance guidance: often provided, but sometimes difficult to find, hard to read, and poorly maintained. Many tracker companies provide additional information on GDPR compliance and consent implementation on a separate website as part of their online documentation. We found some compliance guidance (with varying levels of detail) for all trackers except the Google Tag Manager. Excluding the 3 trackers implementing consent by default, a developer needs an average of 1.56 clicks to reach these compliance guides. For AppLovin, a developer must click “*Help Center*”, then “*Getting started & FAQ*”, and lastly “*User opt-in/opt-out in the AppsFlyer SDK*”. Facebook required developers to click “*Best Practices Guide*” and then “*GDPR Compliance guide*”. While this GDPR compliance guide provides some guidance

on the implementation of consent, the link to Facebook’s “*consent guide*” with practical examples of how to implement consent was broken. Also, the framing as “*Best Practices*” suggests optionality of legal compliance. For OneSignal, developers must first click “*Data and Security Questions*” and then “*Handling Personal Data*”.

The compliance guides (excluding code fragments) reached a mean Flesch readability score [36] of 41.8, as compared to 50.6 for the step-by-step implementation guides (where 100 means “*very easy*”, and 0 “*very difficult*” to read). Both the implementation and compliance guides are “*difficult*” to read, with the compliance guides somewhat more so. For 3 of the 13 trackers, we were directed to broken or outdated links (see Figure 1). Google AdMob linked to an outdated consent strategy, while the Facebook SDK and Flurry linked to non-existing pages (returning 404 errors). We found other pages with compliance information for each of these trackers, but broken guidance can act as a deterrent for developers who want to implement consent and follow their legal obligations. We have notified the companies about this situation.

6 Limitations

It is important to acknowledge some limitations of our methodology. Our analysis in Section 4 used dynamic analysis, and not all tracking might be detected. We only inspected network traffic before and shortly after consent was given. Apps might therefore conduct more tracking during prolonged app use. Besides, we only reported the network traffic that could be clearly attributed to one of the apps we studied, potentially leading to under-reporting of the extent of Google’s

tracking (as explained in Section 4). As for our second study in Section 5, we studied the online documentation of tracker libraries with great care, but did not always find all relevant information, particularly regarding the local storage of data on a user’s device. Where this was the case, we disclosed this (e.g. see Table 2).

7 Discussion and Future Work

Consent is an integral part of data protection and privacy legislation, both in the EU and the UK, and elsewhere. This is all the more so in the context of third-party tracking, for which consent appears the only viable lawful ground under the ePrivacy Directive and the GDPR, as analysed in Section 3. Not only has this been emphasised by multiple DPAs, but is also acknowledged by tracking companies themselves in the documentation they make available to app developers. Relying on the controller’s legitimate interests—the only conceivable alternative to consent under EU and UK data protection law—would likely fail short of passing the balancing test outlined in Article 6(1)f GDPR. This also follows from the requirement to obtain consent prior to storing or accessing information on a user’s device, under the ePrivacy Directive.

Against this backdrop, we analysed 1,297 mobile apps from Google Play in Section 4 and discovered a widespread lack of appropriate mechanisms to gather consent as required under the applicable regulatory framework. We found that, while the guidelines of many commonly used tracker libraries require consent from EU and UK users, most apps on the Google Play Store that include third-party tracking features do not implement any type of consent mechanism. The few apps that require data subjects to consent do so with regard to personalised advertising, but rarely for analytics—despite this being one of the most common tracking practices. Where an opt-out from personalised advertising was possible, the number of tracker domains contacted decreased only slightly after opting-out, hinting to continued data collection when serving contextual advertising. These observations are at odds with the role of consent as the only viable option to justify the processing of personal data inherent to third-party tracking.

As detailed in Section 4, the fact that only 9.9% of the investigated apps request any form of consent already suggests widespread violations of current EU and UK privacy law. This is even before considering the validity of the consent mechanisms put in place by that small fraction of apps. As underlined in Section 3, consent must be “*freely given*”, “*informed*”, “*specific*” and “*unambiguous*”. The findings outlined in Section 4 suggest that most apps that do implement consent force users to grant consent, therefore ruling out its qualification as “*freely given*”. The same goes for the 43.7% of those apps that do not provide data subjects with the possibility to consent separately for each purpose, but instead rely on *bulk* consent for a wide array of purposes.

When considering both the absence of any form of consent

in more than 90% of the investigated apps and the shortcomings inherent to the few consent mechanisms that are implemented by the remaining sample, we infer that the vast majority of mobile apps fail short of meeting the requirements stemming from EU and UK data protection law. Our analysis does not even consider the fact that consent is only one of a variety of legal rules that third-party tracking needs to comply with. Breaches of other legal principles—such as data minimisation, purpose and storage limitation, security and transparency—might be less visible than a lack of consent and harder to analyse, but no less consequential.

We further found that one of the reasons for the lack of consent implementation in apps might be inadequate support by tracker companies [11, 54]. Studying the online documentation of the 13 most commonly used tracker libraries in Section 5, only 3 trackers implemented consent by default, and another 3 disclosed the need to implement consent as part of step-by-step implementation guides. These step-by-step guides serve as a primary resource for app developers, and can give a false impression of completeness when in fact additional code needs to be added for many trackers to retrieve user consent. This is true for at least 6 out of 13 trackers, including Google Analytics and the Facebook App Events SDK, which likely need consent, but neither disclose this in their implementation guides nor implement such consent by default. While most trackers provide some compliance guidance, we found that this can be difficult to find, hard to read, and poorly maintained. Whatever the reasons for the lack of consent, the result is an absence of end-user controls for third-party tracking in practice.

Lastly, it is worth highlighting that Google, which is both the largest tracking company and the main developer of Android, faces conflicts of interest with respect to protecting user privacy in its Google Play ecosystem [26, 41, 70]. The company generates most of its revenue from personalised advertising, and relies on tracking individuals at scale. Certain design choices by Google, including its ban of anti-tracking apps from the Play Store, its recent action against modified versions of Android, and the absence of user choice over AdID access for analytics on Android (as opposed to iOS), create friction for individuals who want to reduce data collection for tracking purposes, and lead to increased collection of personal data, some of which is unlawful as our analysis has shown.

Future work. An overarching question for future work is the extent of the legal obligations faced by the many actors involved in the third-party tracking ecosystem, ranging from app developers to providers of tracker libraries and mobile operating systems. This is inextricably linked to their qualification as “*controllers*”, a legal notion the boundaries of which remain, despite recent jurisprudence and detailed guidance, still controversial. Our analysis highlighted how simple changes in the software design can have significant effects for user privacy.

Moreover, while the US—unlike many developed countries—lack a federal privacy law, there exists a variety of specific privacy laws, such as COPPA to protect children and HIPAA to protect health data, as well as state-level privacy laws, including CCPA in California. Some of these laws foresee consent requirements similar to EU and UK law. We leave it to further work to assess how widely apps comply with the consent requirements of US privacy legislation.

8 Conclusions

Our work analyses the legal requirements for consent to tracking in apps, and finds an absence of such consent in practice, based on an analysis of a representative sample of Google Play apps. This, in turn, suggests widespread violations of EU and UK privacy law. Simple changes by software intermediaries (such as Google and Facebook), including default consent implementations in tracker libraries, better legal guidance for app developers, and better privacy options for end-users, could improve the status quo around app privacy significantly. Yet, there is doubt that these changes will happen without further intervention by independent parties—not only end-users, but also lawmakers and regulators—due to inherent conflicts between user privacy and surveillance capitalism.

While the web has seen a proliferation of deceptive consent banners in recent years, we hope that apps will not see a similar mass adoption. Rather, we aim to provide valuable insights for the current policy discourse around user choice over tracking and ultimately to make such choice more meaningful. As Apple has demonstrated with its recently introduced iOS 14.5, system-level user choices can streamline the process of retrieving user consent and make hidden data collection more transparent.

References

- [1] 9to5mac.com. Apple rebuffs Facebook criticism, says iOS anti-tracking features are about 'standing up for our users'. <https://9to5mac.com/2020/12/16/apple-facebook-app-tracking-transparency/>, 2020.
- [2] 9to5mac.com. Facebook attacks Apple in full-page newspaper ads. <https://9to5mac.com/2020/12/16/facebook-attacks-apple/>, 2020.
- [3] 9to5mac.com. Apple versus Facebook on ad-tracking: Harvard sides with Apple. <https://9to5mac.com/2021/02/05/apple-versus-facebook-harvard/>, 2021.
- [4] Alessandro Acquisti. Nudging Privacy: The Behavioral Economics of Personal Information. *IEEE Security & Privacy Magazine*, 7(6):82–85, 2009.
- [5] AdAway. AdAway. <https://github.com/AdAway/AdAway>, 2021.
- [6] AdGuard. AdGuard for Android. <https://adguard.com/en/adguard-android/overview.html>, 2021.
- [7] Yuvraj Agarwal and Malcolm Hall. 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*, page 97. ACM Press, 2013.
- [8] AppsFlyer. 360° Mobile Attribution. <https://www.appsflyer.com/product/mobile-attribution-for-user-acquisition/>, 2021.
- [9] AppsFlyer. Android SDK integration for developers. <https://support.appsflyer.com/hc/en-us/articles/207032126-Android-SDK-integration-for-developers#integration>, 2021.
- [10] Aurora Open Source Software. Warden : App management utility. <https://gitlab.com/AuroraOSS/AppWarden>, 2021.
- [11] Rebecca Balebako, Abigail Marsh, Jialiu Lin, Jason Hong, and Lorrie Faith Cranor. The privacy and security behaviors of smartphone app developers. In *Proceedings 2014 Workshop on Usable Security*. Internet Society, 2014.
- [12] Solon Barocas and Helen Nissenbaum. On notice: The trouble with notice and consent. In *Proceedings of the engaging data forum: The first international forum on the application and management of personal electronic information*, 2009.
- [13] Elettra Bietti. Consent as a Free Pass: Platform Power and the Limits of the Informational Turn. *Pace Law Review*, page 60, 2020.
- [14] Reuben Binns, Ulrik Lyngs, Max Van Kleek, Jun Zhao, Timothy Libert, and Nigel Shadbolt. Third party tracking in the mobile ecosystem. In *Proceedings of the 10th ACM Conference on Web Science - WebSci '18*, pages 23–31. ACM Press, 2018.
- [15] Reuben Binns, Jun Zhao, Max Van Kleek, and Nigel Shadbolt. Measuring third-party tracker power across web and mobile. *ACM Transactions on Internet Technology*, 18(4):1–22, 2018.
- [16] Marcel Bokhorst. NetGuard. <https://github.com/M66B/NetGuard>, 2021.
- [17] Marcel Bokhorst. XPrivacyLua. <https://lua.xprivacy.eu/>, 2021.

- [18] Bundeskartellamt. B6-22/16 (facebook v bundeskartellamt).
- [19] Kai Chen, Xueqiang Wang, Yi Chen, Peng Wang, Yeon-joon Lee, XiaoFeng Wang, Bin Ma, Aohui Wang, Yingjun Zhang, and Wei Zou. Following Devil’s Footprints: Cross-Platform Analysis of Potentially Harmful Libraries on Android and iOS. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 357–376. IEEE, 2016.
- [20] Saksham Chitkara, Nishad Gothoskar, Suhas Harish, Jason I. Hong, and Yuvraj Agarwal. Does this app really need my location?: Context-aware privacy management for smartphones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):1–22, 2017.
- [21] Commission Nationale de l’Informatique et des Libertés. Décision n° MED 2018-042 du 30 octobre 2018 mettant en demeure la société X. <https://www.legifrance.gouv.fr/cnil/id/CNILTEXT000037594451/>, 2018.
- [22] Competition and Markets Authority. Online platforms and digital advertising. https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final_report_1_July_2020_.pdf.
- [23] Court of Justice of the European Union. Orange România SA v Autoritatea Națională de Supraveghere a Prelucrării. <http://curia.europa.eu/juris/document/document.jsf?docid=233544&doclang=EN>, 2020.
- [24] Datenschutzkonferenz. Orientierungshilfe der Aufsichtsbehörden für Anbieter von Telemedizin.
- [25] Disconnect.me and Mozilla. Firefox Blocklist. <https://github.com/mozilla-services/shavar-prod-lists>.
- [26] Benjamin G Edelman and Damien Geradin. Android and competition law: Exploring and assessing google’s practices in mobile. *European Competition Journal*, 2016.
- [27] Manuel Egele, Christopher Kruegel, Engin Kirda, and Giovanni Vigna. PiOS: Detecting Privacy Leaks in iOS Applications. In *Proceedings of NDSS 2018*, 2011.
- [28] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. “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)*, pages 1–24. ACM Press, 2021.
- [29] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. 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*, pages 393–407, 2010.
- [30] EU Article 29 Data Protection Working Party. Opinion 06/2014 on the notion of legitimate interests of the data controller under Article 7 of Directive 95/46/EC. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp217_en.pdf, 2014.
- [31] EU Article 29 Data Protection Working Party. Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is “likely to result in a high risk” for the purposes of Regulation 2016/679. https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=611236, 2017.
- [32] European Parliament and Council. Regulation 2016/679 (General Data Protection Regulation). <http://data.europa.eu/eli/reg/2016/679/oj>, 4 2016.
- [33] Exodus. Statistics. <https://reports.exodus-privacy.eu.org/en/trackers/stats/>.
- [34] Ronan Ó Fathaigh. Mobile privacy and business-to-platform dependencies: An analysis of SEC disclosures. *Journal of Business*, page 58, 2018.
- [35] Eric Ferrari-Herrmann. Trapped in Google’s safety net: what modders need to know. , 2021.
- [36] Rudolph Flesch. A new readability yardstick. *Journal of Applied Psychology*, 32(3):221–233, 1948.
- [37] Google. Android 7.0 for Developers. https://developer.android.com/about/versions/nougat/android-7.0#default_trusted_ca, 2016.
- [38] Google. Device and network abuse. <https://support.google.com/googleplay/android-developer/answer/9888379>, 2021.
- [39] Google. Get started with AdMob in your Android project. <https://firebase.google.com/docs/admob/android/quick-start>, 2021.
- [40] Google. User Data. <https://support.google.com/googleplay/android-developer/answer/10144311>, 2021.
- [41] Daniel Greene and Katie Shilton. Platform privacies: Governance, collaboration, and the different meanings

- of “privacy” in iOS and android development. *New Media & Society*, 20(4):1640–1657, 2018.
- [42] Irit Hadar, Tomer Hasson, Oshrat Ayalon, Eran Toch, Michael Birnhack, Sofia Sherman, and Arod Balissa. Privacy by designers: software developers’ privacy mindset. *Empirical Software Engineering*, 23(1):259–289, 2018.
- [43] Catherine Han, Irwin Reyes, Amit Elazari, Joel Reardon, Alvaro Feal, Kenneth A. Bamberger, Serge Egelman, and Narseo Vallina-Rodriguez. 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)*, 2019.
- [44] Jin Han, Qiang Yan, Debin Gao, Jianying Zhou, and Robert H Deng. Comparing Mobile Privacy Protection through Cross-Platform Applications. In *Proceedings 2013 Network and Distributed System Security Symposium*, page 16. Internet Society, 2013.
- [45] Information Commissioner’s Office. How do we apply legitimate interests in practice? <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/legitimate-interests/how-do-we-apply-legitimate-interests-in-practice/>, 2021.
- [46] Information Commissioner’s Office. How do we comply with the cookie rules? <https://ico.org.uk/for-organisations/guide-to-pecr/guidance-on-the-use-of-cookies-and-similar-technologies/how-do-we-comply-with-the-cookie-rules/>, 2021.
- [47] Information Commissioner’s Office. What are the rules on cookies and similar technologies? <https://ico.org.uk/for-organisations/guide-to-pecr/guidance-on-the-use-of-cookies-and-similar-technologies/what-are-the-rules-on-cookies-and-similar-technologies/#rules10>, 2021.
- [48] Inmobi. Android Guidelines: Getting Started with Android SDK Integration. <https://support.inmobi.com/monetize/android-guidelines/>, 2021.
- [49] Caitlin Kelleher and Michelle Ichinco. Towards a model of api learning. In *2019 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pages 163–168. IEEE, 2019.
- [50] Julian Andres Klode. DNS-Based Host Blocking for Android. <https://github.com/julian-klode/dns66>, 2021.
- [51] Joseph Lawrance, Christopher Bogart, Margaret Burnett, Rachel Bellamy, Kyle Rector, and Scott D Fleming. How programmers debug, revisited: An information foraging theory perspective. *IEEE Transactions on Software Engineering*, 39(2):197–215, 2010.
- [52] Anh Le, Janus Varmarken, Simon Langhoff, Anastasia Shuba, Minas Gjoka, and Athina Markopoulou. Antmonitor: A system for monitoring from mobile devices. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data*, C2B(1)D ’15, pages 15–20, 8 2015.
- [53] Aleecia M McDonald and Lorrie Faith Cranor. The Cost of Reading Privacy Policies. *I/S: A Journal of Law and Policy for the Information Society*, page 26, 2008.
- [54] Abraham H Mhaidli, Yixin Zou, and Florian Schaub. “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*, page 21, 2019.
- [55] microg. Implementation Status. <https://github.com/microg/GmsCore/wiki/Implementation-Status>, 2020.
- [56] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2017.
- [57] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020.
- [58] Ehimare Okoyomon, Nikita Samarin, Primal Wijesekera, Amit Elazari, Narseo Vallina-Rodriguez, Irwin Reyes, Álvaro Feal, and Serge Egelman. On The Ridiculousness of Notice and Consent: Contradictions in App Privacy Policies. *The Workshop on Technology and Consumer Protection (ConPro ’19)*, 2019.
- [59] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. User Tracking in the Post-cookie Era: How Websites Bypass GDPR Consent to Track Users. *WWW’2021*, 2021.
- [60] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. In *Proceedings of NDSS 2018*, 2 2018.

- [61] Joel R. Reidenberg, Jaspreet Bhatia, Travis D. Breaux, and Thomas B. Norton. Ambiguity in Privacy Policies and the Impact of Regulation. *The Journal of Legal Studies*, 45(S2):S163–S190, 2016.
- [62] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 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*, pages 361–374. ACM Press, 2016.
- [63] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. “Won’t Somebody Think of the Children?” Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies*, 2018(3):63–83, jun 2018.
- [64] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 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*, pages 2347–2356. ACM Press, 2014.
- [65] Anastasia Shuba, Athina Markopoulou, and Zubair Shafiq. Nomoads: Effective and efficient cross-app mobile ad-blocking. In *Proceedings on Privacy Enhancing Technologies 2018*, pages 125–140, 10 2018.
- [66] Daniel J. Solove. Privacy Self-Management and the Consent Dilemma. *Harvard Law Review*, 2012.
- [67] Daniel J Solove. The myth of the privacy paradox. Available at SSRN, 2020.
- [68] Yihang Song and Urs Hengartner. 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, pages 15–26, 2015.
- [69] Juan Carlos Torres. Google SafetyNet update might be the end for Android rooting, custom ROMs. <https://www.slashgear.com/google-safetynet-update-might-be-the-end-for-android-rooting-custom-roms-30627121/>, 2020.
- [70] UK Competition and Markets Authority. Online platforms and digital advertising market study final report. https://assets.publishing.service.gov.uk/media/5efc57ed3a6f4023d242ed56/Final_report_1_July_2020_.pdf, 2020.
- [71] United States Congress. Children’s Online Privacy Protection Act. <https://www.ftc.gov/system/files/2012-31341.pdf>, 1998.
- [72] Max Van Kleek, Iliaria Liccardi, Reuben Binns, Jun Zhao, Daniel J. Weitzner, and Nigel Shadbolt. 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*, pages 5208–5220. ACM Press, 2017.
- [73] Nicolas Viennot, Edward Garcia, and Jason Nieh. A measurement study of google play. In *The 2014 ACM International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '14, pages 221–233, 2014.
- [74] Ari Ezra Waldman. Cognitive biases, dark patterns, and the ‘privacy paradox’. *Current opinion in psychology*, 31:105–109, 2020.
- [75] Haoyu Wang, Zhe Liu, Jingyue Liang, Narseo Vallina-Rodriguez, Yao Guo, Li Li, Juan Tapiador, Jingcun Cao, and Guoai Xu. Beyond google play: A large-scale comparative study of chinese android app markets. In *Proceedings of the Internet Measurement Conference 2018*, IMC '18, pages 293–307, 2018.
- [76] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N Cameron Russell, and Norman Sadeh. Maps: Scaling privacy compliance analysis to a million apps. *Privacy Enhancing Technologies Symposium 2019*, 72, 6 2019.