Catherine Han*, Irwin Reyes, Álvaro Feal, Joel Reardon, Primal Wijesekera, Narseo Vallina-Rodriguez, Amit Elazari, Kenneth A. Bamberger, and Serge Egelman

The Price is (Not) Right: Comparing Privacy in Free and Paid Apps

Abstract: It is commonly assumed that *"free"* mobile apps come at the cost of consumer privacy and that paying for apps could offer consumers protection from behavioral advertising and long-term tracking. This work empirically evaluates the validity of this assumption by comparing the privacy practices of free apps and their paid premium versions, while also gauging consumer expectations surrounding free and paid apps. We use both static and dynamic analysis to examine 5,877 pairs of free Android apps and their paid counterparts for differences in data collection practices and privacy policies between pairs. To understand user expectations for paid apps, we conducted a 998-participant online survey and found that consumers expect paid apps to have better security and privacy behaviors. However, there is no clear evidence that paying for an app will actually guarantee protection from extensive data collection in practice. Given that the free version had at least one third-party library or dangerous permission, respectively, we discovered that 45% of the paid versions reused all of the same third-party libraries as their free versions, and 74% of the paid versions had all of the dangerous permissions held by the free app. Likewise, our dynamic analysis revealed that 32% of the paid apps exhibit all of the same data collection and transmission behaviors as their free counterparts. Finally, we found that 40% of apps did not have a privacy policy link in the Google Play Store and that only 3.7% of the pairs that *did* reflected differences between the free and paid versions.

Keywords: Privacy, mobile applications, measurements, dynamic analysis, static analysis, consumer protection

DOI 10.2478/popets-2020-0050
Received 2019-11-30; revised 2020-03-15; accepted 2020-03-16.

***Corresponding Author: Catherine Han:** University of California, Berkeley, E-mail: catherinehan@berkeley.edu
Irwin Reyes: Two Six Labs / International Computer Science Institute, E-mail: irwin.reyes@twosixlabs.com
Álvaro Feal: IMDEA Networks Institute / Universidad Carlos III de Madrid, E-mail: alvaro.feal@imdea.org
Joel Reardon: University of Calgary / AppCensus, Inc., E-mail: joel.reardon@ucalgary.ca
Primal Wijesekera: International Computer Science Institute / University of California, Berkeley, E-mail: primal@berkeley.edu
Narseo Vallina-Rodriguez: IMDEA Networks Institute / International Computer Science Institute / AppCensus, Inc., E-mail: narseo.vallina@imdea.org
Amit Elazari: University of California, Berkeley, E-mail: amit.elazari@berkeley.edu
Kenneth A. Bamberger: University of California, Berkeley, E-mail: kbamberger@berkeley.edu
Serge Egelman: International Computer Science Institute / University of California, Berkeley / AppCensus, Inc., E-mail: egelman@cs.berkeley.edu

1 Introduction

Mobile app marketplaces offer consumers a large selection of products: as of 2019, the Google Play Store offers approximately 2.8M Android apps [5], while Apple's App Store lists approximately 2.2M iOS apps [9]. Many apps are available free of charge, while others require consumers to pay a one-time fee to download them: roughly 4.4% of apps in the Google Play Store require payment [44], as compared with 6% of iOS apps [31].

Common app pricing models include free, paid, *"freemium,"* and *"paidmium"* [46]. Free apps are available to download and do not offer in-app purchases, while paid apps require the user to pay for the initial download. The *"freemium"* model raises revenue primarily from in-app purchases, while the app itself is free to install. Likewise, the *"paidmium"* model also relies on in-app purchases, though the app itself also costs.

In aggregate, free apps attract over 10 times the volume of downloads as paid apps [34]. Developers of free apps rely on methods to generate revenue besides directly collecting money from paying consumers, such as partnering with advertising networks to serve ads to users. In 2015, the annual net-to-publisher revenue derived from mobile in-app advertising worldwide was $40 billion and is projected to be $117 billion in 2020 [4].

It has become apparent that users often trade their privacy for these *"free"* apps [8]. The question, however, remains unanswered for paid apps: are consumers

of paid apps truly safe from extensive user profiling and tracking? Users paying for apps expect them to be of higher quality compared to free versions [33], and a common selling point to that end is the removal of ads in paid versions. Even media outlets have reinforced these consumer expectations, stating that paid apps have better security and privacy assurances than free apps [3]. The lack of ads, however, might give false assurance that these apps are free of extensive data collection, a practice often associated with user tracking for the purpose of ad targeting. That is, even if an app does not display ads, it could still perform invasive tracking for the purpose of serving highly-targeted ads in *other* apps. This could be through the use of "free" third-party services, like game engines and social networking platforms (*i.e.,* services that typically perform aggressive personal data collection [35, 39]), or trading personal data through data brokers [51].

Meanwhile, regulators have been pushing tech companies for increased transparency about their data collection practices. In a recent landmark ruling against Google, the French data regulator, CNIL, levied a 50 million Euro fine for a breach of the European Union (EU) General Data Protection Regulation's (GDPR) transparency and informed consent requirements concerning data collection for personalized ads [11]. Subsequently, underscoring American frustration with the abuse of consumer data by large technology companies, the Federal Trade Commission (FTC) approved a $5 billion fine against Facebook for its misuse of users' personal information—marking the largest fine levied by the federal government against a tech giant [21]. The recently passed California Consumer Privacy Act (CCPA) [43] further exemplifies rising regulatory and public concern surrounding consumers' ability to make informed decisions about their digital safety.

Exploring if app behaviors comport with user expectations can inform developers, regulators, policymakers, and consumers alike. Potentially misleading representations may run afoul of the FTC's prohibitions against deceptive practices and state laws prohibiting unfair business practices, as well as general privacy regulations, such as the GDPR and CCPA. Finally, such inquiry can also inform economic models exploring the viability of "pay for privacy" consumer protection models [2].

To that end, we explore the facets of consumer expectation of free and paid apps. We surveyed 998 participants about their general expectations for free and paid apps. Our results clearly indicate that respondents were more likely to expect better privacy-preserving behavior from paid versions of apps, relative to their free counterparts. To measure how well these expectations are met in reality, we sought ground truth by comparing the implementation and data collection practices of free Android apps and their paid counterparts offered on the Google Play Store over a corpus of 5,877 pairs of apps. We measured potential differences and similarities along four key aspects: third-party libraries—which may be used for advertising and tracking—bundled within the apps, the nature of the permissions they access, the types of sensitive data shared with third-party services, and differences in privacy policies offered by each version of the app.

We acknowledge that our analysis does not consider or vet any claims (or lack thereof) made by developers about any privacy advantages their paid apps might have over free offerings. The scope of this work instead centers on how free and paid apps themselves differ in implementation and behavior, and how those observations comport with consumer expectations.

From our results, we make four key observations about pairs whose free versions exhibited at least one of each metric:

- 45% of paid versions include all of the same third-party libraries as their free versions.
- 74% of paid versions hold all of the same dangerous permissions (which guard sensitive resources like the contact list and geolocation) as their free versions.
- Most of the app pairs lie on either extreme, in which 32% of paid versions exhibit all the same data collection behaviors as the free versions while 52% have none of them.
- In spite of increasing regulatory pressure to improve transparency, only 55% of the pairs in our corpus provided a privacy policy on their app store pages, and 3.7% of corpus pairs have policies that differ between the free and paid apps.

2 Related Work

Our research draws on and contributes to relevant prior work on the analysis of privacy in mobile apps and comparisons of free and paid apps. We also place our work in the context of consumer expectations of privacy online.

2.1 Analysis of Mobile App Privacy

Previous work has analyzed the collection of personal information through both static and dynamic analysis [1, 10, 24]. Static analysis consists of evaluating soft-

ware without execution [1, 7], whereas dynamic analysis focuses on tracking the transmission of sensitive information at runtime [10, 36, 40, 50]. Runtime behavior is often paired with the observation of network traffic to identify personal data dissemination. To automate the process of such analysis, researchers have developed several tools to not only simulate user interaction, but also give summaries of network traffic [20, 35, 52]. Both approaches have been largely used by the research community to study, for instance, the dissemination of personal data [13, 22, 32], malware behavior on Android [25], and different app ecosystems such as preinstalled software [12] or mobile browsers [23, 26].

The methods in this paper combine static and dynamic analysis methods introduced in previous work [20, 38], broadening the analysis of mobile apps by loosening the constraints on our corpus in two key ways: (1) including both free and paid apps for direct comparison to one another; and (2) having a broader scope than apps designed only for children and families.

2.2 Comparison of Free and Paid Apps

Prior research also sought to examine the relationship between free and paid mobile apps. Researchers have used static analysis to examine the prevalence of tracking libraries and their data collection behaviors in free and paid apps [41]. Other studies have investigated vulnerabilities associated with the maintenance of software and inclusion of third-party libraries in apps with different monetization models [35, 45].

Earlier works center on a broad comparison of a body of free apps with a body of paid apps. Our work offers a novel view on the comparison between free and paid apps by presenting a precise, side-by-side analysis of specifically constructed pairs of apps: a free app and its paid "premium" version. Our approach compares apps that are directly related to one another: the same general app from the same developer, but offered separately as free and paid versions.

2.3 Consumer Expectations and Attitudes

Data Collection and Behavioral Advertising: In the space of consumer attitudes toward behavioral advertising, existing work has revealed that while some users desire the benefits of targeted advertising, a majority (64%) of the survey participants found the idea *invasive* [29]. McDonald and Cranor indicated they saw "signs of a possible chilling effect with 40% self-reporting they would change their online behavior if advertisers were collecting data." The same research also demonstrated that many participants had a poor understanding of how advertising works on the Internet—whether that meant the use of cookies, the presence of behavioral advertising, the usage of tracking and fingerprinting, or the types of legislation and protection there are for consumers on the Web. Overall, this underscores the gap in consumer understanding in a "pay for privacy" model, as it was found that while participants were "comfortable with the idea that advertising supports free online content, ... they do not believe their data are part of that exchange."

Paying for Privacy: Previous work has examined the value of privacy to consumers and how privacy benefits could be leveraged as a selling point by businesses [47]. Survey data indicates that consumers place great importance on having insight and control over how companies handle their personal data [28]. While some may argue that consumers today do not care for their privacy, a past study has shown that half of its participants "disagreed or strongly disagreed that they do not care if advertisers collect search terms, or if advertisers collect data about websites visited, both of which occur regularly for behavioral advertising and analytics data [29]."

With growing concern surrounding online privacy, prior work indicates that some consumers are indeed willing to pay a premium for products that protect privacy—conditioned on the privacy benefits being explicitly listed in a digestible manner for the consumers [47]. To gain improved understanding of users' perceptions of privacy behaviors of paid and free app versions, we take a step further and conduct a survey to study whether advertising apps as *"ad-free,"* would lead most consumers to believe that this is synonymous with *"better privacy"* or not. With this, we can identify ways in which the behavior of apps differ from users' expectations, ultimately determining if the *"ad-free"* representation misleads consumers.

Likewise, research has shown that developers are also aware of consumers' desire for privacy [30]. However, the same study illuminated a gap between developers recognizing this demand and actually meeting it in practice: developers spend little time vetting and configuring the ad networks they integrate into their apps, opting instead for ad libraries that are popular and easy to use. Developers frequently accept the default content and privacy settings of the ad libraries they use, even if changing the configuration would improve user privacy (by reducing the ad network's data collection) or increase revenue (by increasing the data collection).

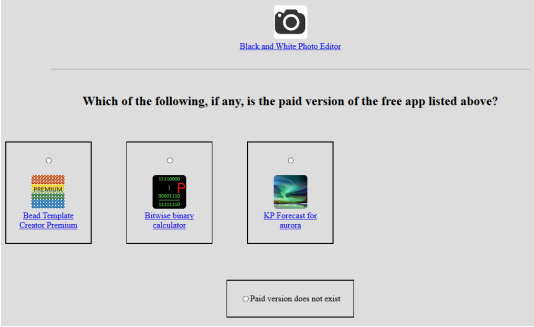

Fig. 1. MTurk task for selecting the paid version of a free app.

3 Methodology

The main objective of this work was to analyze differences in app execution and privacy practices between free apps and their respective paid counterparts. Using a combination of static analysis and dynamic analysis techniques, we examine the implementation details and data collection behaviors of 5,877 pairs of free and paid Android apps, downloaded from the Google Play Store in April 2019. For the purposes of our study, we define *"free apps"* as those that are available for download for no up-front cost and *"paid apps"* as those that require a one-time payment to be downloaded. The paid apps in our analysis are single discrete app purchases with no additional payments made. We acknowledge that apps may employ other monetization strategies, such as the *"fremium"* model—in which an app offers additional features gated behind in-app purchases and recurring subscriptions, but they are outside the scope of this study. Our corpus, however, contains 886 pairs in which at least one of the apps was listed as offering in-app purchases.

3.1 App Corpus

We formed our app corpus by first scraping the *"Top Free"* charts in each of the categories in the Google Play Store. This provided a collection of the most commonly downloaded free apps, for which we attempted to find their respective paid versions, if any existed. Unfortunately, the Google Play Store does not explicitly associate free apps with their paid counterparts (*e.g., "Quick PDF Scanner FREE"* and *"Quick PDF Scanner PRO"*), or even indicate if a paid release exists at all for a given free app. This limitation required us to manually identify the free/paid pairs necessary for this analysis. In order to find these pairs at scale, we crowdsourced this as a labeling task on Amazon Mechanical Turk (MTurk), recruiting 2,944 workers. We presented workers with a free app drawn from our earlier scrape of the Google Play Store and a list of all paid apps from the same developer (Figure 1). The list of paid apps from the same developer was collected by following the link to the developer's catalog, as listed in the free app's Google Play Store page. This method is conservative in that it does not combine multiple developer listings belonging to the same company; for example, *"Comcast"* [19] and *"Comcast Cable Corporation, LLC"* [18]. Hence, if the two versions (free and paid) are under different developer listings, then our approach does not capture such pairs for the analysis. In order to increase the likelihood of a valid free/paid pairing for a given task, we only presented workers free apps whose titles or package names contained the words *"free"* or *"lite"*, as those keywords suggest that a *"paid"*, *"pro"*, or *"full"* version is likely to exist. If a free app did not have an apparent paid version among the choices, then workers were instructed to select the *"Paid version does not exist"* option. Each free app was presented to three different workers, then subsequently adjudicated by researchers for agreement and correctness. Workers were paid $0.10 for each identified pair in consensus with the others.

This yielded 7,023 potential pairs for further inspection. To keep costs manageable, we discarded 182 pairs whose paid app was priced as more than $10. We also discarded 680 pairs in which one or both of the apps were no longer available when we attempted to download them in April 2019. Finally, due to region-locking or incompatibility with our Nexus 5X test devices, we were unable to purchase or download apps in 284 pairs.

We ultimately identified, purchased, and downloaded a final corpus of 5,877 pairs of free apps and their paid versions.[1] We obtained the latest versions of these apps from the Google Play Store in April 2019, at an average cost of $2.35 per paid app. Our corpus represents 3,464 unique developers, and the free apps in this collection had a median install count of 10,000, as reported in the Google Play Store.

3.2 Evaluating Apps

We evaluated pairs of free and paid apps along the following dimensions: (1) third-party SDKs; (2) declared Android permissions; (3) transmissions of sensitive data; and (4) privacy policies. This approach reveals both the

1 https://raw.githubusercontent.com/catherinehan/free-vs-paid/master/corpus-pairs.txt

potential exposure of sensitive data to apps (via bundled third-party code and declared permissions) and the *actual* observed data sharing behavior of those apps—these observations are not meant to be comprehensive, but are instead bounded by the duration and nature of inputs during test execution. We measured these using the following methods:

Static Analysis: We identified third-party SDKs and declared permissions by directly examining app binaries (APKs) without running them, leveraging existing tools to perform this static analysis. The Android Asset Packaging Tool (*aapt*, included in the Android developer tools) reports the permissions declared by apps to access various resources, such as Internet, geolocation services, and the user's identity, among others. Our analysis only considers the *dangerous* permissions as defined by Android itself [16], which protect access to sensitive system resources and user data. These declared dangerous permissions provide insight on apps' potential runtime capabilities. Using *aapt* permissions reports, we measured the differences and similarities in declared permissions within pairs of free and paid apps, revealing the degree to which paid apps are able to access the same or a different set of (potentially sensitive) resources as their free counterparts.

We analyzed APKs for third-party libraries using both Apktool [48] and LibRadar [27]. Third-party libraries are commonly used to expedite software development by providing drop-in functionality, like graphics rendering and crash reporting, as well as ad delivery and audience analytics. Third-party libraries run with the same privileges as their host app and are thus able to collect and transmit sensitive data. We used Apktool to reverse-engineer the class structures of apps; we then analyzed their outputs to identify what third-party packages were present (*i.e.,* by disregarding packages that belong to the core app) and are thus able to access the same data as the core app.

As third-party libraries can provide a broad range of functions, we also relied on LibRadar to gain additional insight into the nature of the third-party libraries we identified. LibRadar detects and identifies many libraries based on stable API features, offering categorizations of these libraries' purposes, such as advertisement, development aid, mobile analytics, payment, social network, and utility. Because most of these categories have use cases that can be associated with the core app features, we focused on categories that could pose privacy threats, such as advertising and analytics.

Data Type	Description
AAID	Android Advertising ID
Android ID	Unique ID created at Android setup
GSF ID	Google Services Framework ID
HW ID	Phone hardware ID (serial number)
IMEI	Mobile phone equipment ID
SIM ID	SIM card ID
MAC Address	MAC address of WiFi interface
Email	Email address of phone owner
Phone #	Mobile phone's number
Latitude, Longitude	User location
Router MAC Address	MAC addresses of nearby hotspots
Router SSID	SSIDs of nearby hotspots

Table 1. The types of personal information that are detected in our analysis logs. The top half of the table lists persistent IDs.

Dynamic Analysis: While the static analysis techniques described above aid in determining the potential behavior of apps, we also utilized dynamic analysis techniques to observe apps' actual runtime behaviors. Using methods derived from prior research [20, 32], we executed the corpus of apps on Nexus 5X smartphones that we instrumented with a custom version of the operating system with our own root certificate, performing MITM on TLS connections. Our system detects certificate-pinned TLS that it cannot decrypt, but we found no instances of this in our corpus. This is consistent with prior work [37] showing that TLS-interception prevention methods, such as cert-pinning, are present in less than 2% of all apps and is likely only among large development factories like Google and Facebook. Our instrumentation captures and records all network activity and OS activities such as resource usage; however, our analysis is based on the captured network logs. These network captures were subsequently analyzed to determine if and when apps transmitted sensitive data to remote services. We searched these network captures for the transmission of sensitive data that can be used to uniquely track users over time and across different services: persistent identifiers, such as the Android Advertising ID (AAID), IMEI, and Wi-Fi MAC address, as well as personally-identifiable information (PII), such as geolocation data and the device owner's name, email address, and phone number (Table 1 lists all the identifiers we monitored). We took extra measures beyond naive string matching to find sensitive data in network traffic: our analysis also detects string transformations (*e.g.,* upper case, lower case, reverse), hashes (*e.g.,* MD5, SHA-1, SHA-256), and encodings (*e.g.,* Base64) of that data, and certain combinations thereof.

This dynamic analysis approach poses a challenge in that it is often insufficient to simply launch an app in the instrumented environment. Apps need user input to progress through different screens and perform a variety of functions that trigger different app behaviors. We used the Application Exerciser Monkey [17] from the Android developer toolbox to simulate user input and automate the dynamic analysis. The Exerciser Monkey is a naive input generator that sends pseudorandom tap and swipe events to the app with no knowledge of what is displayed on the screen. Prior work on the dynamic analysis of children's apps found that random input generators reach, on average, 61% of the unique app screens that a human user would [20]. As such, the results of our dynamic analysis should be interpreted as a lower bound with no false-positives and a potentially high false-negative rate: we are unable to determine if unobserved behaviors simply did not happen in a given execution or if they would never happen.

In order to compare the observed behaviors of free apps to their corresponding paid versions as fairly as possible, we executed pairs simultaneously on identically-configured phones with the same Exerciser Monkey input event sequence (*i.e.,* the Exerciser Monkey was given the same seed to generate the input sequence). This increases the likelihood that observed differences in data collection behaviors arose from implementation differences, not differences in the inputs.

Although our corpus contained 5,877 pairs of free and paid apps, we were unable to perform dynamic analysis on 55 of those pairs because at least one of the apps either failed to install or caused the test device to crash when launched. All dynamic analysis results in this research were drawn from the remaining 5,822 pairs that successfully ran to completion in our test environment.

Privacy Policies: New privacy regulations, like the EU GDPR, enforce the need for user consent before apps can collect, treat, or share users' private data, when there is not another legal basis. To inform this consent, users are expected to turn to privacy policies. Privacy policies *should* inform users of the different types of data collected, the usage of this data, and whether apps share this with any third parties. Since 2004, under the California Online Privacy Protection Act (CalOPPA), online services—including mobile apps—that collect personal information from California consumers have been required to post privacy policies [42]. To examine this, we identified app privacy policies using the "Privacy Policy" link provided in their individual Play Store listings and downloaded the content of the page at that link.

In order to identify potential differences between free and paid apps' privacy policies, we extracted the text body of their respective policies, then performed a text comparison (using `diff`) between the two. As this method is prone to false positives, we manually reviewed the differences and identified which policies reported more data collection or third-party libraries in the free version of the apps.

3.3 Consumer Expectations

In addition to examining the differences in behaviors within pairs of apps, we also wanted to understand consumer expectations surrounding free and paid versions of apps. To this end, we designed a survey to answer the following research questions:

- Do consumers actually expect better privacy practices from paid apps, as some experts suggest?
- How does the monetization model of the app (*i.e.,* free or paid) affect consumer purchasing decisions?

Survey Flow:

Our survey was composed of several subsections (see Appendix A): a mix of open-ended responses, multiple-choice questions, and 5-point Likert scale questions, concluding with a series of demographics questions. The main focus of the survey was to present participants with a hypothetical paid "ad-free" version of a desirable free app, then surveying their expectations for how the paid version might differ from the existing free version. To that end, after obtaining their consent, we presented participants with a list of the most popular free apps: Amazon, Facebook, Instagram, Lucky Go, Pandora, Snapchat, and TikTok. We asked them to imagine that they were setting up a phone anew, and thus must pick which of these apps they would be most likely to install. Those who reported none of the above were disqualified from proceeding.

On the next page of the survey, we presented participants with two mocked-up Google Play Store listings for their selected app and asked which they would be most likely to install: (A) the free version or (B) a hypothetical paid version—priced at $0.99—that was advertised as being "ad-free" (Figure 2). We randomized whether the free version was labeled as A or B, but re-coded this so that in our analysis App A always refers to the free version, while App B refers to the paid version. After participants indicated their preferred ver-

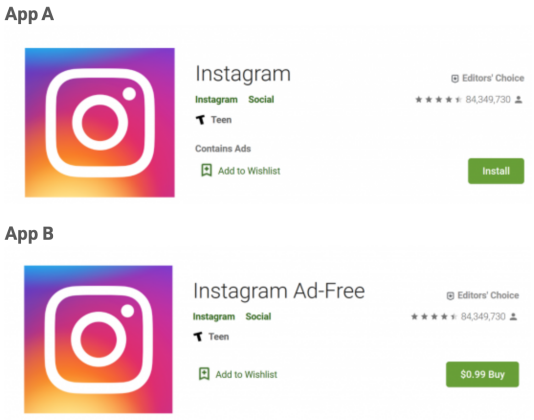

Fig. 2. Participants were shown Play Store listings of their chosen app (in this example, Instagram) as two versions: (A) the free version and (B) a hypothetical paid "ad-free" version.

sion and entered their reasoning in an open-ended text field, we asked participants, *"In what way, if any, would you expect the above two apps to differ?"* Responses to this question were also collected using an open-ended text field. Following best practices for grounded theory coding, the open-ended responses were coded by two researchers who collaborated on a master codebook, and then proceeded to code all of the responses independently, ultimately meeting to agree on the final codes.

After participants answered these questions, they proceeded to the next page of the survey, where they were asked, *"In what ways do you believe a user's personal data would be treated differently between the two apps?"* They responded using another open-ended text field. Following this question, we asked participants to answer several Likert-scale questions using the following 5-point scale: *"Definitely A (1),"* *"Likely A (2),"* *"Equally A and B (3),"* *"Likely B (4),"* and *"Definitely B (5)."* with respect to specific privacy and security-related practices. For example, participants were asked to choose which app they believed was more likely to share their data with third-party services, advertisers, and law enforcement agencies. They were also asked to choose which app they believed was more likely to comply with privacy laws and regulations. The survey then concluded with demographic questions.

Recruitment: We recruited 1,085 participants from the Prolific Academic survey platform, limiting participation to those within the U.S. and having a 95% or greater approval rating. We conducted our study during December 2019, though we piloted an earlier version of the survey with 1,100 participants (we do not report on those responses in this paper). Of the final 1,085 participants that we recruited, we disqualified 87, resulting in a final sample of 998 responses. Our sample was gender-balanced, with 48% self-identifying as male; the median reported age was 34, with the reported ages ranging from 18 to 79. In addition, approximately 60% of our sample had at least a bachelors' degree, and 48% of our sample reported themselves as single (see Appendix B for further demographic details).

The survey took under five minutes to complete, for which we compensated participants $1.00 for their time. Study procedures, the recruitment posting, and consent form were reviewed and approved by the University of California, Berkeley institutional review board.

4 Limitations

We acknowledge that our app analysis is not fully comprehensive, which is largely due to various mobile software development techniques. Recent research has shown some third-party libraries and app developers circumvent the Android permission system, whether through coordination with higher-privileged apps or through bugs in the Android platform [38]. Our static analysis only considered the permissions explicitly declared by apps, not whether apps attempt to gain access to resources for which they do not have privileges. Additionally, our analysis of third-party libraries did not attempt to identify obfuscated libraries: that is, libraries whose names have been stripped of identifying information and replaced with generic symbols at compile time (*e.g.,* `com.adnetwork.sdk` becoming `a.b.c`). This information was used to determine which third-party libraries in free apps were also present in paid versions. As such, we believed that it was more prudent to disregard obfuscated libraries in order to avoid falsely concluding that an obfuscated library in the free app was indeed the same as another obfuscated library in the paid app. From our corpus of 5,877 pairs, obfuscated libraries were present in 645 pairs, accounting for 11% of the free and paid pairs we analyzed. Hence, our analysis yields a lower bound of third-party libraries found.

The dynamic analysis relied heavily on a random input generator to drive apps and execute different functions. While we made a best effort to ensure that the free and paid versions of a given app were run as identically as possible—by executing the two versions on two identically-configured phones, both receiving the same input sequences, and performing the tests simultaneously—we note that this does not guarantee that differences in observed behaviors came solely from how the two versions were implemented. The free and

paid version of a given app could, for example, have subtle UI differences between them and could therefore proceed down entirely different execution paths, even under the same input sequence. Likewise, app execution may not be entirely deterministic: for instance, apps may communicate with different ad networks depending on the state of the marketplace at any given time, due to the effects of real-time and parallel bidding [6].

We imposed a $10 price cap when purchasing paid apps, as some apps were unreasonably costly (*e.g.,* more than $100) and thus had few actual downloads. Under this policy, we excluded 182 valid pairs identified in the MTurk matching task. We acknowledge the possibility of bias in our results arising from this price cap. We believe, however, that any such bias would be minimal. Our analysis mainly presents aggregate results. Even if we had included all 182 overpriced pairs, they would only represent a 3% expansion of the corpus, and the remaining 5,877 pairs that comprise our corpus would dominate the findings. One aspect in which our analysis may be incomplete due to this is in evaluating the specific privacy characteristics of paid apps above the price cap; for example, a professional software tool compared to its feature-limited free version.

We investigate the differences between free and paid apps with respect to how they are implemented and how they behave at runtime. This analysis uses observations in free apps as a baseline: all pairwise comparisons presented are conditioned on the free app having at least one observation for any of the corresponding metrics; therefore, we disregard pairs in which the free app had no third-party packages, no permission requests, or shared no sensitive data with a remote service, respectively. Note that there were indeed observations along these dimensions that were exclusive to the paid app within a pair. For example, a paid app may have used a third-party library that was not found in the free app. We discuss these in addition to the pairwise results.

4.1 Declared Android Permissions

The Android permission system serves to protect user privacy by requiring apps to hold appropriate permissions to use various device capabilities (*e.g.,* Internet connectivity and persistent disk storage) and access sensitive user data (*e.g.,* phone number, unique identifiers, and geolocation). The Android Open-Source Project (AOSP) defines a set of default Android permissions that should be supported by every Android Play Protect-certified device [12]. Android has designated different security levels to permissions based on their risk. A subset of them, are labeled as *dangerous* because they guard sensitive functionality that directly affects user security and privacy [14]. Apps that attempt to access resources guarded by these dangerous permissions must prompt the user for approval before they can do so. We focus our study primarily on dangerous permissions given their potential impact on users' privacy.

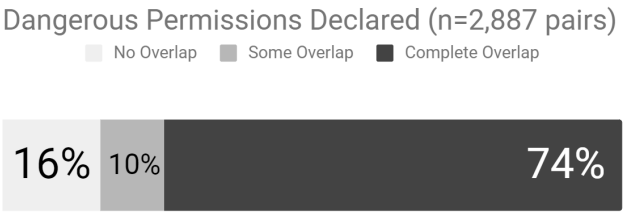

Fig. 3. Frequency of dangerous Android permissions in common between free and paid versions, given that the free version had at least one dangerous permission.

Table 2. Dangerous permissions declared by free apps, paid apps, and both, from corpus N = 5,877 pairs

Dangerous permission	Only in free	Both free and paid	Only in paid
WRITE_EXTERNAL_STORAGE	325	1835	107
READ_EXTERNAL_STORAGE	128	706	76
READ_PHONE_STATE	209	559	53
ACCESS_FINE_LOCATION	136	409	72
ACCESS_COARSE_LOCATION	208	328	54
CAMERA	21	262	33
GET_ACCOUNTS	68	259	49
RECORD_AUDIO	22	231	21
READ_CONTACTS	4	77	6
READ_CALENDAR	3	48	6
WRITE_CALENDAR	3	47	7
CALL_PHONE	3	32	3
WRITE_CONTACTS	1	27	3
BODY_SENSORS	0	5	0
RECEIVE_MMS	0	1	0
RECEIVE_SMS	0	1	0
RECEIVE_WAP_PUSH	0	1	0
PROCESS_OUTGOING_CALLS	0	0	1

5 App Behavior

Out of our corpus of 5,877 pairs, 2,887 (49%) contained free apps that declared at least one Android-defined dangerous permission. In 74% of these pairs, the corresponding paid version (Figure 3) declared the same dangerous permissions as the free version. That is, when a free version holds privileges to access sensitive device

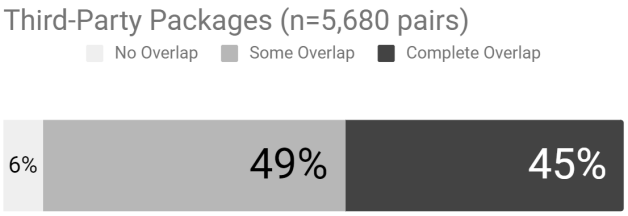

Fig. 4. Frequency of third-party package reuse among free/paid pairs, where the free app had at least one third-party package.

resources and user data, its corresponding paid version has the same capabilities nearly $\frac{3}{4}$ of the time. The most common dangerous permissions shared by both free and paid versions of an app were those to use shared—amongst all other apps with the permission—disk storage, get information about the phone's state (*e.g.,* phone number, cellular network information, call status), and access the device's geolocation (Table 2).

We also observed 350 pairs in which the paid apps held dangerous permissions *not* declared by their respective free versions. As before, the most common dangerous permissions held exclusively by paid apps were those to access shared disk space, which is notable as this is a known side channel for circumventing the permission system [38]. Taken together with the high likelihood that a paid app will have all the same dangerous permissions as its free counterpart, these findings suggest that paying for an application does not guarantee a reduction in users' exposure to data collection, and in some cases could even increase their exposure via access to shared disk space.

5.1 Bundled Third-Party Packages

The use of third-party code is common practice in software engineering to expedite development. In mobile apps, third-party libraries allow for pre-built functionality like graphics rendering, advertising, or analytics. Third-party code bundled in apps gains the same privileges as the host app, including permissions.

Of the 5,877 pairs in our corpus, 5,680 (97%) had at least one third-party package in the free version according to LibRadar [27]. Of these (Figure 4), we observed that 45% of paid apps contained the same third-party libraries as the free versions, while only 6% of paid apps showed no third-party libraries carried over from the free version. The remaining 49% of paid apps had varying degrees of third-party library reuse from the free version to the paid version. This suggests that paid apps are likely to contain most, if not all, of the same third-party libraries as the free versions. This may leave paying consumers exposed to the same potential for third-party data collection as found in free apps. Although we acknowledge that our static analysis does not account for third-party libraries included but not actually executed (*i.e.,* dead code), these results show that developers have little motivation to remove externally-produced code in paid apps.

Based upon the library categorizations of LibRadar, we analyzed the types of third-party libraries present in free and paid versions of apps, focusing our attention on libraries associated with labels such as *"Advertisement"* and *"Mobile Analytics."* Focusing on advertising libraries specifically, LibRadar detected at least one ad library present in either the free or paid release (or both) in 3,043 pairs (52% of the overall corpus). The most commonly observed advertising library was `com.google.ads` at 2,623 (45% of the entire corpus) occurrences within our corpus, followed by `com.unity3d.ads`, `com.mopub`, and `com.chartboost.sdk`, with 237, 235, and 234 occurrences, respectively. Of those pairs, there were 2,918 (50% of the entire corpus) free apps where ad libraries were detected, while 1,488 (25%) paid apps were found to contain ad libraries; there were 1,320 (22%) pairs where there was at least one advertising library that was present in both the free and the paid version. Furthermore, 209 paid apps even bundled at least one advertising library that was not present in the free counterparts, suggesting that some paid apps will not only share some of the same advertising libraries included in the free version, but also introduce new ones.

Analytics libraries are also known to collect a wide array of information about users [35]. Our data, however, showed that analytics libraries are less prevalent than ad libraries across all pairs. Only 831 (14% of the entire corpus) pairs had at least one analytics library in either version of the pair. The most common analytics library was `com.flurry.android` (owned by Verizon) followed by `com.crashlytic.android` (owned by Google) and `com.google.analytics`. Overall, 754 (13%) free versions had an analytics library and 86% of their paid counterparts shared at least one analytics library with their free version. We also found that 96 paid versions introduced at least one new analytics library which was not present in the free version.

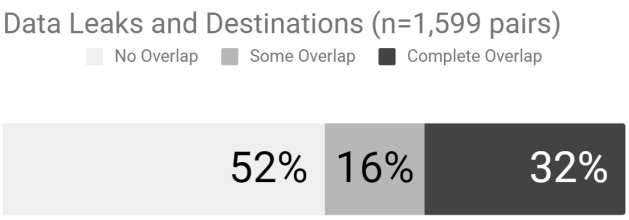

Fig. 5. How often paid apps shared *the same* sensitive data with the same destinations as their free versions, if the free version exhibited that behavior.

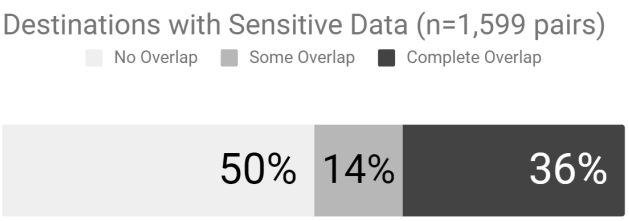

Fig. 6. How often paid apps shared *any* sensitive data with the same destinations as their free versions, if the free version exhibited that behavior.

5.2 Network Transmissions

Third-party services bundled in apps routinely collect data about users and their devices [20, 35]. For example, crash reporting services might gather hardware specifications and usage telemetry to help developers debug their apps, along with users' unique identifiers. Likewise, advertising networks collect persistent identifiers and personal information to better match users with ads relevant to their interests (commonly known as "ad targeting"). The data collected by a given third-party could also be packaged and traded with other organizations, including data brokers [35, 51].

Our dynamic analysis environment captured all network traffic—including TLS-protected transmissions—originating from and received by the apps we tested. We used these network captures to identify when apps share sensitive data (*e.g.,* persistent identifiers and personal information) and to where that data is sent. We also mapped domains to organizations using the methodology proposed by Razaghpanah et al. [35].

There were 1,599 (27%) pairs from our corpus of 5,877 in which the free app transmitted sensitive data to remote servers over the Internet. Among these pairs, we identified how many of their corresponding paid versions also sent the *same* pieces of information to the *same* destinations (Figure 5). For example, we noted when paid app Y sent location data to `somedomain.com` and whether that behavior was observed in its corresponding free app X. Our analysis, however, omitted transmissions to IP addresses that were not resolved to domain names, and all resolved domains were converted to second-level domains. In 516 (32%) of these pairs, we found the paid app transmitted all the same pieces of sensitive data to the same destinations as what was observed in the free app. An additional 255 (16%) of pairs showed the paid app exhibiting some of their respective free apps' data collection behaviors. While further analysis is needed to determine the purpose and necessity of sending this data with respect to apps' core functionality, the most frequently-observed sensitive data type shared by both free and paid apps were those that enable persistent tracking—*i.e.,* Advertising ID in 651 pairs (41%), Android ID in 570 pairs (36%), IMEI in 65 pairs (4%) and Location in 39 pairs (2%).

Using the 771 pairs in which the paid app transmitted at least some of the same data to the same destinations as the free version, we examined if these transmissions were more (or less) likely to use TLS in the paid version, as compared to the free version. There were 250 free apps that sent sensitive data using a mix of unencrypted and TLS-protected transmissions, and 521 that sent data using TLS only. Likewise, there were 257 paid apps using mixed transmissions, and 514 that used TLS exclusively. 758 of these pairs (98%) had no difference in TLS use (or mixed-use) between the free and paid apps. In 3 pairs, the paid app improved on the free version by using TLS in all observed transmissions. In 10 pairs, the paid app was observed sending some unencrypted data when its free version only used TLS transmissions.

We also considered the destinations themselves that received various types of sensitive data: given that the free version transmitted sensitive data to a particular domain, we measured how frequently the corresponding paid version transmitted sensitive data (of any type) to the same destination as well (Figure 6). The non-empty set of pairs (2%) that shifted from having no sensitive transmissions in common between the free and paid versions (when we earlier required the destinations *and* data types to match) to having some overlap. This suggests that there are cases where free and paid versions collect different kinds of data for the same service.

We further investigated the particular *domains* that lie on the two extreme ends of observed data sharing over the Internet. In particular, we observed which domains are more likely to *receive* data from both the paid and the free versions, versus those more likely to be *removed* or deactivated in the paid version. We narrowed the search to those domains receiving persistent identifiers, a type of PII, which allow users to be tracked over

Table 3. Third-party domains contacted by at least 50 pairs, and which tended to be deactivated in the paid version. In brackets, we report the parent company when it is different from the domain name.

Domain name	Purpose	Only in free	Both free and paid	All identifiers %free	All identifiers %paid	Immutable IDs %free	Immutable IDs %paid
adadapted.com	Advertising / Analytics	117	1	98	2	0	0
adcolony.com	Advertising	94	16	96	17	5	0
adjust.com	Advertising	85	14	96	15	3	1
adrta.com (Pixalate)	Advertising / Anti-fraud	139	1	94	2	0	0
adsmoloco.com	Advertising / User engagement	121	1	69	3	0	0
adsrvr.org (theTradeDesk)	Advertising	72	3	84	2	0	0
amazon-adsystem.com (Amazon)	Advertising	117	3	61	0	21	0
applovin.com	Advertising / User engagement	93	20	85	18	1	0
appsflyer.com	Analytics / Attribution	108	22	92	18	1	2
chartboost.com	Advertising	118	41	90	30	85	29
criteo.com	Advertising	195	4	93	2	0	0
kochava.com	Advertising	47	2	92	7	5	1
manage.com (Criteo)	Advertising	73	0	61	1	0	0
mopub.com (Twitter)	Advertising	112	9	57	7	12	5
sitescout.com (Centro)	Advertising	124	1	96	2	0	0
startappservice.com (StartApp)	Advertising	58	1	76	1	10	0

Table 4. Third-party domains contacted by at least 50 pairs, and which tended to remain active in the paid version. In brackets, we report the parent company when it is different from the domain name.

Domain name	Purpose	Only in free	Both free and paid	All identifiers %free	All identifiers %paid	Immutable IDs %free	Immutable IDs %paid
coronalabs.com	Development framework / Analytics	9	47	36	34	36	34
crashlytics.com (Alphabet)	Crash reporting / Analytics	60	346	38	39	30	31
facebook.com	Social networking / Advertising / Analytics	106	220	68	48	0	0
flurry.com (Verizon)	Analytics	14	34	76	76	76	76
google-analytics.com	Analytics	35	73	13	12	13	12
unity3d.com	Game engine / Advertising / Analytics	89	252	89	73	79	70

time and across services. We report on domains that received data from at least 50 pairs. We counted whether the domain was observed only in the free version or both. Table 3 shows the domains that were most likely to be deactivated in the paid version. Table 4 shows the domains that were likely to remain in the paid version based on observed network transmissions. We did not observe any domains contacted by more than one app pair that were more likely to be exclusive to the paid version. As with the earlier analysis, we converted fully-qualified domain names to second-level domains (*e.g.,* `mpx.mopub.com` and `ads.mopub.com` are both counted as `mopub.com`.

In Tables 3 and 4, we also characterize the use of persistent identifiers by these domains. For each domain, we report the percentage of free and paid apps that sent persistent identifiers, such as e-mail or advertising ID, as well as the percentage that sent *immutable* (*i.e.,* non-trivially resettable) persistent identifiers, such as the device's IMEI or Android ID. These percentages are relative to the total number of pairs for which at least one app in the pair communicated with the particular domain.

Table 4 presents a subset of third-party domains that received persistent identifiers from apps in our study. This list contains only domains contacted by apps from at least 50 different pairs and that tend to *remain* in the paid version. This means that more free/paid pairs communicated with these domains in the paid version than exclusively in the free version. Unity is a game engine with associated advertising and analytics capabilities, and we see that it sends the Android ID, the AAID, the IMEI, and the Wi-Fi MAC addresses in free and paid versions of apps. Likewise, Crashlytics is a Google-owned event-reporting service that typically

collects both the Android ID and AAID, and sometimes the device serial number and precise geolocation, as well.

Nearly half of the paid apps had a varying degree of overlap with their free versions, in terms of their exfiltration of PII to third parties. These behaviors might be viewed as privacy violations, if consumers expect the paid version to be more privacy-protective. Our traffic analysis, however, paints a gloomy picture: there is, again, no guarantee that by paying for a mobile app, it will be less intrusive and more privacy-protective in terms of sharing data with third parties.

5.3 Popular Paid Apps

There were 27 paid apps in our corpus with high download counts (*i.e.,* 500K+ Google Play Store installs).[2] These popular pairs are similar to the corpus as a whole. For highly-downloaded pairs whose free app had at least one dangerous permission (N = 23), 70% of the paid apps declared all the dangerous permissions from the free app, 26% had some, and 4% had none. For third-party packages (N = 25), 48% of the paid apps had all the free version's packages and 52% had some. And, for free apps that transmitted at least one sensitive data type (N = 20), we observed 40% of paid versions sharing all the same data to the same destinations, 25% with some of the same sharing behaviors, and 35% with none.

We examined the Google Play Store listings for these popular paid apps and found that certain apps tout being ad-free as one of the benefits of paying for the app. The paid app `me.dreamsky.leagueofstickmanzombie` (1M+ installs) lists "No advertisiment [sic]" among the "privileges of this version."[3] Still, we observed that it transmitted all the same tracking data as its free counterpart: the Android ID to `chartboost.com`, and the AAID to `adcolony.com`, `facebook.com`, and `vungle.com`, among others. Similarly, the paid app `se.maginteractive.rumble` (500K+ installs) also says it's an "ad-free premium experience"[4] yet its APK contains some of the same ad libraries as its free version: `com.unity3d.ads`, `com.adcolony.sdk`, and `com.inmobi.sdk`, etc.

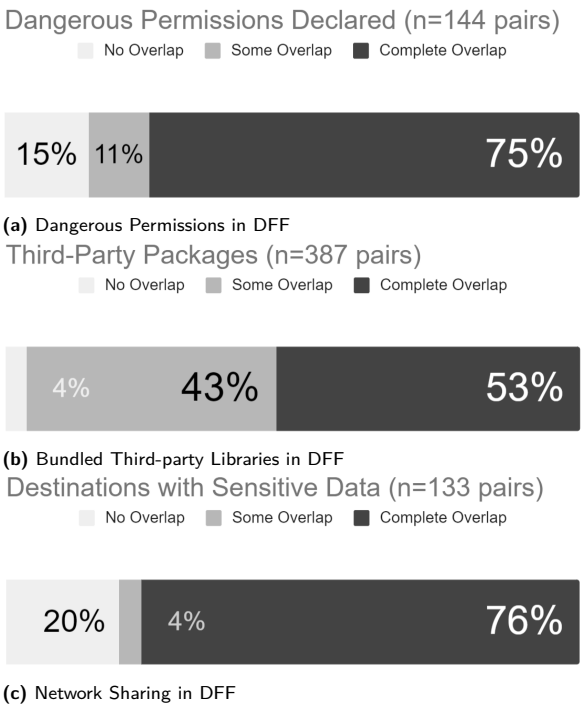

(a) Dangerous Permissions in DFF

(b) Bundled Third-party Libraries in DFF

(c) Network Sharing in DFF

Fig. 7. Designed for Family (COPPA) Analysis

5.4 COPPA Compliance

In the United States, data collection in children's apps is governed by the the Children's Online Privacy Protection Act (COPPA) [49], which has similar provisions to the European Union General Data Protection Regulation's children's provisions ("GDPR-K"). Both require parental consent prior to many types of data collection. Prior research has examined how mobile apps meant for children comply with COPPA [20]. While that work found widespread potential COPPA violations in free apps for children, it did not examine any paid apps, much less compare them to their free counterparts.

We examine apps that are subject to COPPA and analyze how the paid and free versions differ in their implementation and behavior. From the 5,877 free/paid app pairs in our corpus, we searched for pairs in which *both* versions were listed as "Designed for Families" (DFF) on the Google Play Store. App developers listing their software under the DFF program must state that their apps are compliant with COPPA [15]. We found 387 pairs of apps in the DFF program, for which we repeated the analysis presented in the previous subsections. While determining legal liability under COPPA is well beyond the scope of this work, we compare the free

2 https://raw.githubusercontent.com/catherinehan/free-vs-paid/master/popular-pairs.txt

3 https://raw.githubusercontent.com/catherinehan/free-vs-paid/master/me.dreamsky.leagueofstickmanzombie-20200213-playstore.png

4 https://raw.githubusercontent.com/catherinehan/free-vs-paid/master/se.magicinteractive.rumble-20200213-playstore.png

and paid versions in terms of their declared dangerous permissions (Figure 7a), bundled packages (Figure 7b), and network transmissions (Figure 7c).

Under all three analyses, the majority of the paid versions had at least some degree of overlapping behavior with their free versions, if not identical. While having the same third-party libraries or sharing data with the same remote domains are not COPPA violations by themselves, prior research has shown widespread potential violations [20] arising from these types of data collection behaviors. Given these observations, it is unlikely that paid versions of DFF apps differ from their free counterparts in terms of data access and collection, and thus, paid versions are unlikely to be better than free versions in terms of protecting children's privacy.

5.5 Privacy Policies

Many new privacy regulations (e.g., GDPR, CCPA, etc.) require app developers to post privacy policies that disclose apps' data collection and processing practices. The Google Play Store allows developers to link apps' privacy policies in their public listings. We implemented a crawler to automatically fetch the privacy policies for apps in our corpus. While 6% of the apps were no longer listed in the Google Play Store at the time of the crawl, we found that 40% of the remaining available apps did not have a privacy policy link in their Google Play Store listing. Additionally, for 5% of the available apps, privacy policy links were provided, but resulted in HTTP 404 errors. Ultimately, we were only able to download privacy policies for 55% of the corpus still listed in the Google Play Store. These results alone illustrate how impractical it is to expect users to use privacy policies to make informed decisions about their online privacy: almost half of app developers are likely not meeting their legal obligations to post those policies.

Of the 5,877 pairs in our corpus, 2,646 pairs (45%) had separate policies for the paid and free version. In order to examine differences in each pair of policies, we first performed a `diff` on the policy text, and then manually examined the differences. We labeled policies to highlight the differences in behavior on the free and paid version, such as requesting more permissions, accessing more types of personal data, or adding new third-party libraries. We discarded the pairs in which the only difference between the two policies was the stated name of the app (*e.g., Privacy Policy for App Lite* and *Privacy Policy for App Pro*). We also eliminated pairs in which the privacy policy for one of the apps led to an entirely different document; for instance, links leading to the terms of service, or a FAQ.

Table 5. Differences found in the privacy policy of the free version of an app pair

Differences observed in the privacy policy of the free version	# of apps (% of total)
Has more 3rd parties and gets more data	47 (1.8%)
Has more 3rd parties	43 (1.6%)
Gets more data	5 (0.2%)
Request more permissions	1 (0.1 %)

Out of the remaining 2,499 pairs, 92 (3.7%) contained differences in the privacy policies of the free and paid version of the app, Table 5 shows the type of differences observed in the policies of the free app in comparison with its paid counterpart. We also searched the policies for mentions of COPPA, GDPR, or CCPA to determine if app publishers take privacy legislation into account when informing users of their rights. Only 1% of the apps directly mentioned at least one of these laws.

The analysis of inconsistencies between the privacy policies and the behavior of a given app is beyond the scope of this study. In fact, a majority of free and paid app pairs lack privacy policies for at least one of the apps. Furthermore, we acknowledge that our methods do not take into account the possibility that a single policy discloses behaviors for both free and paid versions (*i.e.,* in different sections within the same text), so further text similarity analysis is required.

6 Consumer Expectations Survey

Prior research has shown that consumers value privacy and prefer control over how companies handle their personal data [28]. This, coupled with the assumption that consumers can avoid the "hidden cost" of a free app by purchasing the paid version, suggests that consumers believe in a "pay for privacy" model of the app ecosystem. We examined this by testing whether consumers are likely to believe that free and paid versions of the same app offer the same privacy and security protections. To that end, we constructed a survey with a mix of open-ended, multiple-choice, and 5-point Likert scale questions, followed by a set of demographics questions.

To probe further into consumers' expectations of free and paid app behavior, we asked participants to choose an app that they would be likely to install, and then showed them Google Play Store listing mockups for free and paid versions of their chosen app. Overall, 387

participants (38.8% of 998) indicated that they would prefer to install the hypothetical paid version of the app, often citing the removal of ads as the primary incentive.

Next, we asked participants, in what way, if any, they would expect the two apps to differ. A significant majority (85.8%) mentioned the inclusion or exclusion of ads between versions. One participant stated, *"I would expect App A [free version] to have ads within the app, and I would expect App B [paid version] to not have any ads"* (P77). This reflects the sentiment of the majority of participants, who correctly expected the paid version to be ad-free compared to its free counterpart. Their expectations about what this meant, in practice, can be further subdivided into sentiments related to annoyance and privacy:

- *"I would expect app B to run better and not track my data for advertising purposes since I paid to avoid them."*
- *"The first one would show ads which would take longer to view my content on Facebook since I have to wait for the ad to be over."*

A minority mentioned differences in app features (2.3%): *"I would expect App A to contain obnoxious ads and potentially mine personal data to provide those ads. I'd expect App B to be a smoother user experience."* A common functionality difference was a better user experience and UI to interact with the app.

To probe further on their expectations on user tracking, we asked participants if they believed that a user's personal data would be treated differently between the two apps. Many participants (28.7% of 998) believed that there would be a difference between the two apps. While explaining how they perceive the difference, about half of these participants mentioned user tracking (46.0%) and targeting for advertisements (56.5%) as ways that the free version would differ from its paid counterpart. Many participants explicitly associated the presence of ads with a degree of tracking and data collection they would expect from an app:

- *"Without ads I would assume that an ad free version wouldn't really have as much use for my personal data because it isn't trying to sell me anything."*
- *"I believe App A's user data may be sold to advertisers to target ads, App B's information would less likely be sold as it's ad free."*

Thus, without the visual cue of an ad, users are likely to incorrectly assume that their personal information is not being collected.

Using a Likert scale, we asked participants to quantitatively express their expectations on the data collection behaviors between the free and paid versions of their chosen app.[5] Overall, we observed that participants expected better privacy and security practices from the paid version of the app. Because our sample size was so large, every comparison was statistically significant, which is why a more meaningful statistic is the effect size. In order of decreasing effect size, we found that participants expected that the paid app is:

- less likely to "share your data with advertisers" ($Z = 18.086$, $p < 0.0005$, $r = 0.573$)
- less likely to "share your data with third-party services" ($Z = 15.305$, $p < 0.0005$, $r = 0.485$)
- less likely to "use your data for secondary purposes" ($Z = 14.008$, $p < 0.0005$, $r = 0.444$)
- less likely to "access more resources than it needs for its functionality (i.e., more permissions)" ($Z = 10.797$, $p < 0.0005$, $r = 0.342$)
- more likely to "have effective privacy controls (features that allow you to specify which data types you do not want the app to collect)" ($Z = 10.024$, $p < 0.0005$, $r = 0.317$)
- more likely to "no longer retain your data after you uninstall the app" ($Z = 9.947$, $p < 0.0005$, $r = 0.315$)
- less likely to "retain your data when no longer needed for the functionality of the app" ($Z = 9.577$, $p < 0.0005$, $r = 0.303$)
- more likely to "protect the data you gave it permission to access" ($Z = 9.298$, $p < 0.0005$, $r = 0.294$)
- more likely to "comply with privacy laws and regulations" ($Z = 8.092$, $p < 0.0005$, $r = 0.256$)
- more likely to "be transparent with you about its data collection and sharing behaviors" ($Z = 7.839$, $p < 0.0005$, $r = 0.248$)
- more likely to "store your data securely to protect it from potential breaches" ($Z = 7.022$, $p < 0.0005$, $r = 0.222$)
- less likely to "share your data with law enforcement agencies" ($Z = 6.593$, $p < 0.0005$, $r = 0.209$)
- more likely to "collect and upload your data securely to maintain its confidentiality" ($Z = 6.382$, $p < 0.0005$, $r = 0.202$)

5 Scale and statements detailed in Section 3.3; all comparisons were made using the one-sample Wilcoxon signed rank test and corrected for multiple testing using the Bonferroni method.

Our results indicate large effect sizes associated with participants' belief that paid apps are less likely to disclose personal data to third-parties or otherwise use it for secondary purposes. Unfortunately, as our analysis shows, this expectation is not met in practice.

7 Discussion

We present the first large-scale privacy analysis of paid mobile apps in direct comparison with their free counterparts. While prior research has looked into consumers' privacy expectations when paying for online services, we present ground truth on how exactly paid versions of free apps are, and in many cases are not, living up to these expectations. Using both static and dynamic analysis, we quantified the differences between declared permissions, bundled third-party libraries, and data sharing over the network. In all three perspectives, we categorized our comparisons of the free vs. paid app pairs into the following: complete overlap, some degree of overlap, and no overlap. We also examined the privacy policies of the pairs to better understand what disclosures, if any, are being made to consumers between versions. To establish a basis of consumers' expectations, we surveyed participants to measure their expectations on the trade-offs between free and paid app versions.

Based on our analysis, free and paid versions of the same app share the same dangerous permissions the vast majority of the time, have the same third-party libraries about half of the time, and collect and send the same sensitive data to the same third parties a third of the time. These findings run counter to the general belief that paying for the app protects the consumer from extensive data collection and tracking. It is even more troubling that there is no easy way for consumers to determine whether a given paid app actually affords greater privacy protections than its free counterpart. In a small minority of free and paid pairs, however, the paid app did request fewer dangerous permissions (Figure 3). This suggests that at least some app developers have made an effort to make sure that paid versions are less intrusive and are more likely to meet consumers' expectations.

Of the three metrics we used in our analysis, the nature of the sensitive data being shared with third parties is the most critical. During our analysis we found that 32% of paid apps shared the same sensitive data with the same third parties as their free counterparts. This data shows that many apps, regardless of their monetization model, disseminate sensitive and personal data with third parties (*e.g.,* advertising networks or analytics services), potentially defying consumer expectations.

7.1 Increased Transparency

Our most significant finding is that paying for an app does not guarantee better privacy or anonymity for the consumer. At the moment, neither platforms nor apps provide any mechanism to inform the consumer about the behavioral differences between paid and free versions of the same app. Thus, more transparency about these differences—if any—are required: consumers have a right to know whether they will obtain any privacy benefits in exchange for a payment.

One would assume that one of the roles of privacy policies is to elucidate these differences. However, we found that a large number of apps—40% of the apps that were still available—do not have a link to their privacy policy in the Google Play Store, demonstrating the diminished user agency in the ability to make informed decisions based on privacy policies—which due to poor readability, are often misunderstood even when read. Furthermore, in a more detailed analysis of the privacy policies' text, we found that only 3.7% of apps exhibit differences in the type of data collected or the number of third-parties embedded in the app, indicating that in most cases, users cannot rely on policies to decide whether the paid version will be less data-hungry.

7.2 Consumer Rights

When consumers pay money for apps, they should be entitled to know what exactly they will receive in exchange for that payment. In Section 6, we demonstrated that many consumers *do* expect to receive better privacy protections in exchange for these payments and that these expectations are not currently being met.

Thus, this research should serve as a wake-up call to app developers, app markets, and regulators. Developers need to do a better job of communicating the privacy practices of their apps, including disabusing users of the notion that paid versions provide better privacy protections when that is not the case. App markets need to make greater efforts to highlight the privacy practices of apps, so that consumer expectations can be better shaped to match reality. Finally, regulators need to bring enforcement actions when consumers are materially misled about apps' privacy practices.

Acknowledgements

We would like to thank our anonymous reviewers and our shepherd Prof. Alan Mislove (Northeastern University) for their detailed and valuable feedback on the preparation of the final version of this paper. The authors also thank Noura Alomar for her feedback on the survey. This work was supported by the U.S. National Security Agency's Science of Security program (contract H98230-18-D-0006), the Department of Homeland Security (contract FA8750-18-2-0096), the National Science Foundation (grants CNS-1817248 and CNS-1564329), the European Union's Horizon 2020 Innovation Action program (grant Agreement No. 786741, SMOOTH Project), the Rose Foundation, the Data Transparency Lab, and the Center for Long-Term Cybersecurity at U.C. Berkeley.

References

[1] Domenico Amalfitano, Anna Rita Fasolino, Porfirio Tramontana, Bryan Dzung Ta, and Atif M. Memon. MobiGUITAR: Automated Model-Based Testing of Mobile Apps. *IEEE Software*, pages 53–59, 2015.

[2] Amina Wagner, Nora Wessels, Peter Buxmann, Hanna Krasnova. Putting a Price Tag on Personal Information - A Literature Review. In *Proc. of the 51st Hawaii International Conference on System Sciences*, pages 3760–3769, 2018.

[3] Sara Angeles. Are Free Apps Safe? https://www.businessnewsdaily.com/4868-free-app-security-risk.html. Archived at https://web.archive.org/web/20181129010454/https://www.businessnewsdaily.com/4868-free-app-security-risk.html. Last Accessed: November 28, 2018.

[4] App Annie. Digital App Economy Forecast: App Annie's App Monetization Report. https://web.archive.org/web/20200218001956/https://www.appannie.com/en/insights/market-data/app-monetization-report-2016/. Last Accessed: February 17, 2020.

[5] AppBrain. Number of Android apps on Google Play. https://www.appbrain.com/stats/number-of-android-apps. Archived at https://web.archive.org/web/20181129003859/https://www.appbrain.com/stats/number-of-android-apps. Last Accessed: August 26, 2019.

[6] Appodeal. Now That In-App Header Bidding Is Finally Here, Is The Waterfall Era Truly Over? (Part 1). https://blog.appodeal.com/waterfall-parallel-bidding-part-one/. Archived at https://web.archive.org/save/https://blog.appodeal.com/waterfall-parallel-bidding-part-one/. Last Accessed: August 27, 2019.

[7] Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick McDaniel. FlowDroid: Precise Context, Flow, Field, Object-sensitive and Lifecycle-aware Taint Analysis for Android Apps. In *Proc. of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 259–269, 2014.

[8] Brian X. Chen. How to Protect Your Privacy as More Apps Harvest Your Data. https://www.nytimes.com/2017/05/03/technology/personaltech/how-to-protect-your-privacy-as-more-apps-harvest-your-data.html. Archived at https://web.archive.org/web/20181129005245/https://www.nytimes.com/2017/05/03/technology/personaltech/how-to-protect-your-privacy-as-more-apps-harvest-your-data.html. Last Accessed: November 28, 2018.

[9] Arytom Dogtiev. App Download and Usage Statistics (2018). http://www.businessofapps.com/data/app-statistics/. Archived at https://web.archive.org/web/20181130221155/http://www.businessofapps.com/data/app-statistics/. Last Accessed: November 30, 2018.

[10] W. Enck, P. Gilbert, B. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. TaintDroid: An Information-flow Tracking System for Realtime Privacy Monitoring on Smartphones. In *Proc. of the 9th USENIX conference on Operating systems design and implementation (OSDI)*, page 393–407, 2010.

[11] Chris Fox. Google hit with £44m GDPR fine over ads. https://www.bbc.com/news/technology-46944696. Archived at https://web.archive.org/save/https://www.bbc.com/news/technology-46944696. Last Accessed: January 21, 2019.

[12] Julien Gamba, Mohammed Rashed, Abbas Razaghpanah, Juan Tapiador, and Narseo Vallina-Rodriguez. An analysis of pre-installed android software. In *Proc. of 41st IEEE Symposium on Security and Privacy (S&P)*, 2020.

[13] Clint Gibler, Jonathan Crussell, Jeremy Erickson, and Hao Chen. AndroidLeaks: Automatically Detecting Potential Privacy Leaks in Android Applications on a Large Scale. In *Proc. of the 5th international conference on Trust and Trustworthy Computing (TRUST)*, pages 291–307. Springer-Verlag, 2012.

[14] Google, Inc. Dangerous permissions. https://developer.android.com/guide/topics/permissions/requesting.html#normal-dangerous. Accessed: August 17, 2017.

[15] Google, Inc. Families - developer policy center. https://play.google.com/about/families/. Accessed: August 31, 2019.

[16] Google, Inc. Permissions overview | Android Developers. https://developer.android.com/guide/topics/permissions/overview.html#normal-dangerous. Last Accessed: August 31, 2019.

[17] Google, Inc. UI/Application Exerciser Monkey. https://developer.android.com/tools/help/monkey.html.

[18] Google Play Store. Comcast cable corporation profile. https://play.google.com/store/apps/developer?id=Comcast+Cable+Corporation,+LLC.

[19] Google Play Store. Comcast profile. https://play.google.com/store/apps/developer?id=Comcast.

[20] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, Serge Egelman. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. In *Proceedings on the 2018 Privacy Enhancing Technologies Symposium (PETS)*, pages 63–83, 2018.

[21] Cecilia Kang. F.T.C. Approves Facebook Fine of About $5 Billion. https://www.nytimes.com/2019/07/12/technology/

facebook-ftc-fine.html. Archived at https://web.archive.org/web/20190817002726/https://www.nytimes.com/2019/07/12/technology/facebook-ftc-fine.html. Last Accessed: August 16, 2019.

[22] Jinyung Kim, Yongho Yoon, Kwangkeun Yi, and Junbum Shin. ScanDal: Static Analyzer for Detecting Privacy Leaks in Android Applications. *IEEE Workshop on Mobile Security Technologies (MoST)*, 2012.

[23] Jeffrey Knockel, Adam Senft, and Ronald Deibert. Privacy and security issues in bat web browsers. In *6th {USENIX} Workshop on Free and Open Communications on the Internet (FOCI)*, 2016.

[24] Li Li, Alexandre Bartel, Tegawendé F Bissyandé, Jacques Klein, Yves Le Traon, Steven Arzt, Siegfried Rasthofer, Eric Bodden, Damien Octeau, and Patrick McDaniel. Iccta: Detecting inter-component privacy leaks in android apps. In *Proc. of the 37th International Conference on Software Engineering (ICSE) Volume 1*, pages 280–291. IEEE Press, 2015.

[25] Martina Lindorfer, Matthias Neugschwandtner, Lukas Weichselbaum, Yanick Fratantonio, Victor van der Veen, and Christian Platzer. ANDRUBIS - 1,000,000 Apps Later: A View on Current Android Malware Behaviors. In *Proc. of the Third International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, pages 3–17, 2014.

[26] Meng Luo, Oleksii Starov, Nima Honarmand, and Nick Nikiforakis. Hindsight: Understanding the evolution of ui vulnerabilities in mobile browsers. In *Proc. of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 149–162. ACM, 2017.

[27] Ziang Ma, Haoyu Wang, Yao Guo, and Xiangqun Chen. LibRadar: fast and accurate detection of third-party libraries in Android apps. In *Proc. of the 38th International Conference on Software Engineering Companion (ICSE-C)*, pages 653–656. ACM, 2016.

[28] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. Internet Users' Information Privacy Concerns (IUIPC): The Construct, The Scale, and A Causal Model. *Information Systems Research*, pages 336–355, 2004.

[29] Aleecia M. McDonald and Lorrie Faith Cranor. Americans' attitudes about internet behavioral advertising practices. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society (WPES)*, pages 63–72. ACM, 2010.

[30] 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)*. USENIX Association, 2019.

[31] Leo Mirani. The amount most people are willing to pay for an app is $0 - until they've actually downloaded it. https://qz.com/129699/the-amount-most-people-are-willing-to-pay-for-an-app-is-0-until-theyve-actually-downloaded-it/. Archived at https://web.archive.org/web/20181114231539/https://qz.com/129699/the-amount-most-people-are-willing-to-pay-for-an-app-is-0-until-theyve-actually-downloaded-it/. Last Accessed: November 14, 2018.

[32] Elleen Pan, Jingjing Ren, Martina Lindorfer, Christo Wilson, and David Choffnes. Panoptispy: Characterizing audio and video exfiltration from android applications. *Proceedings on Privacy Enhancing Technologies*, 2018(4):33–50, 2018.

[33] Matthew Panzarino. Why you should want to pay for apps. https://thenextweb.com/apps/2011/04/24/why-you-should-want-to-pay-for-apps/. Archived at https://web.archive.org/web/20181129005820/https://thenextweb.com/apps/2011/04/24/why-you-should-want-to-pay-for-apps/. Last Accessed: November 28, 2018.

[34] Rajiv Garg and Rahul Telang. Inferring App Demand from Publicly Available Data, 2013.

[35] A. Razaghpanah, R. Nithyanand, N. Vallina-Rodriguez, S. Sundaresan, M. Allman, C. Kreibich, and P. Gill. Apps, Trackers, Privacy, and Regulators: A Global Study of the Mobile Tracking Ecosystem. In *Proc. of NDSS Symposium*, 2018.

[36] A. Razaghpanah, N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, P. Gill, M. Allman, and V. Paxson. Haystack: In Situ Mobile Traffic Analysis in User Space. *arXiv preprint arXiv:1510.01419*, 2015.

[37] Abbas Razaghpanah, Arian Akhavan Niaki, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Johanna Amann, and Phillipa Gill. "studying tls usage in android apps". In *Proceedings of CoNEXT*, New York, New York, December 2017. Association for Computing Machinery.

[38] 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, Santa Clara, CA, August 2019. USENIX Association.

[39] I. Reyes, P. Wijesekera, A. Razaghpanah, J. Reardon, N. Vallina-Rodriguez, S. Egelman, and S. Kreibich. "Is Our Children's Apps Learning?" Automatically Detecting COPPA Violations. In *IEEE ConPro*, 2017.

[40] E.J. Schwartz, T. Avgerinos, and D. Brumley. All You Ever Wanted to Know About Dynamic Taint Analysis and Forward Symbolic Execution (but Might Have Been Afraid to Ask). In *Proc. of the IEEE Symposium on Security and Privacy (SP)*, Oakland '10, 2010.

[41] Suranga Seneviratne, Harini Kolamunna, and Aruna Seneviratne. A Measurement Study of Tracking in Paid Mobile Applications. In *Proc. of ACM WiSec*, 2015.

[42] State of California. Codes display text: Business and professions code - bpc division 8. special business regulations [18400 - 22948.25], chapter 22. internet privacy requirements [22575 - 22579]. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=8.&chapter=22.&lawCode=BPC. Accessed: March 25, 2020.

[43] State of California Department of Justice. California Consumer Privacy Act (CCPA). https://www.oag.ca.gov/privacy/ccpa.

[44] Statista. Distribution of free and paid Android apps in the Google Play Store as of June 2019. https://www.statista.com/statistics/266211/distribution-of-free-and-paid-android-apps/. Archived at https://web.archive.org/web/20190818205551/https://www.statista.com/statistics/266211/distribution-of-free-and-paid-android-apps/. Last Accessed: August 18, 2019.

[45] Takuya Watanabe, Mitsuaki Akiyama, Fumihiro Kanei, Eitaro Shioji, Yuta Takata, Bo Sun, Yuta Ishi, Toshiki Shiba-

hara, Takeshi Yagi, Tatsuya Mori. Understanding the origins of mobile app vulnerabilities: a large-scale measurement study of free and paid apps. In *Proceedings of the 14th International Conference on Mining Software Repositories*, pages 14–24, 2017.

[46] Ailie K. Y. Tang. Mobile App Monetization: App Business Models in the Digital Era, 2016.

[47] Janice Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. In *Proceedings of the 2007 Workshop on the Economics of Information Security (WEIS'07)*, Pittsburgh, PA, USA, 2007.

[48] Connor Tumbleson and Ryszard Wiśniewski. Apktool - A tool for reverse engineering 3rd party, closed, binary Android apps. https://ibotpeaches.github.io/Apktool/.

[49] U.S. Federal Trade Commission. How to comply with the children's online privacy protection rule. http://www.ftc.gov/bcp/conline/pubs/buspubs/coppa.htm.

[50] Eline Vanrykel, Gunes Acar, Michael Herrmann, and Claudia Diaz. Leaky birds: Exploiting mobile application traffic for surveillance. In *International Conference on Financial Cryptography and Data Security*, pages 367–384. Springer, 2016.

[51] Giridhari Venkatadri, Piotr Sapiezynski, Elissa M Redmiles, Alan Mislove, Oana Goga, Michelle Mazurek, and Krishna P Gummadi. Auditing offline data brokers via facebook's advertising platform. In *The World Wide Web Conference*, pages 1920–1930. ACM, 2019.

[52] Yuta Ishii, Takuya Watanabe, Fumihiro Kanei, Yita Takata, Eitaro Shioji, Mitsuaki Akiyama, Takeshi Yagi, Bo Sun, Tatsuya Mori. Understanding the security management of global third-party Android marketplaces. In *Proceedings of the 2nd ACM SIGSOFT International Workshop on App Market Analytics*, pages 12–18, 2017.

A Consumer Expectations Survey

1) Imagine that you were setting up your phone anew and needed to install apps from the Google Play Store. Which of the following, if any, would you be most likely to install?
- Amazon
- Facebook
- Instagram
- Lucky Go
- Pandora
- Snapchat
- TikTok
- I would not install any of the above apps

In the previous question, you selected [App Name]. Now imagine that there are two versions of this app, labeled App A and App B, available for installation.

[Image (Appendix C)]

2) Which app would you be more likely to install?
- App A
- App B

3) Why?

4) In what way, if any, would you expect the above two apps to differ?

5) Do you believe a user's personal data would be treated differently between the two apps?
- Yes
- No
- I'm not sure

6) [If Yes] In what ways do you believe a user's personal data would be treated differently between the two apps?

Consider again the previously presented apps, App A and App B, to answer the following questions.

[Image (Appendix C)]

7) Based on the images, which app do you believe is **more likely** to... *(5-point Likert scale: Definitely A, Likely A, Equally A and B, Likely B, Definitely B)*
- share your data with third-party services?
- share your data with advertisers?
- share your data with law enforcement agencies?
- store your data securely to protect it from potential breaches?
- collect and upload your data securely to maintain its confidentiality?
- be transparent with you about its data collection and sharing behaviors?
- comply with privacy laws and regulations?
- no longer retain your data after you uninstall the app?
- retain your data when no longer needed for the functionality of the app?
- have effective privacy controls (features that allow you to specify which data types you do not want the app to collect)?
- access more resources than it needs for its functionality (i.e., more permissions)?
- protect the data you gave it permission to access?
- use your data for secondary purposes?

8) What is your age?

9) What is your gender?

10) What is the highest degree or level of school you have completed? If you are currently enrolled in school, please indicate the highest degree you have received.
- Less than a high school diploma
- High school degree or equivalent
- Bachelor's degree (e.g. BA, BS)
- Master's degree (*e.g.,* MA, MS)
- Doctorate (*i.e.,* PhD, EdD)
- Other (please specify)
- Prefer not to say

11) What is your marital status?
- Single (never married)
- Married
- In a domestic partnership
- Divorced
- Widowed
- Prefer not to say

12) Do you have children?
- Yes
- No
- Prefer not to say

13) What is your household income?
- Below $10k
- $10k - $50k
- $50k - $100k
- $100k - $150k
- Over $150k
- Prefer not to say

14) What type of smartphone do you own?
- Android
- Apple iPhone
- Windows Phone
- Other (please specify)
- I do not own a smartphone

B Survey Participant Demographics

Our 998 participants identified themselves as belonging to each of the following categories:

Gender
- Female (501 participants)
- Male (484 participants)
- Other (please specify)
  - Non-binary (6 participants)
  - Gender-queer (2 participants)
  - Gender-fluid (1 participant)
  - Two-spirit (1 participant)
- Prefer not to say (3 participants)

Education
- Less than a high school diploma (5 participants)
- High school degree or equivalent (355 participants)
- Bachelor's degree (e.g. BA, BS) (417 participants)
- Master's degree (*e.g.,* MA, MS) (150 participants)
- Doctorate (*i.e.,* PhD, EdD) (30 participants)
- Other (please specify)
  - Associate's degree (22 participants)
  - Some college, no degree (12 participants)
  - Vocational school (2 participant)
- Prefer not to say (5 participants)

Marital status
- Single (never married) (477 participants)
- Married (354 participants)
- In a domestic partnership (104 participants)
- Divorced (55 participants)
- Widowed (5 participants)
- Prefer not to say (3 participants)

If they have children
- Yes (356 participants)
- No (633 participants)
- Prefer not to say (9 participants)

Household income
- Below $10k (69 participants)
- $10k - $50k (360 participants)
- $50k - $100k (354 participants)
- $100k - $150k (128 participants)
- Over $150k (60 participants)
- Prefer not to say (27 participants)

Type of smartphone owned
- Android (526 participants)
- Apple iPhone (469 participants)
- Windows Phone (1 participant)
- Other (please specify)
  - Blackberry (1 participant)
- I do not own a smartphone (1 participant)

C App Survey Screenshots

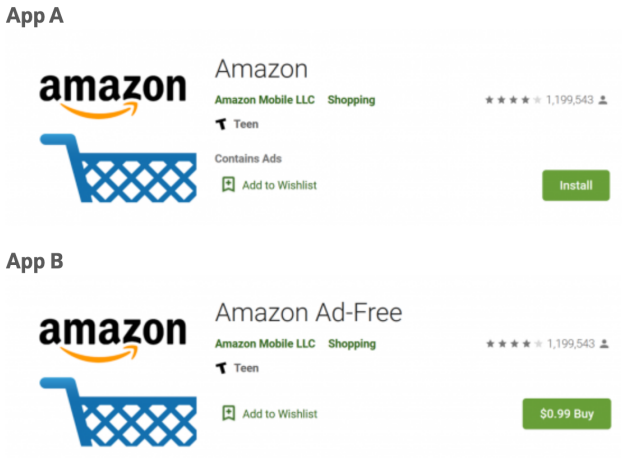

Fig. 8. Amazon

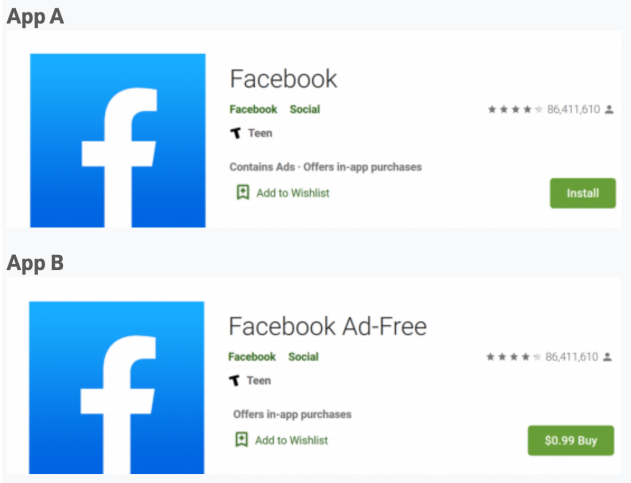

Fig. 9. Facebook

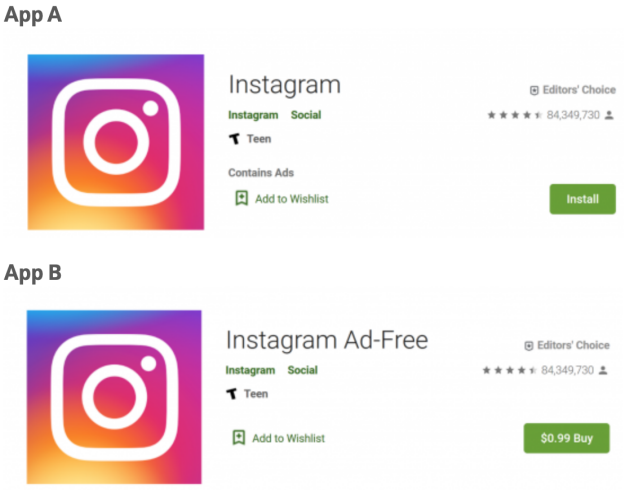

Fig. 10. Instagram

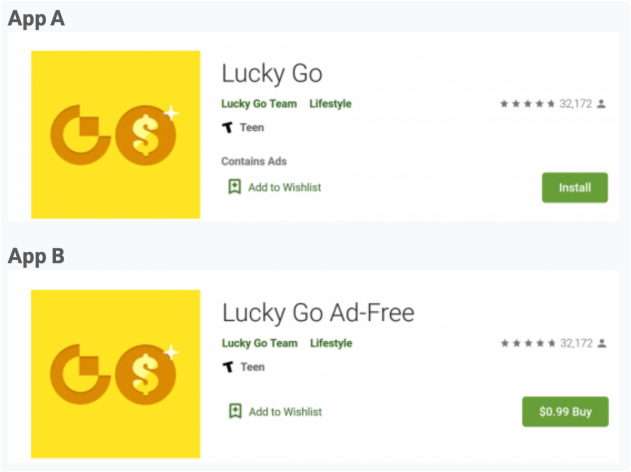

Fig. 11. Lucky Go

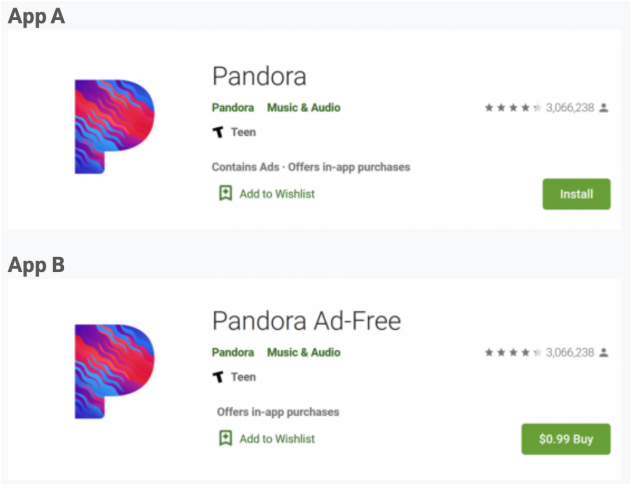

Fig. 12. Pandora

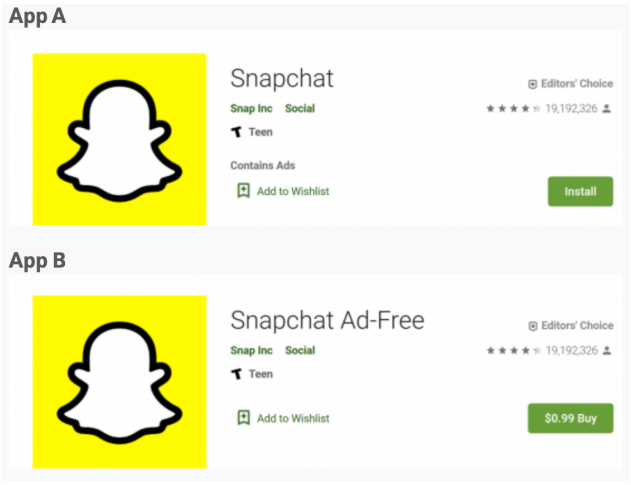

Fig. 13. Snapchat

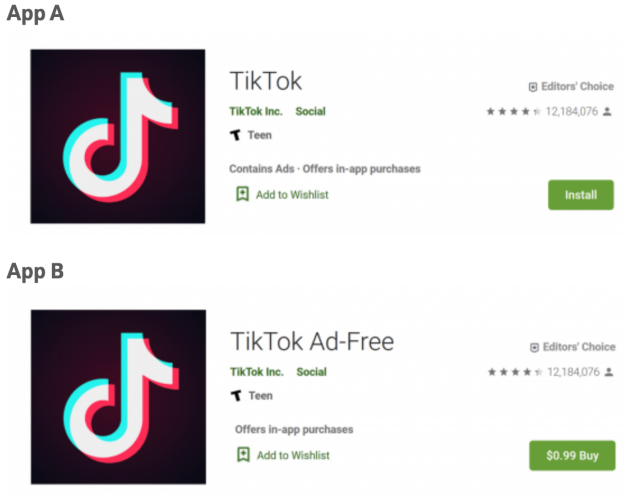

Fig. 14. TikTok