

Noura Alomar* and Serge Egelman

Developers Say the Darnedest Things: Privacy Compliance Processes Followed by Developers of Child-Directed Apps

Abstract: We investigate the privacy compliance processes followed by developers of child-directed mobile apps. While children’s online privacy laws have existed for decades in the US, prior research found relatively low rates of compliance. Yet, little is known about *how* compliance issues come to exist and how compliance processes can be improved to address them. Our results, based on surveys ($n = 127$) and interviews ($n = 27$), suggest that most developers rely on app markets to identify privacy issues, they lack complete understandings of the third-party SDKs they integrate, and they find it challenging to ensure that these SDKs are kept up-to-date and privacy-related options are configured correctly. As a result, we find that well-resourced app developers outsource most compliance decisions to auditing services, and that smaller developers follow “*best-effort*” models, by assuming that their apps are compliant so long as they have not been rejected by app markets. We highlight the need for usable tools that help developers identify and fix mobile app privacy issues.

Keywords: Privacy compliance, Software developers, Children privacy, Google Play, SDKs

DOI 10.2478/popets-2022-0108

Received 2022-02-28; revised 2022-06-15; accepted 2022-06-16.

1 Introduction

Since 2000, online services directed at children under 13 in the US have had to comply with the Children’s Online Privacy Protection Act (COPPA) [32, 34]. Understanding the successes and failures of differing compliance processes is important, especially as more and more jurisdictions adopt privacy regulations with re-

quirements similar to COPPA’s. For example, the General Data Protection Regulation (GDPR) [26] in the EU and the California Consumer Privacy Act (CCPA) [7], like COPPA, require verifiable parental consent when the data subject is below a certain age (§2).

US regulators regularly pursue enforcement actions against those violating COPPA, including developers of mobile apps and third-party SDKs. For example, HyperBeard, a developer of child-directed apps available on the Google Play Store, was alleged to have shared users’ persistent identifiers with third-party SDKs for tracking purposes [9]. Both InMobi and OpenX, providers of advertising SDKs, reached settlements with the U.S. Federal Trade Commission (FTC) over separate allegations that they had collected location data to track child users without parental consent [2, 20]. Despite ongoing public enforcement actions over COPPA’s nearly 25 years of existence, prior research on COPPA compliance found potential violations in a majority of child-directed Android apps [49, 60, 64]. While many of these issues were attributed to third-party SDKs that developers had embedded within their apps [63, 64], little is known about *why* these issues persist, the effectiveness of various compliance processes, or how processes can be improved.

We investigate the privacy compliance processes followed by developers of child-directed apps.¹ We distributed two surveys to developers of apps available to children on Google Play, followed by semi-structured interviews to gain additional insights. We paid particular attention to developers’ perspectives on the requirements of COPPA, GDPR, and CCPA, their experiences with app market policies, their data-collection and consent-handling procedures, and the processes followed to select and configure third-party SDKs. We also performed technical analyses of some of the apps published by our participants, so that we could discuss the results during the interviews.

*Corresponding Author: Noura Alomar: University of California, Berkeley, E-mail: nnalomar@berkeley.edu

Serge Egelman: University of California, Berkeley and International Computer Science Institute, E-mail: egelman@cs.berkeley.edu

¹ We use “developer” to refer to an app development organization. Our recruitment targeted anyone involved in developing child-directed apps, including technical and managerial roles.

Our sample included software engineers, CEOs, and product managers who work for app development companies of varying sizes, as well as independent developers. We found that most developers place a large degree of trust in the policy enforcement performed by app markets and consider their apps compliant with privacy laws once accepted for inclusion on the stores (§5.3.3). We also found that many developers are unaware of relevant SDK privacy settings, which are necessary to configure SDKs for use in compliance with applicable privacy laws (§5.7.3). Our findings highlight the need for usable tools that allow developers to identify potential privacy violations and guide them towards improving their apps' privacy compliance. We discuss the challenges that developers of child-directed apps face when trying to comply with privacy regulations and present recommendations to make it easier for developers to identify gaps in their privacy-compliance processes (§6).

2 Kids' Privacy Regulations

The Children's Online Privacy Protection Act (COPPA) [33] is a federal law that protects the privacy of children residing in the US from inappropriate collection of their personal data by online services, which includes websites and mobile apps [32]. Prohibited practices include collecting contact information (including location data), as well as collecting identifiable data for profiling or behavioral advertising, without verifiable parental consent. It applies to services whose target audiences *include* children—defined as those under 13—even if children are not their “*primary audience*.” In the case of mixed audiences that include children, if users' data may be used for COPPA-prohibited purposes, child users must be identified (e.g., using age gates) so that their data receives COPPA-compliant treatment.

The EU's General Data Protection Regulation (GDPR) [26] is a privacy law that applies to online services available to individuals in Europe [26]. Under Article 6 [5], online services are required to have a legal basis for processing their users' data, and obtaining informed consent is one of the legal bases recognized in the law. The GDPR expressly recognizes that “[c]hildren merit *specific protection*” [13]. Thus, when consent is the legal basis for data processing, that consent must be provided or approved by a parent, when the user is a child. The child provisions of the GDPR apply to users under the age of 16, though any EU member state is given the flexibility to lower it to 13 [6].

The California Consumer Privacy Act (CCPA) [7] is a state privacy law that is intended to protect the privacy rights of California residents [7]. Not all online businesses serving Californians are required to comply, and this determination can be made based on the size of the business, the number of users in California, and the amount of revenue produced as a result of data collection [7]. The CCPA requires obtaining verifiable consent from parents before selling data collected from children under the age of 13. However, children who are between 13 and 16 years of age may consent [27]. That is, under the CCPA, those over 16 must explicitly opt out of data sales, those 13–16 must explicitly opt in to data sales, and the parents of those under 13 must choose to opt their children in to data sales.

When COPPA, GDPR or the CCPA apply to a developer of child-directed app(s), and if consent is relied upon as a basis for lawful data processing, the developer would be required to be transparent with regards to the types of data they collect from their users, the kinds of third parties that receive users' data, and the purposes of their data processing activities [7, 12, 27]. Data subjects are provided with other rights, including the ability to withdraw their consent, request that the developer destroy the data they have about them, inquire about how their children's data was used or processed, and stop any potential sales of their data (e.g., to advertising networks) [7, 26].

3 Related Work

In 2012, the FTC released two reports that investigated the privacy-related behaviors and disclosures of child-directed mobile apps and identified a range of potential COPPA violations [21, 22]. More than 10 years earlier, the FTC showed that enacting COPPA had an effect on encouraging websites to be transparent and more protective of their users' data [24]. These reports spurred further research that aimed to understand the privacy behaviors of mobile apps at a larger scale [43, 47, 54, 57, 62–64, 76], examine consumer expectations of privacy [48, 56, 58, 66, 71, 75], and understand app developers' privacy practices [53, 61, 65, 67].

Studies have automated the process of examining actual privacy behaviors of apps and identified a range of potential compliance issues, including the unlawful exfiltration and transmission of children's personal data, not obtaining verifiable parental consent prior to sharing children's data with third parties for prohibited pur-

poses, and not being transparent to parents about how their children's data might be used or shared [62, 64, 76]. The large-scale investigation of potential COPPA violations conducted by Reyes et al. [64] attributed many potential issues to third-party SDKs, and found that even though many SDKs offered privacy configurations that developers could use to make their apps compliant, they were often not used. Similarly, both Kollnig et al. and Nguyen et al. tested apps' GDPR compliance and found that third parties received data from apps without using appropriate consent mechanisms [49, 59].

There is also evidence of third-party SDKs surreptitiously attempting to exfiltrate users' personal data from their devices by exploiting side and covert channels or initiating API calls to other SDKs within the same apps to circumvent consent mechanisms [63, 72]. Binns et al. observed that child-directed apps were *more* likely to transmit personal data to third parties, as compared to general-audience apps [40]. To address the privacy risks introduced by third-party SDKs, Bhoraskar et al. proposed a system to detect their malicious behaviors at runtime, which when used on child-directed apps uncovered that most prompted kids to disclose their personal data to advertisers [39]. Researchers who tested apps targeting wider audiences (e.g., parental control apps [44] and smart TV apps [57]) similarly identified behaviors that threaten user privacy and make the apps potentially non-compliant with applicable regulations.

Other researchers focused on investigating users' privacy expectations [36, 56, 58, 66, 71], awareness of their rights under applicable privacy regulations [75], and abilities to make informed privacy decisions [48, 66]. Sun et al. interviewed child users and identified gaps in their mental models, showing that they are often incapable of recognizing privacy threats [66]. Apthorpe et al. found that parents' expectations of how their children's privacy should be protected were consistent with the provisions of COPPA [36].

Researchers have analyzed developers' posts in online forums and consistently found that getting apps approved by app markets is one of the main reasons why developers might seek privacy-related guidance [53, 65, 69]. Mhaidli et al. [55] qualitatively investigated the factors that impact how app developers select and configure third-party advertising SDKs and found that developers' decisions can mostly be explained by their desire to select prominent SDKs and use their default configurations. Bamberger and Mulligan [38] interviewed privacy practitioners in multiple countries and found that having well-structured privacy-oriented processes is essential for steering employees towards consid-

ering privacy implications. Prior work also found that not having robust processes for handling security issues, and lack of access to security expertise or organizational leadership willing to provide necessary resources are some of the main barriers to adopting secure engineering practices [35, 37, 41, 52, 70, 73, 74].

Most relevant to our study is the work of Ekambaranathan et al. [42], who targeted developers of child-directed apps in their interviews and survey. Although they confirmed that app developers rely on app markets for privacy-related advice and trust popular third-party SDKs, our investigation differs in that it is specifically geared towards obtaining a holistic process-oriented understanding of how developers of child-directed apps address privacy compliance. We gave particular attention to investigating app developers' consent handling procedures, their awareness of specific privacy regulations and perceptions of their compliance obligations, and whether they utilize privacy compliance options offered by third-party SDKs. We additionally analyzed the privacy behaviors of developers' apps so that we can compare developers' self-reported responses about their processes and perceived compliance with ground truth data: whether their apps *actually* comply.

4 Methodology

Our primary goal was to understand the organizational privacy practices of developers of children's apps. To that end, we designed an initial survey and recruited respondents by identifying the official contact emails for developers of children's apps published on the Google Play store (§4.1). The results inspired us to identify questions for individual developers, allowing us to better contextualize different privacy compliance processes and developers' roles. We therefore targeted our second survey at personnel within organizations and used it to learn about their specific roles (§4.2). To capture additional process-related details, we subsequently recruited app developers for semi-structured interviews (§4.3).

We collected responses for the first survey² from April–August 2021, whereas the responses for the second survey were collected from August–October 2021, and the interviews were conducted from September–

² Our Institutional Review Board (IRB) did not consider the first survey on organizational processes to be human subjects research; they approved our second survey and interviews.

December 2021. As an incentive to participate, we offered drawings for five \$200 Amazon gift cards at each of the three stages of the study. We promised study participants anonymity: participants wishing to participate in the raffle provided email addresses, which were stored separately from—and unlinkable to—the study data. We collected no other personally-identifiable information to preserve participants' privacy. In the two surveys, we used identifiers coded to organizations, allowing us to link survey responses to apps available in the Google Play store, which we separately analyzed.

4.1 Survey 1: Organizational Processes

Our initial survey had 37 questions across five sections focusing on understanding developers' awareness of the data transmission and collection behaviors of their apps, privacy compliance processes, and familiarity with relevant children's privacy laws (Appendix A):

- A section on understanding whether developers are fully aware of their apps' behaviors, including: (1) the types of data collected, (2) the purposes of data collection, (3) whether parental consent is obtained before any data collection or transmission, and (4) the means by which consent is obtained.
- A section on understanding: (1) whether developers have formal processes for addressing privacy issues, (2) whether they are familiar with COPPA, GDPR, and CCPA, and believe that these regulations apply to their apps, and (3) the processes followed to vet third-party code for inclusion in their apps.
- Three more sections for each privacy regulation (COPPA, GDPR, and CCPA) that probed organizational processes to address each.

For recruitment, we scraped Google Play to collect email addresses of developers who had apps listed in the Designed for Families (DFF) program [46], the section specifically directed at children. To find additional child-directed apps, we also performed keyword searches, such as “toddler,” “kid,” “preschool,” and “children” to identify developers of apps that target child users. We identified 1,903 developers of child-directed Android apps, and received 50 responses from developers whose apps were available to users in California and the EU, therefore subjecting them to COPPA, GDPR's children's provisions, and potentially the CCPA.

We reviewed the responses to the 23 open-ended questions and then the first author built an initial codebook, which was shared with a second coder who of-

fered no changes. It consisted of 18 codes, selected based on the requirements of privacy laws and Google Play store policies (e.g., “parental-consent,” “children-data,” “opt-out,” “age-gates,” and “privacy-policies”). Since the responses across all open-ended questions frequently overlapped, we elected to use all 18 codes to code all the questions, rather than choosing distinct codes for each. Two coders independently coded all open-ended responses and then met to reconcile differences (Cohen's $\kappa = 0.74$, “substantial” agreement [51]), such that the resulting coded data represented unanimity.

4.2 Survey 2: Developer Perspectives

This survey consisted of 33 questions (Appendix B) across four sections to understand respondents':

- roles in the development of child-directed apps;
- personal experiences with the Designed for Families (DFF) program and Google Play Store policies;
- familiarity with COPPA, GDPR, and CCPA;
- resources for compliance guidance, and their confidence in their abilities to build compliant apps; and
- efforts to select and properly configure the privacy settings of third-party SDKs.

We also asked respondents about their challenges and the type of support they think they need from their organizations and the Google Play Store to make it easier for them to comply with applicable privacy regulations. We recruited participants through social media, the Google Play Store, and additionally used online search tools to find publicly-available email addresses of personnel working at organizations that develop child-directed apps (e.g., searching developers' websites to look for information on how to reach their technical and privacy contacts). We sent 3,202 invitation emails for the second survey, and received 77 responses from individuals involved in developing child-directed apps, corresponding to 72 different organizations. At the end of each survey, we asked whether we could contact participants in the future for follow-up questions; we conducted follow-up semi-structured interviews with those who agreed. For this survey, we built a codebook using the same process as in the first survey which used additional codes (e.g., “SDK-process,” “privacy-laws-resources,” “compliance-challenges,” and “Google-Play-compliance”). It was then used to code 16 open-ended questions. As before, two coders independently coded the responses, and then met to reconcile disagreements (Cohen's $\kappa = 0.88$, “almost perfect” agreement [51]).

4.3 Semi-Structured Interviews

A subset of survey respondents returned for follow-up interviews, during which we shared the results of our app analyses with them (§4.4). (Our second survey contained a question at the end asking respondents if we could follow up with them for interviews.) Our interview guide (Appendix C) consisted of questions about the processes participants' organizations follow to comply with privacy regulations, select and configure third-party SDKs, comply with the Google Play Store policies, and their experiences publishing apps.

We conducted 27 semi-structured interviews with individuals from 24 organizations, each lasting 60–90 minutes. For participants who consented to audio recording (all except two),³ we transcribed the recordings and the first author used inductive coding to build a new codebook. It was then used by the first author and one additional coder to independently code a subset of the data, and then they met to discuss their initial results and made any necessary updates to it. The interview codebook consisted of 40 codes that covered many themes, including privacy processes, third-party SDKs, and experiences with the Google Play Store. As before, the coders independently coded the responses and then met to reconcile differences (Cohen's $\kappa = 0.70$, “substantial” agreement [51]).

4.4 App Analysis

Prior to each interview session, we analyzed child-directed mobile apps published by participants' organizations to detect potential privacy issues (i.e., non-compliance with privacy regulations). We used their apps as a child would and then used dynamic analysis tools validated by prior work to examine apps' transmissions [64]. To ensure that we understand how apps behave when child users use them, we proceeded through age gates as a child would, and did not unlock any functionality accessible only to adult users.

We were interested in not just the exfiltration of users' personal data, but also whether any exfiltrations were accompanied by use limitations. As noted in prior work [64], many apps bundle third-party SDKs that provide configuration options that allow them to be deployed in compliance with various privacy laws.

Thus, for each third-party SDK identified in participants' apps, we researched how to configure it to comply with COPPA, CCPA, or GDPR. This included integrating these SDKs in test apps we developed, and then analyzing the network traffic to identify whether apps in our dataset were correctly using those configurations.

For example, Google's Admob SDK requires developers of child-directed apps to use a client-side function that sets a flag, `tag_for_child_directed_treatment`, in outbound traffic to indicate that the corresponding data should be treated in accordance with COPPA [16]. For CCPA, Admob also has another flag, `rdp`, which app developers can set to 1 to restrict the processing of data collected from their users (e.g., after an adult user has made a “do not sell request” or before a parent has consented) [15]. Other third-party SDKs, such as Unity Ads [28, 64], have options on their online dashboards that ask app developers whether their apps target children under the age of 13 and then send flags in inbound traffic—to the app—indicating that the app must comply with COPPA (e.g., Unity Ads sets a flag called `coppa` to true). Thus, we built a list of the privacy configuration options provided by each of the SDKs in question (Appendix E). Our analysis included Google Admob [15, 16], Adcolony [1], Vungle [30], Chartboost [8] and a few other popular advertising SDKs. Finally, we read each app's privacy policy and determined whether they accurately described the observed data flows. When there were inconsistencies between apps' observed behaviors and either their privacy policies or applicable laws, we presented participants with our findings during the interviews.

5 Results and Analysis

In our organizational survey, we asked organizations about their compliance processes and received responses from 50 organizations who collectively published more than 900 apps on Google Play (72% published < 20 apps, 14% published 20–39 apps, and 14% published > 40 apps). We received 77 responses to the second survey from individuals working in 72 organizations and who varied in terms of their roles and the number of apps they were involved in developing. Individual respondents in the second survey had roles in developing >10 apps (36%), 5–10 apps (17%), or 1–5 apps (43%), whereas the remaining 4% had no technical app development roles. As for their job functions, 78% were either software developers, test engineers, team leads, or prod-

³ We interviewed the two participants who did not consent to audio-recording and the notes we took during their interviews were subsequently used in our data analyses.

uct managers; 11% had leadership roles (e.g., CEOs); and the remaining 11% were either business/marketing employees, user-experience engineers, or customer service representatives. According to the DFF badging on Google Play, 56% of our interview participants and 72% of our survey respondents developed apps targeting kids who are up to 12 years of age. The rest developed apps that were not enrolled in DFF, but featured descriptions indicating the apps were directed at children.

5.1 Data Collection Procedures

In our organizational survey, we asked about the types of user data their apps collect and found that 38% of the 50 responding organizations reported none. Of the 31 organizations who reported collecting user data, they reported collecting: behavioral data about user interactions (42%), advertising IDs (35%), location data (32%), email addresses (16%), names of users (6%), IMEIs (6%), WiFi MAC addresses (3%), and other persistent identifiers (10%). Only 26% (of the 50) reported transmitting data to third parties, whereas 72% were certain that their apps did not (one was unsure). However, our app testing showed that 25% were incorrect about their apps not transmitting identifiers to third parties. The main purposes for data collection (open-ended) mentioned by the 31 respondents were:

1. Advertising and analytics: serving personalized ads, understanding user acquisition, analyzing app usability, collecting crash reports and getting insights about how to improve user engagement
2. User account creation/recovery or providing users with prior app usage data
3. Fulfilling developers' contractual obligations
4. Providing support for internal operations

Of those reporting collecting personal information, only 52% were “extremely” or “somewhat” confident that their apps always used encryption to do so (e.g., COPPA requires data be transmitted securely). In response to whether or not their privacy policies fully disclosed their apps' data collection and sharing behaviors, 92% answered affirmatively; dynamic analysis indicated that 17% of these respondents published child-directed apps that transmitted persistent identifiers to third-party domains, despite posting privacy policies that stated they do not collect children's identifiers.

We also found that certain interview participants' privacy policies either did not list all third parties receiving data from their apps, or said that an app is not

directed to kids when the Play Store description indicated otherwise. We discussed these issues with participants during the interviews to understand why their privacy policies were not fully transparent. We identified a few explanations: 1) one developer indicated that they offer apps directed to adults, as well, and they used the same privacy policies for their child-directed apps without modification, 2) another developer mentioned that their privacy policies were copied from other developers, and 3) two developers mentioned that the privacy policies were prepared by separate app publishers and they did not have visibility into their specifics.

One participant explained: “*We took it from other apps that have several million downloads and just changed the company name*”(C4).⁴ We also observed that participants incorrectly thought that certain SDKs were removed from all their apps in previous releases. Therefore, the takeaway from this section is that developers might not always have complete understandings of their apps' behaviors, which could affect their ability to make accurate disclosures in their privacy policies.

5.2 Consent Processes

Sixty percent of the 50 responses to the initial organizational survey indicated that their organizations believe they need to obtain parental consent before collecting data from child users. However, only 36% (of 50) reported doing so across *all* of their organizations' child-directed apps, prior to data collection (two reported doing so in *some* of their apps). In terms of *how* parental consent is obtained, 16% mentioned they use age gates to separate child users from adult users, 10% indicated they present users with knowledge-based questions that are not easy for kids to answer, 6% have certain features that are not available for children users, two ask for payment details that only a parent would be able to provide, one indicated that they ask all users to agree to their terms of service and/or privacy policy, and another indicated that they ask a parent to sign a consent form before letting their children use their apps.

We tested each organizations' apps from IP addresses in California to examine in-app consent screens, though did not observe any mechanisms for obtaining verifiable parental consent, including apps that we found sharing users' identifiers with third-party adver-

⁴ A: a respondent from the first survey, B: a respondent from the second survey, and C: interview participant.

tisers. We found that 26% of app developers hid certain features behind simple knowledge-based questions (e.g., basic arithmetic) that could likely be answered by children, whereas 6% displayed age gates that could be trivially bypassed by providing a birth year. None of these methods appear to meet the FTC’s requirements for “verifiable parental consent” [33].

Despite the absence of consent screens in the apps that we examined, many developers seemed to understand the importance of obtaining parental consent:

- “...can’t go into a legal agreement with a child.” (A12)
- “Because Google Play requires parental consent before collecting data from children.” (A9)
- “Most children could not understand what it means to give/deny consent to collect data. As such, if one were to collect data, parental consent would be the minimal ethical responsibility.” (A14)

Respondents who said that their organizations do not need parental consent explained:

- “Our app is not dangerous when used by children so there is no need for parental approval first.”(A34)
- “It’s listed in the app that it collects data and its listed in the privacy policy it’s doing so.”(A32)
- “Data collection is done by analytics SDK and it collects general data only, like location of devices.”(A23)
- “All apps in the Google Play store must be evaluated and rated before they are published.”(A22)

Only 28% of the 50 responding organizations indicated that they allow users to opt out of data collection, and 48% explicitly stated they do not. With regards to the specific opt-out mechanisms used, 14% indicated that they rely on the system-wide “Opt out of Ads Personalization” [14] setting, 10% indicated that child users are opted out by default (even though we observed that all of them had at least one app that transmitted persistent identifiers to third parties without parental consent and without having any obvious opt-out mechanism), 8% said they have opt-out settings within their apps, 4% said they use opt-out configurations offered by the SDKs integrated in their apps (we confirmed this for one developer by identifying the corresponding SDK configuration flags in their apps’ traffic; however, the other developer had an app that sent advertising IDs to AppsFlyer without opting child users out), and two said they allow users to contact them to opt out.

When running the apps corresponding to our interview participants, only a few had verifiable consent methods. When we explained what “verifiable consent”

means, many participants mentioned that they were not aware of that legal requirement:

- “I did not know about the thing where it makes sure that it is the parent and not the kid.”(C15)
- “actually this is very profitable for developers that it is not verified because children are clicking on ads all the time, .. and kids like this and they can click on purchase and you have a sale.”(C14)
- “...cannot even imagine how I would do it like give your consent that I can show you ads and then users immediately remove the app.”(C2)

Some participants stated that it is the responsibility of parents to review their apps before letting their kids use them. Others thought that the app stores should be responsible for explaining to developers how to obtain verifiable parental consent:

- “I don’t know how the child deals with that but I think that will be the responsibility of the parent.”(C10)
- “I think that is too big of a problem for our little studio to sort it out in a meaningful way. I think the logical step should be for the app stores to be responsible for developing sure ways to do this.”(C12)

The takeaway from this section is that many developers are not relying on verifiable methods for obtaining parental consent, and it appears that there is a general lack of awareness of this legal requirement.

5.3 Specific Privacy Laws

We report on developers’ awareness of their legal obligations. We believe those who participated in our study were subject to COPPA and GDPR, based on their apps’ availability in the EU and US. Since their apps were available in California, we asked them about CCPA, as well, but cannot be certain that they were subject to it (e.g., we do not know their revenues).

5.3.1 Awareness

Only 52% of the 50 organizations in the first survey stated that they have formal processes for addressing privacy issues during the development process, whereas 42% stated that they did not (6% preferred not to answer). Regarding familiarity with specific laws, 80%, 76%, and 52% stated their organizations were familiar with COPPA, GDPR, and CCPA, respectively. This is

Table 1. Reasons why respondents in the first survey believed that they were exempt from COPPA, GDPR, or CCPA.

Reason	COPPA	GDPR	CCPA
SDK companies take responsibility for protecting the received data	4%	8%	8%
Users can opt out using the system-wide “Limit Ad Tracking” setting	4%	6%	12%
We don’t collect data from our users	10%	4%	12%
We are covered under the internal operations exemption	4%	none	4%
Our apps are not targeted at kids	6%	6%	10%
Our apps don’t target users in US, EU/UK or California (one region was displayed for each regulation)	none	4%	2%

consistent with the 82%, 78%, and 62% of the organizational survey respondents who indicated they believed their apps are required to comply with these laws.

Those claiming exemptions used a multiple-choice question to provide justifications (Table 1). For those who indicated that their apps did not collect user data, we tested their apps and found that all but three were correct; all three published at least one app that transmitted identifiers to advertising networks. Another 10% indicated that kids were not among their target audiences, though we found that they either had apps that were listed in the “families” section of Google Play, at least one of their apps was described as child-directed, or they responded to another question in the survey that their app(s) were designed for children. Two respondents who claimed that data collection to support their organizations’ internal operations is exempt from CCPA—an exemption that only exists for COPPA—also indicated in their responses to another question that they are not familiar with the CCPA. Three respondents also mentioned that they were not required to comply with the GDPR or CCPA, because their apps are not directed at users in the EU or California. However, in all cases, we were able to download and use their apps from IP addresses in both places.

In contrast to the organizational survey, which asked whether organizations have processes for GDPR/COPPA/CCPA compliance, the second survey asked about individual developers’ awareness of these laws in their day-to-day work, as reported on a 5-point Likert scale (from “expert” to “unaware”). We additionally asked about respondents’ awareness of the California Privacy Rights Act (CPRA), which will replace

Table 2. Wilcoxon Signed Ranks tests results for participants’ familiarity with the different privacy regulations.

Hypothesis	(means, SDs), Z, effect size
More familiar with COPPA than GDPR ($p = 0.189$)	COPPA(3.26, 1.081), GDPR(3.40, 1.029), $Z = 1.31$, $r = 0.10$
More familiar with COPPA than CCPA ($p < 0.001$)	COPPA(3.26, 1.081), CCPA(2.12, 1.203), $Z = 5.91$, $r = 0.47$
More familiar with COPPA than CPRA ($p < 0.001$)	COPPA(3.26, 1.081), CPRA(1.74, 1.018), $Z = 6.74$, $r = 0.54$
More familiar with GDPR than CCPA ($p < 0.001$)	GDPR(3.40, 1.029), CCPA(2.12, 1.203), $Z = 6.16$, $r = 0.49$
More familiar with GDPR than CPRA ($p < 0.001$)	GDPR(3.40, 1.029), CPRA(1.74, 1.018) $Z = 6.79$, $r = 0.54$
More familiar with CCPA than CPRA ($p < 0.001$)	CCPA(2.12, 1.203), CPRA(1.74, 1.018), $Z = 3.67$, $r = 0.29$

the CCPA. To be able to evaluate respondents’ confidence in their knowledge of the laws in question, we added a fictitious privacy law, “FLIRPPA,” as a quality control. Among the 77 respondents, familiarity with COPPA and GDPR was higher than familiarity with CCPA and CPRA (see Figure 1). As expected, familiarity with FLIRPPA was lowest; however, 18 respondents claimed familiarity. Having 23% claiming familiarity with a fictitious law could demonstrate the level of uncertainty of developers in their compliance obligations, given the increasing number of laws that they are asked to comply with and that developers might not be based in the US or the EU. A Wilcoxon Signed Ranks test showed that differing levels of familiarity with the different laws were statistically significant, except when comparing COPPA to GDPR (see Table 2).

In the second survey, we also asked respondents to use an open-ended response to explain how they determine the privacy laws that are applicable to them:

- 36% primarily use information from Google Play;
- 21% said they either relied on their own online research or asked friends with more experience;
- 12% mentioned that they ignore privacy laws either because they assume that their apps are compliant or because they are small developers;
- 9% have consulted lawyers;
- 8% believed that they are required to comply with the laws applicable to the countries where their apps are available;
- 8% stated they follow specifications from app publishers or managers in their organizations; and
- around 3% mentioned working with third-party compliance services and said they trust their advice.

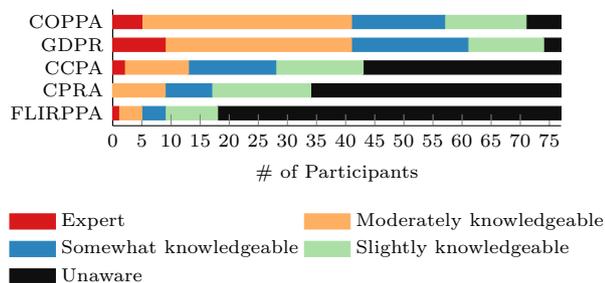


Fig. 1. Number of participants who described themselves as experts, knowledgeable, or unaware of the various privacy laws.

We used a follow-up multiple-choice question to quantify how frequently developers consult various information sources and found that 52% (of 77) rely on Google Play, 26% rely on other online sources, 16% have in-house counsel, and 9% consult external lawyers.

5.3.2 Understanding Obligations

We found that 80%, 70%, and 60% of the 50 responding organizations believed that all their apps were compliant with COPPA, GDPR, and CCPA, respectively. For each of these laws, we asked participants about their familiarity with their obligations and their organizations' compliance processes for addressing them.

When asked about COPPA requirements, open-ended responses included not collecting personal data (42%), using parental consent screens or age gates (14%), avoiding user targeting or tracking (10%), and posting privacy policies (4%). At least 12% did not know. On a subsequent page, we explained that COPPA applies to apps available to kids under the age of 13 in the US and that developers are required to obtain verifiable parental consent before collecting certain data through these apps. After receiving this prompt, 26% of developers stated that this information changed their beliefs about their obligations. Respondents also explained that they learn about their privacy obligations via enforcement performed by app markets:

- “Google routinely verify the respect of COPPA. They send notifications in case of violations, so we can fix them.”(A15)
- “Our app was pulled from the store, it was easier to change the target audience to 13 and above.”(A22)

As for respondents' self-described GDPR obligations, the main themes were not collecting data from users (34%), obtaining consent before collecting data (8%),

posting privacy policies (8%), and limiting tracking or advertising (8%). At least 16% either did not know what GDPR is or were not sure about what it required of them. As before, a subsequent page informed respondents that the children's provisions of GDPR (Article 8) apply to apps directed at kids in the EU under the age of 16. Subsequently, 30% said this information changed their beliefs about GDPR.

Finally, when we asked about CCPA, 28% said that it requires that they do not collect data, only 6% provided explanations related to the “right-to-know” requirements, and 4% indicated their apps should have age gates or age-appropriate ads. At least 34% either could not state any CCPA requirements or stated that they were not sure. As before, after subsequently explaining CCPA's requirements, over a third (36%) stated this information changed their understanding.

5.3.3 Compliance Processes

Forty-four percent (of 50 organizational survey respondents) stated that their approach to privacy compliance is to collect no data from users. Though, when we tested their apps, 27% transmitted advertising IDs to third parties, such as Unity Ads, Facebook, and Chartboost, and open-ended responses raised questions about developers' understandings of data collection generally:

- “We never collect any data from users and we always use Admob for ads.”(A16)
- “We do not collect personal data. We disclose the use of ads SDKs, and follow their guidelines.”(A15)
- “Not sure which ones [data types] exactly collected by ad networks we have used. I hope it is only Ad identifier.”(A2)

Consistent with prior findings, 18% rely on Google Play for guidance, with several indicating that their compliance processes consist of waiting to see if their apps are rejected:

- “...when the submission time comes we fill in those questionnaires. If it turns out we violate some regulation, we go back and fix it.”(A45)
- “Once we have a reject from Google Play, we fix it.”(A4)

As for access to legal resources, only 10% of the 50 organizations indicated that they get help from legal experts, whereas two said they relied on information from privacy-related news, conferences, or forums.

Only 8% of the 77 responses to our followup survey indicated that privacy compliance was not a challenge for them (see Appendix D). However, 68% of respondents stated that it is unlikely that they would be investigated by regulators if privacy issues are found in their apps, and 44% felt extremely or moderately concerned by the possibility of having their apps removed from the Play Store due to policy violations. We ran a Wilcoxon Signed Ranks test to examine whether participants were more concerned about the prospect of removal from Google Play than being investigated by privacy regulators and found the result to be statistically significant ($Z = 3.914$, $p < 0.001$, $r = 0.315$), which further highlights the major role of app stores in the privacy compliance ecosystem. As for whether they are concerned about the possibility of losing users due to publicized privacy issues or losing revenue if they avoid collecting personal data from their users, 38% and 31% were extremely or moderately concerned, respectively.

During the followup interviews, most correctly understood that they are required to comply with privacy regulations applicable to any region where their apps are available. However, they indicated that they use the store policies as their main source for guidance and trust that they will be notified if their apps have potential violations of relevant regulations. This was particularly prevalent among developers who work independently:

- “I don’t think that we’d think a lot about COPPA, it was just assumed that the Google store would approve and then we were okay with everything.”(C15)
- “...if my app is not complying with COPPA, GDPR or the Google Play policies, then my apps should not be on the store. I am not liable because Google is checking and so the responsibility is on Google.”(C9)
- “...I can confidently say that all my apps are compliant with all the standards because it was verified by Google before they published it.”(C26)

A participant claimed their company does not have to comply with COPPA, CCPA, or GDPR because they are not based in the US or EU—even though their apps are available in both—explained, “...we deal with Google and so Google is responsible for all the data”(C10).

Participants who worked with publishers or within large companies said that their superiors instruct them on what privacy features to implement or SDKs to integrate, and that even though their role is to ensure that apps do not have inappropriate content, they are not involved in privacy-related discussions:

- “We have our law team for that and have tons of professionals that know specific details about privacy, but I as a developer am not sure about what this process exactly is” (C18).
- “To be completely honest with you, there aren’t really any [process followed by the developer], it is totally handled by them [the publisher].” (C3)

Employees of large organizations said they rely on third parties, such as PRIVO [23] and kidSAFE [18], to audit their apps (e.g., recommending SDKs to remove and providing feedback on how to design consent screens and age gates):

- “Having a compliance agency is important. I don’t think that you’d ever have all the answers you need internally.” (C23)
- “...kidSAFE would do that for us. We have a yearly check with them that audits everything.”(C13)

Since our sample did not include a large number of developers whose apps were approved by these safe harbor programs, however, we leave investigating the efficacy of their auditing processes to future work.

Some mentioned contacts at major app markets who they consult whenever they need advice on how to comply, one explained: “we have a manager for our accounts at the Google store, they notify us of every policy change and give us an update on what we should do” (C17).

Others mentioned that their technical teams are separated from their privacy and legal teams, and that, if they were to be involved in privacy compliance decisions, they expect that to help them improve the privacy of their apps. One explained: “we have around 50 people and there is only one legal guy. I think that if everyone has knowledge, then it would help” (C23).

Participants who had legal teams highlighted that these teams are mostly responsible for writing their privacy policies, and helping engineers understand how to comply with the policies of the stores. One mentioned: “We have a legal team. Their responsibility is to comply as much as possible with current Google Play, and App-Store policies” (C17). Therefore, the general takeaway from this section is that even though the majority of developers have awareness that this is a regulated area, they lack complete understanding of what their compliance obligations are and rely on external feedback to determine whether they are in compliance.

5.4 Platform Policies

Apps that are directed to children are required to participate in Google's Designed for Families (DFF) program [46]. Eighty-three percent of the 77 who responded to the second survey indicated that they were aware of the DFF program and around 57% indicated that they developed apps that participate. Several respondents said they would prefer not to participate in DFF, if given a choice, for the following reasons: (1) not being able to use certain analytics or advertising services; (2) having to wait for a long time before their apps get approved for inclusion; (3) incurring additional costs to make apps compliant with DFF requirements; (4) not having enough clarity on how to make apps compliant with DFF; and (5) not being able to obtain clear enough explanations on why an app was removed from the store. For example, some respondents mentioned:

- *“With DFF, the adverts don't work well (50% of our revenues lost).”(B42)*
- *“Hard to follow all DFF rules. Google Play can disable app without clean explanation on what exactly was wrong.”(B62)*
- *“[The] main problem is the review time, Google take too long to review the app, could reach 5 to 7 days. Also due to legal requirements we can't use attribution solutions to analyze our marketing campaigns.”(B2)*
- *“Enrolling in the program incurs additional costs for our development process, and adds additional burden in terms of factors we have to consider, yet potentially questionable whether it has any positive impacts for us.”(B35)*

In contrast, others liked having their apps participate in DFF because they felt it provides credibility for their apps and gives them more visibility:

- *“The requirements are pretty simple when we already have a more strict policy.”(B30)*
- *“I want the store to be able to confidently share my app with kids, if participating in the DFF helps with this, then I am on board.”(B34)*
- *“Kids safety is very important, despite the revenue for us is dramatically low. This program should give some advantage to companies that really are taking care of the content.”(B17)*

Several app developers who are parents themselves perceived the importance of DFF:

- *“As a parent I think the DFF is a great idea. I wish it would be handled by an independent org.”(B45)*
- *“As a parent, and developer myself, I think it's great to have a separate set of rules for software companies that target under 13 users.”(B27)*
- *“Because as a dad of 4 year kid, I don't prefer that my child watch/play anything that is disturbing for him. I think Google policy is good for developers who want to make healthy apps for kids.”(B65)*

Our results suggest that for programs like DFF to become widely used and successful, app stores should enforce participation, while at the same time providing more useful guidance to developers.

5.5 Non-Compliance Warnings

As noted previously, developers rely on warnings from Google to inform them of privacy issues within their apps. Sixty-eight percent of the 77 respondents to our followup survey reported previously receiving these notifications, covering topics such as the following: 1) age-inappropriate content, 2) collection or sharing of personal data, 3) selecting an inappropriate target audience when submitting apps to the store, or 4) broken links to privacy policies or privacy policies lacking sufficient detail.

At the same time, 55% stated that app markets should take a greater role in evaluating mobile apps for compliance (only 5% said they should perform a smaller role) (see Appendix D). In our follow-up interviews, most of our participants said that they have received these notifications, and indicated challenges addressing them because they lacked detail, are received too often, include false positives that waste developers' time, and developers cannot readily obtain support from Google. One explained, *“they usually send us very standard texts like “you are not complying” and usually we have no idea about what the problem is”(C27).*

Two engineers working at companies that did not have dedicated teams to advise them explained:

- *“I got an app rejected because the message said there was an SDK that wasn't complying, but it did not say which one, I thought let me get rid of this one and then reapply and it got rejected again, and I would remove another one and reapply and it got rejected again, and it was tiring.”(C15)*
- *“...Apps gets rejected for reasons that do not make sense actually.”(C20)*

Other participants complained that app stores force them to seek outside legal advice: “If we talk about anything that has a potential COPPA implication, they say we cannot answer your question, go to your compliance agency and that’s just [a] pain when you are working closely on releases” (C23). This therefore shows that while getting app stores to notify developers about issues that exist in their apps can be a powerful tool to influence compliance rates, developers have to be provided with proper advice on what the issues are and how to fix them.

5.6 Developers’ Self-Efficacy

We asked developers in our second survey to rate their confidence in their abilities to:⁵

1. identify all types of data collected by their app(s) ($\mu = 3.64$, $\sigma = 1.327$);
2. identify all third parties receiving data from their app(s) ($\mu = 3.51$, $\sigma = 1.420$);
3. comply with applicable privacy regulations ($\mu = 3.82$, $\sigma = 0.996$);
4. configure third-party SDKs to be compliant ($\mu = 3.58$, $\sigma = 1.080$);
5. identify all privacy issues their apps have ($\mu = 3.55$, $\sigma = 0.967$);
6. fix privacy issues reported to them by external parties ($\mu = 4.03$, $\sigma = 0.959$);
7. envision how a privacy vulnerability that exists in their app(s) can be exploited ($\mu = 3.60$, $\sigma = 1.042$);
8. control the kinds of data accessed or transmitted by third-party SDKs ($\mu = 3.19$, $\sigma = 1.257$); and
9. limit the sharing of data by SDKs with other parties ($\mu = 3.17$, $\sigma = 1.250$).

The mean scores for the nine items in our scale show that most survey respondents were “Somewhat Confident” or “Confident” in their abilities to handle privacy issues. However, our interviews showed that many developers were not fully aware of the privacy practices of their own apps and relied mainly on the feedback received from the stores for determining whether their apps are sufficiently compliant. Because the nine items exhibited high internal consistency (Cronbach’s $=0.869$), we used them as a scale ($\mu = 32.08$, $\sigma =$

7.271). Using a Mann-Whitney U test, we examined whether respondents whose apps did not transmit personal data to third-party domains scored higher on our scale than those whose apps did, but we did not find a statistically significant difference ($Z = 0.552$, $p = 0.581$, $r = 0.045$). Similarly, those participating in DFF did not score significantly higher than those who did not ($Z = 0.319$, $p = 0.750$, $r = 0.026$).

5.7 Processes for Third-Party SDKs

This section summarizes our findings on how third-party SDKs are selected for inclusion, and the extent to which developers are familiar with SDK privacy settings.

5.7.1 Third-Party SDK Selection

In terms of the processes used to vet third-party SDKs, 22% of the 50 who responded to our organizational survey claimed to not have SDKs in their apps (our testing showed that they were correct except for two developers who integrated the AdColony and Adjust SDKs, which collected users’ advertising IDs), 20% said that they check the documentation about data collected by SDKs before integrating them, 10% stated that they only choose SDKs that are certified by Google, and 10% said they only work with trusted SDK companies:

- “We only work with big name SDK data collectors e.g. Facebook, Apps Flyer - and believe they comply with the law.”(A32)
- “We don’t incorporate 3rd party code because it is too difficult to be assured of their practices. Reading their privacy policies is rarely sufficient.”(A35)
- “We allow only Google SDK on our apps, which we believe they follow all regulations. But honestly, we cannot know what they do with the data.”(A6)

Absent from most responses were indications that they understood that both Google-approved SDKs and Google’s own SDKs are not necessarily compliant by default and require developers to still configure them for use in child-directed apps [45]. Only 6% mentioned using technical testing to identify the kinds of data collected by SDKs or to configure them for compliance (one of them published an app that transmitted advertising IDs to AppsFlyer). Furthermore, in our second survey, only 31% of 77 had processes to vet third-party SDKs, whereas 48% did not. When asked to explain how they decide whether to use an SDK or not in an

⁵ Respondents used a 5-point Likert scale ranging from “Very Confident” (5) to “Not Confident at All” (1).

open-ended response, several themes were identified: 1) independent developers doing their own research, asking friends and deciding based on the information obtained from available resources; 2) developers who used SDKs made available by big players they trust to be compliant; 3) small companies making decisions based on internal discussions, delegating these decisions to their CEOs or just relying on their developers to make their best judgements; and 4) medium-to-large companies delegating these tasks to a dedicated team. They explained:

- “We stick with well known players in the space.”(B4)
- “Developers suggest the SDKs. They need approval from CEO to include them.”(B51)
- “As we’re a small team, we discuss things and decide together based on the information collected by developers.”(B33)
- “I decide which SDK to include, my org lacks the technical experience and has a lot of juniors.”(B3)

In our interviews, we confirmed that developers, and particularly those who did not have dedicated teams to help, trusted that popular SDKs are legal to use. Additional factors in deciding whether to use an SDK included: 1) the functionality provided; 2) cost vs. expected revenue from using it; 3) ease of integration; and 4) availability of technical support.

Participants who had multiple teams said that such decisions are made by the app development team first, with the option to contact their legal teams later. Others said that engineers can propose an SDK, which then gets reviewed by other dedicated teams in order to approve whether to integrate it. In cases where there are different teams for different apps, each would have their own tech lead who handles SDK-related decisions. This shows that such decisions can be handled in a decentralized way, which may explain why a single organization might have apps with varying states of compliance. In our interviews, we also confirmed that well-resourced developers rely on third-party services for advice on SDK-related decisions: “We will run our proposed implementation by PRIVO just to make sure that we are sending data in a COPPA-compliant way, and we are not sharing identifiers of kids”(C23). Contract developers mentioned that SDKs often get chosen by their clients: “some clients told me they want Facebook Audience and so I must integrate it” (C10).

Those working with publishers mentioned that they are asked to integrate their publishers’ SDKs, and that they need to consult publishers before integrating additional SDKs (§5.8). As noted earlier, many partici-

pants appreciated having strict rules set by app markets, which they believe help them identify safe SDKs for child users: “...The stores force you to be compliant, which is great, because now you cannot mess up. They won’t publish unless you are compliant” (C8).

We also investigated how frequently third-party SDKs get updated by developers and found that 88% of the 77 respondents frequently (43%) or occasionally (45%) check for SDK updates after their initial integrations. Using a Spearman’s correlation test, we examined whether developers who are more confident of their abilities to control the types of data accessed by SDKs (Section 5.6) are more likely to make sure that SDKs integrated in their apps are up-to-date and found that there is a small positive correlation ($p = 0.041$, $\rho = 0.233$). In terms of releasing updates to fix privacy issues, 18% (of 77 respondents) claimed that they do so often or very often, whereas 52% rarely or never released privacy updates. We used a Mann-Whitney U test to examine whether those who claimed to release privacy updates are less likely to have sharing of personal data issues in their app(s) compared to respondents whose app(s) did not have issues; the difference was not statistically significant ($Z = 0.556$, $p = 0.578$, $r = 0.453$).

5.7.2 SDK Documentation

A majority of the 77 who responded to the second survey were confident or very confident that someone from their company has looked at all of their third-party SDKs’ privacy policies, terms of service, quickstart guides, or configuration documentation before integrating them. At least 58% of the 77 respondents also indicated that their organizations do check that the SDKs they use comply with applicable privacy regulations and their organizations’ privacy policies. Specifically, 18% mentioned that they look for information about the kinds of data collected by SDKs, 5% mentioned that they look for opt-out options and privacy flags that can be set to limit tracking children (one of them had an app that transmitted advertising IDs to Facebook without setting the corresponding privacy flag correctly), and a few others said that they use these to try to understand whether SDKs are compatible with their development frameworks. However, in our follow-up interviews, many of our participants mentioned that reading SDK documentation is not part of their process, except for a few who worked for large organizations. One explained: “I don’t read them because I rely on the fact that these SDKs are widely spread” (C4).

This is likely why many of our participants did not know that they needed to configure SDKs for compliance (§5.7.3). When we asked a participant why she did not use them, she said: “I was expecting that they do this automatically” (C2). Other participants mentioned that they only read parts of SDK documentation that are relevant to the app stores to make sure that their apps are not rejected. One participant justified why she does not proactively read SDK documentation: “If I have a problem, Google asks me to read them to fix the problem and in that case, I read them” (C10). Others mentioned that they look specifically for licensing and copyright information only, and some said that they tried reading SDK documentation, but found it difficult to understand.

5.7.3 SDK Misconfigurations

In terms of testing whether SDKs are configured correctly, only 5% of the 77 individual survey respondents indicated that they look at their apps’ network traffic to check for configuration issues. Most respondents indicated that they do functional or user experience tests or that they do not test SDKs after their initial integrations. Few interview participants mentioned that they consider looking at the types of data collected by third-party SDKs before deciding to integrate them; all who do work at large corporations. One participant explained how they modified an SDK: “we have access to the source code, we manually removed the IDFA-related code, and then we were able to use it” (C24). Another explained one of their test cases: “In order to trigger CCPA, we use a VPN to go to LA, and I see the logs from the device that are sent in real time” (C11).

Most of the 27 developers we interviewed had at least one third-party SDK privacy misconfiguration (§4.4). We discussed our findings with the participants and identified a number of reasons that might have led to these misconfigurations. The majority of them did not seem to know about the SDK privacy settings related to children generally or COPPA, GDPR or CCPA in particular. Although we had participants who claimed that they knew about some, but not all, of the SDK settings they were supposed to configure, they expressed that they expected the stores to reject their apps if they had any SDK-related issues. They assumed that their apps were compliant because the stores accepted them:

- “It sounds to me that this could be detected automatically by Google, like: ‘Hey, you have marked your app for kids but you did not configure your SDK properly.’” (C2)

- “...we didn’t use these functions before because they did not force us to use them.” (C10)
- “When the app being uploaded to their servers is using AppsFlyer, they could easily check if there is a consent form.” (C3)

Additionally, many participants indicated that the apps that had SDK configuration errors had not been updated for a long time, and that they lacked documentation to help other developers work on updating them, one mentioned: “...they are legacy, it is hard for us to know when this code exists for I don’t know how many years” (C24). Furthermore, when we told some of our participants that we found transmissions from their apps to certain SDKs’ servers, they indicated that they had stopped using these SDKs and thought that they were removed from their apps accordingly, which made them realize that they still needed to remove them. Other participants either: knew how to use some, but not all of the configurations that they were supposed to use; or used certain configurations correctly in some, but not all, of their apps. After raising the issue with them, they indicated that their teams do not have the time or resources that would allow them to keep checking whether SDK privacy options are used correctly:

- “...We are a small team, we try to develop an app and then we are busy with another one, and we end up in a circle running around.” (C6)
- “...We try our best but sometimes apps can stay on the market for a long time and not get updated and not get deleted and not comply with the policies.” (C17)

Furthermore, some participants indicated that their companies have different teams that are responsible for different apps, and that is another reason why some of their apps used SDK privacy options correctly, while they also have other apps that failed to use the same options. Others mentioned that they follow SDK integration instructions they receive from their superiors, and that privacy configurations were not part of what they received: “The publisher has very specific instructions on how to configure SDKs and I followed them to the letter. They didn’t mention this to me” (C1).

The general takeaway from this section is that vetting third-party SDKs for inclusion in child-directed apps and correctly configuring them for compliance are examples of tasks where developers could use better guidance, and that many look to app stores.

5.8 Work Models

Developers' work models have implications for how privacy-related responsibilities are handled. Some developers were fully responsible for their apps' development, release, and maintenance. Others worked with publishers who asked them to ensure that their apps are functioning correctly, by first releasing them on Google Play under their own accounts for testing purposes. Once the publishers are satisfied with how the apps perform, they acquire the apps, add consent screens, privacy policies, and any SDKs they need before releasing the apps under the publishers' accounts. Thus, apps are initially made available to users without sufficient measures to protect their privacy. One explained: *"the apps in our account, they are just tests. When the game gets to a stage where privacy becomes relevant, we would hand over the game to our publisher"* (C3). Participants believed that this helps them become compliant, since publishers handle privacy compliance to maintain their reputations.

Medium-size companies could be either startups, who have small development teams and no access to lawyers, or companies that are owned by larger companies that offer them access to their legal teams. Participants who followed this model mentioned that they have apps published under their own accounts, and they also work on a work-for-hire basis, where they develop apps and publish them under their clients' accounts. In these cases, privacy considerations could be handled by the publisher or the developer, depending on the terms of their contracts: *"If there are technical tasks to implement privacy, we take care of it, like permissions. Other times, we don't have any say in what they do"* (C21).

6 Discussion

Through this research we show that most developers lack a good understanding of the data collection behaviors of their own apps, and as a result, often end up violating privacy laws. Many of these issues stem from the use of third-party SDKs. While code reuse is a good and accepted practice, it adds a layer of complexity that often leads to "supply chain issues." Given developers' lack of awareness of the behaviors of various third-party SDKs, it is clear that they need access to more usable compliance-checking and auditing tools. The lack of proper guidance led developers to search for workarounds to satisfy app store requirements, such as using copied privacy policies, removing apps from

DFF or delegating compliance decisions to app publishers who lacked sufficient understanding of compliance requirements. Indeed, a common theme amongst small-to-medium size organizations was *the burden of compliance* and being ill-equipped to meet it:

- *"It is overwhelming for small businesses to check all the regulations. It is not difficult to comply if you know that you ticked all the boxes, but you need to understand which boxes you need to tick."* (C27)
- *"Nobody wants not to be compliant with this stuff, but everything requires technical development which can take time and effort."* (B6)
- *"It is a lot of work for small companies as we are limited in time & money for lawyers."* (A38)

We also observed that app stores play a substantial role in how developers both learn about relevant privacy laws and assess whether their apps are in compliance. For example, the fact that Google Play currently does not check whether child-directed apps obtain "verifiable" parental consent before collecting personal data for prohibited purposes could explain why developers were not aware of that legal requirement. The lack of a rigorous auditing process by the store left developers with the belief that their apps were in compliance when: 1) their methods to obtain parental consent could be bypassed by children, or 2) the third-party SDKs embedded in their apps were collecting identifiers without appropriate privacy controls. Nonetheless, developers believed that they were in compliance because they had not been told otherwise. Because participating in the DFF program was not actively enforced for child-directed apps, many developers also simply ceased participating in the program when they faced notices of noncompliance, rather than fixing the underlying issues.

Participants who also developed iOS apps found Apple's review process to be more thorough and recommended that Google adopt a similar approach. Prior work found that child-directed apps available on the iOS App Store shared data with fewer advertisers compared to Android apps [50], which could be due to differences in app review processes. App review should include a mandatory checklist that helps developers understand how to comply and provides pointers to SDK compliance options that need to be configured. While Google Play now provides a list of acceptable child-directed advertising SDKs, it does not enforce configuring them for compliance with relevant privacy laws, nor does it provide information about how to correctly do so. In the absence of this, developers appear to believe—

incorrectly—that by choosing SDKs from this list, compliance is a given. SDK providers should coordinate with app stores to enable providing timely and up-to-date advice to developers on how to exclude data collected from child users from being used for prohibited activities. This would relieve developers from having to search SDK documentation for compliance advice, which was shown to be challenging [68]. Since developers were found to keep using default settings of third-party SDKs [55], modifying these settings to be privacy-preserving (e.g., serving only contextual ads) would also help reduce the burden of compliance.

Developers require actionable guidance, but for that guidance to be productive, it has to not introduce more burden on developers by making sure that it is accurate, timely, and easy to understand. Notifying developers about issues that turn out to be false positives not only negatively impacts revenues and reputations, but might also lead developers to not act on future recommendations. Several developers stated that they wanted to be provided with guidance soon after they upload their apps to the stores, since they expressed difficulties with resolving privacy issues after having moved on to working on other apps or when some engineers are no longer available to maintain existing apps. They also preferred that app stores be more transparent about the criteria they use to audit apps so that they can adjust their processes accordingly: *“we’re forced to submit our app for review, get rejected several times, and figure out what’s wrong on our own”* (B70).

App store messaging at the various stages of the app submission process should therefore be leveraged to educate developers about their compliance responsibilities. While preparing apps for submission, developers could be presented with information about their compliance status with applicable laws and educated about the potential consequences of not complying. The stores could also reward developers by including badges that show their apps’ compliance status to users. Effective implementation of compliance indicators could motivate developers to work harder on improving the privacy of their apps in order to improve their reputations on the store and be assured that they are indeed in compliance.

In conclusion, while app developers are ultimately responsible for their apps’ compliance, far greater levels of compliance could be achieved if the central app stores provided both better guidance and enforcement, as they are best positioned to do so. While we offer recommendations about how centralized app stores could offer better guidance and centralized compliance check-

ing, future research is needed to iteratively examine *how* best to offer this guidance.

Limitations. This work was not without limitations. The most obvious of which is our reliance on self-reported data (as is the case with all surveys and interviews). While the response rate to our unsolicited emails was relatively low, it was consistent with response rates to other unsolicited surveys. Even though our sample included individual developers as well as developer organizations of various sizes, it is possible that those who chose to respond were not representative of all developers of child-directed Android apps. To compare our sample to the broader developer population, we selected a random sample of developers of child-directed apps who did not participate in our study, and then tested one app per developer, which we selected at random. We compared these results with one app selected at random from the developers who participated. We found that 14.6% of the apps included in the random sample contained potential issues (95% CI=8.9%–22.1%), whereas we identified potential issues with 26% of study participants’ apps (95% CI=18.5%–34.7%). Despite the appearance of more potential privacy issues amongst our participants’ apps, the developers who participated were amongst the more popular developers of child-directed apps. This could suggest that popular apps contain more privacy issues, but a measurement study that includes a much larger sample of apps is needed to validate this relationship, which we leave to future work.

Our results should not be interpreted as suggesting prevalence of specific developer privacy misconceptions or behaviors. We believe that our sample is representative of the broader developer population, but much more insight would have been uncovered had we recruited more individuals within each of the developer organizations that participated. This was not possible, however, given how challenging it is to recruit all those who are involved in compliance processes. We promised to analyze collected data anonymously, but given that privacy compliance is considered a sensitive topic, it is likely that many developers preferred not to participate. For the same reason, it is also possible that some of those who participated could have provided inaccurate information that cannot be validated by analyzing their apps or reading their privacy policies. Our observations can therefore be used as likely explanations of how privacy compliance is being handled by developers of child-directed apps, though it remains unclear how representative these explanations are of the industry as a whole.

Acknowledgements

We thank Hailey Jang and Ruba Abu-Salma for their help with qualitatively coding the study data and their feedback on our study instruments, and Edmund Wang for his help with testing third-party SDKs. We are also grateful to Jordan Fischer for her valuable comments on the legal part of our work. Our thanks extends to our shepherd, Yaxing Yao and the anonymous reviewers for their insightful feedback on our paper. This work was supported by NSA contract H98230-18-D-0006, NSF grant CNS-1817248, and a gift from Google.

References

- [1] Adcolony - Privacy Laws Settings. <https://github.com/AdColony/AdColony-Android-SDK/wiki/Privacy-Laws>. [Online; accessed: 19-November-2021].
- [2] Advertising Platform OpenX Will Pay \$2 Million for Collecting Personal Information from Children in Violation of Children's Privacy Law. <https://www.ftc.gov/news-events/press-releases/2021/12/advertising-platform-openx-will-pay-2-million-collecting-personal>. [Online; accessed: 24-January-2022].
- [3] Amplitude - Android SDK. <https://www.docs.developers.amplitude.com/data/sdks/android/>. [Online; accessed: 05-July-2022].
- [4] Appodeal - Android documentation. <https://wiki.appodeal.com/en/android/get-started>. [Online; accessed: 05-July-2022].
- [5] Article 6 of GDPR. <https://gdpr-info.eu/art-6-gdpr/>. [Online; accessed: 02-January-2022].
- [6] Article 8 of GDPR. <https://gdpr-info.eu/art-8-gdpr/>. [Online; accessed: 27-February-2022].
- [7] California Consumer Privacy Act (CCPA). <https://oag.ca.gov/privacy/ccpa>. [Online; accessed: 02-November-2021].
- [8] Chartboost - Android Privacy Methods. https://answers.chartboost.com/en-us/child_article/android-privacy-methods. [Online; accessed: 19-November-2021].
- [9] Developer of Apps Popular with Children Agrees to Settle FTC Allegations It Illegally Collected Kids' Data without Parental Consent. <https://www.ftc.gov/news-events/press-releases/2020/06/developer-apps-popular-children-agrees-settle-ftc-allegations-it>. [Online; accessed: 17-November-2021].
- [10] Facebook - COPPA. <https://developers.facebook.com/docs/audience-network/optimization/best-practices/coppa/>. [Online; accessed: 05-July-2022].
- [11] Flurry - Android SDK. <https://developer.yahoo.com/flurry/docs/integrateflurry/android/>. [Online; accessed: 05-July-2022].
- [12] GDPR - Consent. <https://gdpr-info.eu/issues/consent/>. [Online; accessed: 01-November-2021].
- [13] GDPR - Recital 38. <https://gdpr-info.eu/recitals/no-38/>. [Online; accessed: 01-November-2021].
- [14] Google - Advertising ID. <https://support.google.com/googleplay/android-developer/answer/6048248#zippy=%2Ctargeting-devices-without-an-advertising-id%2Cchow-to-opt-out-of-personalized-ads>. [Online; accessed: 31-January-2022].
- [15] Google Admob: CCPA Preparation. <https://developers.google.com/admob/android/ccpa>. [Online; accessed: 19-November-2021].
- [16] Google Admob: Child-Directed SDK Settings. https://developers.google.com/admob/android/targeting#child-directed_setting. [Online; accessed: 19-November-2021].
- [17] ironSource - Android SDK - Regulation Advanced Settings. <https://developers.is.com/ironsource-mobile/android/regulation-advanced-settings/#step-1>. [Online; accessed: 05-July-2022].
- [18] kidSAFE. <https://www.kidsafeseal.com/aboutourprogram.html>. [Online; accessed: 21-February-2022].
- [19] Mobfox - Android SDK. <https://mobfox.atlassian.net/wiki/spaces/PUMD/pages/1205239832/Integrate+via+Android+SDK>. [Online; accessed: 05-July-2022].
- [20] Mobile Advertising Network InMobi Settles FTC Charges It Tracked Hundreds of Millions of Consumers' Locations Without Permission. <https://www.ftc.gov/news-events/press-releases/2016/06/mobile-advertising-network-inmobi-settles-ftc-charges-it-tracked>. [Online; accessed: 17-November-2021].
- [21] Mobile Apps for Kids: Current Privacy Disclosures are Disappointing. https://www.ftc.gov/sites/default/files/documents/reports/mobile-apps-kids-current-privacy-disclosures-are-disappointing/120216mobile_apps_kids.pdf. [Online; accessed: 05-January-2022].
- [22] Mobile Apps for Kids: Disclosures Still Not Making the Grade. <https://www.ftc.gov/sites/default/files/documents/reports/mobile-apps-kids-disclosures-still-not-making-grade/121210mobilekidsappreport.pdf>. [Online; accessed: 05-January-2022].
- [23] PRIVO. <https://www.privo.com/welcome>. [Online; accessed: 21-February-2022].
- [24] Protecting Children's Privacy Under COPPA: A Survey on Compliance. <https://www.ftc.gov/sites/default/files/documents/reports/protecting-childrens-privacy-under-coppa-survey-compliance/coppasurvey.pdf>. [Online; accessed: 05-January-2022].
- [25] Tapjoy - Android SDK. <https://dev.tapjoy.com/en/android-sdk/Quickstart>. [Online; accessed: 05-July-2022].
- [26] The General Data Protection Regulation (GDPR). <https://gdpr-info.eu/>. [Online; accessed: 01-November-2021].
- [27] TITLE 1.81.5. California Consumer Privacy Act of 2018. https://leginfo.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5. [Online; accessed: 02-November-2021].
- [28] Unity Ads - COPPA Compliance. <https://docs.unity3d.com/Manual/UnityAnalyticsCOPPA.html>. [Online; accessed: 19-November-2021].
- [29] User opt-in/opt-out in the AppsFlyer SDK. <https://support.appsflyer.com/hc/en-us/articles/360001422989-User-opt-in-opt-out-in-the-AppsFlyer-SDK>. [Online; accessed: 27-January-2022].

- [30] Vungle - Advanced Settings. <https://support.vungle.com/hc/en-us/articles/360047780372>. [Online; accessed: 19-November-2021].
- [31] Yandex - Android SDK. <https://yandex.com/dev/mobile-ads/doc/android/quick-start/gdpr-about.html>. [Online; accessed: 05-July-2022].
- [32] Children's Privacy. <https://www.ftc.gov/tips-advice/business-center/privacy-and-security/children%27s-privacy>, 2020. [Online; accessed: 31-October-2021].
- [33] Complying with COPPA: Frequently Asked Questions. <https://www.ftc.gov/tips-advice/business-center/guidance/complying-coppa-frequently-asked-questions-0>, 2020. [Online; accessed: 31-October-2021].
- [34] New Rule to Protect Children's Online Privacy Takes Effect April 21, 2000. <https://www.ftc.gov/news-events/press-releases/2000/04/new-rule-protect-childrens-online-privacy-takes-effect-april-21>, 2020. [Online; accessed: 26-January-2022].
- [35] Noura Alomar, Primal Wijesekera, Edward Qiu, and Serge Egelman. "You've Got Your Nice List of Bugs, Now What?" Vulnerability Discovery and Management Processes in the Wild. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, pages 319–339, 2020.
- [36] Noah Aporthe, Sarah Varghese, and Nick Feamster. Evaluating the contextual integrity of privacy regulation: Parents' IoT toy privacy norms versus COPPA. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 123–140, 2019.
- [37] Hala Assal and Sonia Chiasson. 'Think secure from the beginning' A Survey with Software Developers. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–13, 2019.
- [38] Kenneth A Bamberger and Deirdre K Mulligan. *Privacy on the ground: driving corporate behavior in the United States and Europe*. MIT Press, 2015.
- [39] Ravi Bhoraskar, Seungyeop Han, Jinseong Jeon, Tanzirul Azim, Shuo Chen, Jaeyeon Jung, Suman Nath, Rui Wang, and David Wetherall. Brahmastra: Driving apps to test the security of third-party components. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 1021–1036, 2014.
- [40] 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*, pages 23–31, 2018.
- [41] Constanze Dietrich, Katharina Krombholz, Kevin Borgolte, and Tobias Fiebig. Investigating system operators' perspective on security misconfigurations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1272–1289, 2018.
- [42] 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 *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2021.
- [43] Ming Fan, Le Yu, Sen Chen, Hao Zhou, Xiapu Luo, Shuyue Li, Yang Liu, Jun Liu, and Ting Liu. An empirical evaluation of gdpr compliance violations in android mhealth apps. In *2020 IEEE 31st International Symposium on Software Reliability Engineering (ISSRE)*, pages 253–264. IEEE, 2020.
- [44] Álvaro Feal, Paolo Calciati, Narseo Vallina-Rodriguez, Carmela Troncoso, and Alessandra Gorla. Angel or devil? a privacy study of mobile parental control apps. *Proc. Priv. Enhancing Technol.*, 2020(2):314–335, 2020.
- [45] Google. Comply with google play's families policy using admob. <https://support.google.com/admob/answer/6223431?hl=en>, 2022.
- [46] Google. Participate in designed for families. <https://support.google.com/googleplay/android-developer/answer/9859555?hl=en>, 2022.
- [47] Catherine Han, Irwin Reyes, Álvaro Feal, Joel Reardon, Primal Wijesekera, Narseo Vallina-Rodriguez, Amit Elazar, Kenneth A Bamberger, and Serge Egelman. The price is (not) right: Comparing privacy in free and paid apps. *Proceedings on Privacy Enhancing Technologies*, (3), 2020.
- [48] Julie Haney, Yasemin Acar, and Susanne Furman. "It's the Company, the Government, you and I": User perceptions of Responsibility for Smart Home Privacy and Security. In *30th USENIX Security Symposium (USENIX Security 21)*, 2021.
- [49] Konrad Kollnig, Reuben Binns, Pierre Dewitte, Max Van Kleek, Ge Wang, Daniel Omeiza, Helena Webb, and Nigel Shadbolt. A fait accompli? an empirical study into the absence of consent to third-party tracking in android apps. *Proceedings of the 7th Symposium on Usable Privacy and Security (SOUPS 2021)*, 2021.
- [50] Konrad Kollnig, Anastasia Shuba, Reuben Binns, Max Van Kleek, and Nigel Shadbolt. Are iphones really better for privacy? a comparative study of ios and android apps. *Proceedings on Privacy Enhancing Technologies*, 2022(2), 2022.
- [51] J. Richard Landis and Gary G. Koch. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174, 1977.
- [52] Frank Li, Lisa Rogers, Arunesh Mathur, Nathan Malkin, and Marshini Chetty. Keepers of the machines: Examining how system administrators manage software updates for multiple machines. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 273–288, 2019.
- [53] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason I Hong. How developers talk about personal data and what it means for user privacy: A case study of a developer forum on reddit. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3):1–28, 2021.
- [54] Minxing Liu, Haoyu Wang, Yao Guo, and Jason Hong. Identifying and analyzing the privacy of apps for kids. In *Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications*, pages 105–110, 2016.
- [55] 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. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 225–244, 2019.
- [56] Kristopher Micinski, Daniel Votipka, Rock Stevens, Nikolaos Kofinas, Michelle L Mazurek, and Jeffrey S Foster. User interactions and permission use on android. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 362–373, 2017.
- [57] Hooman Mohajeri Moghaddam, Gunes Acar, Ben Burgess, Arunesh Mathur, Danny Yuxing Huang, Nick Feamster, Edward W Felten, Prateek Mittal, and Arvind Narayanan.

- Watching you watch: The tracking ecosystem of over-the-top tv streaming devices. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 131–147, 2019.
- [58] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. Privacy expectations and preferences in an iot world. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 399–412, 2017.
- [59] Trung Tin Nguyen, Michael Backes, Ninja Marnau, and Ben Stock. Share first, ask later (or never?)-studying violations of gdpr's explicit consent in android apps. In *USENIX Security Symposium*, 2021.
- [60] Ehimare Okoyomon, Nikita Samarin, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodriguez, Irwin Reyes, Álvaro Feal, Serge Egelman, et al. On the ridiculousness of notice and consent: Contradictions in app privacy policies. In *Workshop on Technology and Consumer Protection (ConPro 2019)*, in conjunction with the 39th IEEE Symposium on Security and Privacy, 2019.
- [61] Mariana Peixoto, Dayse Ferreira, Mateus Cavalcanti, Carla Silva, Jéssyka Vilela, João Araújo, and Tony Gorschek. On understanding how developers perceive and interpret privacy requirements research preview. In *International Working Conference on Requirements Engineering: Foundation for Software Quality*, pages 116–123. Springer, 2020.
- [62] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, Phillipa Gill, et al. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. In *The 25th Annual Network and Distributed System Security Symposium (NDSS 2018)*.
- [63] Joel Reardon, Álvaro Feal, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodriguez, and Serge Egelman. 50 ways to leak your data: An exploration of apps' circumvention of the android permissions system. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 603–620.
- [64] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, Serge Egelman, et al. "Won't somebody think of the children?" examining COPPA compliance at scale. In *The 18th Privacy Enhancing Technologies Symposium (PETS 2018)*.
- [65] Katie Shilton and Daniel Greene. Linking platforms, practices, and developer ethics: Levers for privacy discourse in mobile application development. *Journal of Business Ethics*, 155(1):131–146, 2019.
- [66] Kaiwen Sun, Carlo Sugatan, Tanisha Afnan, Hayley Simon, Susan A Gelman, Jenny Radesky, and Florian Schaub. "They See You're a Girl if You Pick a Pink Robot with a Skirt": A Qualitative Study of How Children Conceptualize Data Processing and Digital Privacy Risks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–34, 2021.
- [67] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. Privacy champions in software teams: Understanding their motivations, strategies, and challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2021.
- [68] Mohammad Tahaei, Kopo M Ramokapane, Tianshi Li, Jason I Hong, and Awais Rashid. Charting app developers' journey through privacy regulation features in ad networks. *Proceedings on Privacy Enhancing Technologies*, 1:24, 2022.
- [69] Mohammad Tahaei, Kami Vaniea, and Naomi Saphra. Understanding privacy-related questions on stack overflow. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2020.
- [70] Christian Tiefenau, Maximilian Häring, Katharina Kromholz, and Emanuel von Zezschwitz. Security, availability, and multiple information sources: Exploring update behavior of system administrators. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*, pages 239–258, 2020.
- [71] Daniel Votipka, Seth M Rabin, Kristopher Micinski, Thomas Gilray, Michelle L Mazurek, and Jeffrey S Foster. User comfort with android background resource accesses in different contexts. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, pages 235–250, 2018.
- [72] Jice Wang, Yue Xiao, Xueqiang Wang, Yuhong Nan, Luyi Xing, Xiaojing Liao, JinWei Dong, Nicolas Serrano, Hao-ran Lu, XiaoFeng Wang, et al. Understanding Malicious Cross-library Data Harvesting on Android. In *30th USENIX Security Symposium (USENIX Security 21)*.
- [73] Charles Weir, Ben Hermann, and Sascha Fahl. From needs to actions to secure apps? the effect of requirements and developer practices on app security. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 289–305, 2020.
- [74] Flynn Wolf, Adam J Aviv, and Ravi Kuber. Security obstacles and motivations for small businesses from a CISO's perspective. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1199–1216.
- [75] Leah Zhang-Kennedy and Sonia Chiasson. "Whether it's moral is a whole other story": Consumer perspectives on privacy regulations and corporate data practices. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, pages 197–216, 2021.
- [76] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel R Reidenberg, N Cameron Russell, and Norman Sadeh. Maps: Scaling privacy compliance analysis to a million apps. In *The 19th Privacy Enhancing Technologies Symposium (PETS 2019)*, 2019.

A Survey 1

This section lists all the questions we included in the survey that we sent to developer organizations.

Section 1: Data Collection and Transmission

1. Approximately how many mobile apps does your organization develop or distribute that are available through Google Play?
2. Are any of these apps designed for children?
3. Do any of your apps transfer data to third parties (e.g., ad networks or other domains)?

4. If Yes, please list the third parties that receive data from your apps.
5. What types of data do you collect from the users of your apps?
6. If your app collects data from users, for what purposes do you use the data that you collect? (If your app(s) do not collect data from users, please indicate that in your answer)
7. When your app transmits data over the Internet, how confident are you that it always uses encryption (i.e., SSL, TLS)?
8. Do your apps have privacy policies?
9. Do your app(s)' privacy policies mention all the data collection and sharing behaviors of your app(s)?
10. Do your app(s) obtain parental consent before collecting certain types of data from child users?
11. If Yes, how do your apps obtain parental consent before collecting users' data?
12. Does your organization think that parental consent needs to be obtained before collecting data from children users?
13. Do any of your apps use age gates to separate child users from adult users?
14. If Yes, how is data from child users treated differently, after a child user is identified using your age gate?
15. Does your app obtain users' consent prior to sending personal information to third parties?
16. If Yes, what mechanism do you use to obtain user consent?
17. Does your app allow users to opt out of the collection of their personal information?
18. If Yes, what mechanism(s) does your app use to allow users to opt out of data collection?

Section 2: Privacy Compliance Processes

1. Does your organization have a formal process for addressing privacy issues during the software development process?
2. Is your organization familiar with any of the following laws? (Choices: COPPA, CCPA, GDPR, None)
3. Which of the following laws does your organization believe your apps are subject to compliance with? (Choices: COPPA, CCPA, GDPR, None)
4. If your organization believes that your app is exempt from COPPA/GDPR/CCPA, which of the following does your organization think apply to your app(s)? (Choices: Our app(s) is not directed to kids, The users can just opt out by enabling the Opt out

of Ads Personalization on their phones, Our app does not target users in the US, We are allowed under the law to collect data for our internal operations [please explain], We believe that SDK companies are responsible for protecting data they receive from our app, Other [please explain])

5. Please describe the processes your organization has in place for complying with various privacy regulations.
6. When your organization integrates third-party code into its apps (e.g., SDKs), what processes does it follow to ensure that those third-party components do not violate your organization's privacy obligations?
7. Do you offer both free and paid versions of your app(s)?
8. If Yes, does your organization treat paid apps differently than free apps in terms of the types of data you collect from users or the types of privacy features offered in your apps?

Section 3: Privacy Laws - COPPA

1. Does your organization believe that your apps are compliant with COPPA?
2. What obligations does your organization believe it has under COPPA?
3. COPPA applies to apps that are directed to kids (under the age of 13) in the US. Therefore, your apps may be subject to COPPA if you make them available to kids in the US through Google's Play store. Under COPPA, developers are required to obtain verifiable parental consent before collecting certain types of data from kids. Does this change your organization's beliefs/understanding about your organization's obligations under COPPA?

Section 4: Privacy Laws - GDPR

1. Does your organization believe that your apps are compliant with GDPR?
2. What does your organization believe their obligations are under GDPR?
3. GDPR applies to all users in Europe, and apps or websites that target users under the age of 16 are subject to the kids provisions of GDPR. Therefore, if you offer your app(s) to users under the age of 16 in any European country through Google's Play store, then your app(s) may be subject to compliance with GDPR. Does that change your organization's beliefs about your organization's obligations under GDPR?

Section 5: Privacy Laws - CCPA

1. Does your organization believe that your apps are compliant with CCPA?
2. What does your organization believe their obligations are under CCPA?
3. CCPA applies to all users in California, and apps or websites that target users under the age of 16 may be subject to the kids provisions of CCPA. Therefore, if your app is targeting children and is available to users in California through Google's Play store, then it may be subject to compliance with CCPA. Under CCPA, verifiable parental consent must be obtained in order to opt-in children users (whose ages are under 13) to the sale of their personal information. Does that change your organization's beliefs about your organization's obligations under CCPA?
2. If Yes, please explain the reasons why Google thought that your app was not compliant with Play Store policies, including what policy they believed your app violated.
3. If you were asked to list all the third parties that get contacted by your apps, how confident are you that you would be able to identify all of them?
4. If you were asked to list all the types of personal data that your app(s) collect from users (including persistent identifiers), how confident are you that you will be able to identify all of them?
5. How knowledgeable are you about each of the following regulations: COPPA, GDPR, CCPA, CPRA and FILRPPA
6. How do you determine what laws/regulations are applicable to the mobile apps that you have a role in developing?

Section 6: Final Reflections

1. Based on your organization's experience, what would make it easier for app developers to comply with privacy regulations?
2. Is there anything else that you would like to add?

B Survey 2

This section lists the questions we sent to personnel within developer organizations.

Section 1: General Questions

1. Which of the following roles best describes your job function?
2. How many mobile apps have you had a role in developing?
3. Do you individually develop or have you ever worked at an organization that develops mobile apps directed at kids under the age of 13?
4. Are you aware of the Designed for Families (DFF) program in the Google Play store?
5. Have you contributed to the development of any mobile apps that participate in the DFF program, now or previously?
6. If Yes: if given the choice, would you prefer to not participate in the DFF program?

Section 2: Privacy Regulations

1. Have you (or someone else at your organization) ever been notified by Google that one of your apps is not compliant with the Play Store's policies?
7. Which of the following resources do you rely on for guidance on how to comply with applicable privacy regulations? (please select all that apply) (Choices: In-house legal counsel, Online sources (e.g., websites and social media), Regular advice from external lawyers, The Google Play store, None of the above, Other [please explain])
8. How confident are you of your ability to comply with privacy regulations?
9. How confident are you that all third-party code embedded in your app(s) is correctly configured to comply with applicable privacy regulations?
10. How confident are you of being able to identify all the potential privacy issues that exist in your app(s)?
11. How confident are you of your ability to fix a privacy issue found in your app by an external tester?
12. How confident are you of your ability to envision the risks your users might experience should they exploit a privacy vulnerability that exists in your app (e.g., an attacker using location data collected by your app)?
13. How challenging is it to make your app(s) compliant with privacy regulations?
14. Please explain your answer to the previous question.
15. What role should central app markets, like the Google Play store or Apple App Store, take in determining whether mobile apps comply with applicable privacy regulations? (Choices: Smaller role, Stay the same, Larger role)
16. What, if anything, should central app markets, like the Google Play store or Apple App Store, do dif-

ferently to help developers comply with applicable privacy regulations?

17. How likely do you think you or your company are to be investigated by privacy regulators as a result of violating privacy regulations?
18. How frequently do you check for whether SDKs integrated in your app(s) have released new versions after your initial integrations?
19. How often do you release updates aimed at fixing a privacy issue found in your app?
20. How concerned are you about the possibility of your app(s) being removed from the Play Store as a result of a discovered privacy issue?
21. How concerned are you about the possibility of experiencing a decline in the number of users who use your app(s) once they learn about a privacy issue (e.g., unexpectedly sharing their data with third parties)?
22. How concerned are you about the possibility of experiencing a decline in revenue if you must limit your apps' data collection practices?

Section 3: SDKs

1. Does your organization have a formal process that it uses to vet potential SDKs for inclusion in your app(s)?
2. Please explain the process your organization uses to decide whether or not to include a given SDK. For example, is this a decision left to individual developers? Is there a designated person or department that vets new SDKs?
3. Once a decision has been made to include a new SDK, is there anything that you or your organization does to ensure that the SDK has been configured correctly?
4. When integrating an SDK in your app(s), how confident are you that you or someone within your organization has read that SDK's: (a) Privacy policy?, (b) Terms of service?, (c) Quickstart guide? and (d) Configuration documentation?
5. What kind of information do you look for in SDK documentation when trying to make your app(s) compliant with privacy regulations?
6. Is there someone within your organization who is responsible for ensuring that the use of third-party SDKs complies with your organization's privacy policy? Please explain your answer.
7. Is there someone within your organization who is responsible for ensuring that the use of third-party

SDKs complies with applicable laws? Please explain your answer.

8. How confident are you of your ability to control the types of data accessed and/or transmitted by SDKs integrated in your app(s)?
9. How confident are you of your ability to limit the transmission of personal data collected by an SDK integrated in your app with other third parties?

Section 4: Challenges and Recommendations

1. Briefly summarize the challenges you experience when trying to make changes in your app(s) to comply with privacy regulations.
2. What kind of support from the Google Play store do you think would make it easier for you to make your app(s) compliant with privacy regulations?
3. What kind of support from your organization's leaders do you think would make it easier for you to develop apps that are compliant with privacy regulations?
4. What kind of support do you need from other people in your organization to make it easier for you to develop apps that are compliant with privacy regulations?

C Interview Questions

Below, we list the questions we used for our semi-structured interviews.

1. Are you aware of the following privacy regulations: COPPA, GDPR, CCPA?
2. If yes, do you believe that your child apps are compliant with COPPA, CCPA, and/or GDPR? Please explain.
3. Can you sketch the processes you and your team follow to comply with COPPA, CCPA, and/or GDPR?
4. Do these processes change when designing apps specifically targeted at children? If yes, how?
5. Would you change the design or add/remove features when designing apps specifically targeted at children? Please explain.
6. Have you ever been told by Google that your app is not compliant with the Play Store rules?
7. What tools do you use to test and/or monitor your compliance with COPPA, CCPA, and/or GDPR? How effective are these tools in helping you comply with these privacy regulations?

8. Do you think that you have the resources and/or (technical and legal) expertise that could help you become compliant with COPPA, CCPA, and/or GDPR?
9. Please explain what you think you would need to improve the compliance of your apps with these privacy regulations.
10. Do you believe it is difficult to comply with privacy regulations such as COPPA, CCPA and/or GDPR? What are the challenges that you experience when trying to comply with them?
11. What kind of third-party code (SDKs) do you integrate into the child apps you develop?
12. Can you explain the reasons why you need to integrate SDKs in your apps?
13. How do you decide whether it is ok to integrate an SDK in your apps or not? Can you explain your thought/decision process?
14. Do you read the privacy policies of SDKs before deciding to integrate them in your apps? If yes, what kind of information do you look for in such privacy policies? And do you find it helpful to do so?
15. Once you read SDK documentations, do you do any kind of testing to determine what kind of data is collected by SDKs?
16. Do you know the types of data collected by each SDK you integrate in your apps?
17. What challenges does your organization face when you integrate SDKs in your apps?
18. Are there any other third-party services that you use in your apps? And if so, how do you make sure that they are COPPA, CCPA, and/or GDPR compliant before you start using them?
19. What kind of questions do you consider asking third-party providers before you use their services?
20. Please provide some examples of privacy concerns that were raised to you by end-users (if any).
21. What are the reasons that may have caused the privacy concerns or issues we have just discussed?
22. What kind of support do you think developers need in order to comply with privacy regulations (COPPA, CCPA, GDPR)?
23. What do you think the implications are of not complying with COPPA, CCPA, and GDPR?
24. Are you concerned about violating the policies of Google Play?
25. Did you know that your app collects this kind of data from end users? Please explain the purpose of this data collection.

26. Did you know that your app communicates with these domains? Please explain the purpose of such communication.
27. Were you aware that these SDKs (refer to specific SDKs we detected in their app) collect this data?
28. Do any of these behaviors contradict what is mentioned in the privacy policy of your app?
29. Did you inform end users about these kinds of behaviors (collection and transmission of user data) in your privacy policy or any other means?
30. Do you have any plans for limiting data collection and data sharing with third parties?
31. Do you have any plans for removing specific SDKs from your app? If yes, please explain.

D Additional Figures

Below, we include more figures that aggregate developers' responses regarding how challenging privacy compliance is for them and the role they expect app markets to perform in the process.

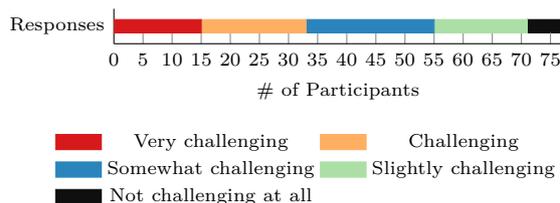


Fig. 2. Developers' opinions on how challenging privacy compliance is for them.

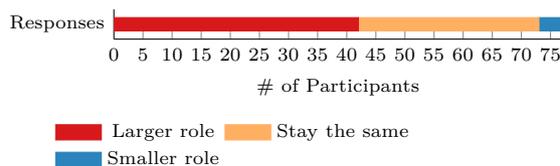


Fig. 3. Developers' opinions on how app markets should be involved in determining apps' compliance with privacy regulations.

E SDK Flags

Table 3 shows the privacy compliance flags we identified when we tested third-party SDK privacy compliance settings.

Table 3. Third-Party SDK Flags

SDK	Flags and privacy compliance settings
Unity Ads [28]	coppa (true or false), appLevelCoppa (true or false), calculatedCoppa (true or false)
Google Admob [15, 16]	tag_for_child_directed_treatment (0 or 1), rdp (0 or 1)
Adcolony [1]	gdpr_required (true or false), coppa_required (true or false), ccpa_required (true or false)
Vungle [30]	user_gdpr_consent_status (opted_in or opted_out), user_ccpa_status (opted_in or opted_out)
Chartboost [8]	pidatauseconsent (-1, 0, 1), privacy_us_privacy (1NY-, 1NN-)
Flurry [11]	fl.ccpa.optout (true, false) Note: Flurry does not allow integrating its SDK in apps before developers “certify” that their apps are not directed at children.
Mobfox [19]	coppa (0, 1), gdpr (0, 1), gdpr_consent (0, 1)
Appodeal [4]	coppa (true, false) [outbound traffic], for_kids (true, false) [inbound traffic] , consent (true, false) [outbound traffic] (0, 1)
Appsflyer [29]	SDK documentation expects developers to not initialize the SDK before obtaining user consent [29].
Amplitude [3]	The flags city, ip_address and lat_lng are set to false after using a method offered by the SDK which is called enableCoppaControl.
Yandex [31]	user_consent (0, 1)
Facebook [10]	COPPA (true, false)
Supersonic ads/IronSource [17]	is_coppa (true, false), consent (0, 1), gdprConsentStatus (true, false), do_not_sell (true, false)
Tapjoy [25]	cgdpr (0, 1), gdpr (0, 1), below_consent_age (0, 1)