

# Do Privacy Labels Answer Users’ Privacy Questions?

Shikun Zhang  
Carnegie Mellon University  
shikunz@cs.cmu.edu

Norman Sadeh  
Carnegie Mellon University  
sadeh@cs.cmu.edu

**Abstract**—Inspired by earlier academic research, iOS app privacy labels and the recent Google Play data safety labels have been introduced as a way to systematically present users with concise summaries of an app’s data practices. Yet, little research has been conducted to determine how well today’s mobile app privacy labels address people’s actual privacy concerns or questions. We analyze a crowd-sourced corpus of privacy questions collected from mobile app users to determine to what extent these mobile app labels actually address users’ privacy concerns and questions. While there are differences between iOS labels and Google Play labels, our results indicate that an important percentage of people’s privacy questions are not answered or only partially addressed in today’s labels. Findings from this work not only shed light on the additional fields that would need to be included in mobile app privacy labels but can also help inform refinements to existing labels to better address users’ typical privacy questions.

## I. INTRODUCTION

The current legal approach to privacy in the United States concentrates on the concept of “Notice and Choice,” namely the expectation that people are provided sufficient information about the collection and use of their data, and are offered meaningful choices about these practices (e.g., opt-out, opt-in, deletion). Today, “notice” is typically addressed through the publication of a privacy policy. However, there is ample evidence that privacy policies fall short when it comes to informing the public—they are simply too long, too complicated, and often also too vague [20], [6], [21], [27], [25], [26], [16].

For the past dozen years, privacy researchers have advocated the adoption of privacy nutrition labels [9], [10], [11], [12]. Standardized privacy nutrition labels, which succinctly summarize those data practices that

people are most commonly concerned about, offer the promise of providing users with more effective privacy notices than full-length privacy policies. Inspired by this earlier academic research [9], [12], Apple and Google respectively introduced iOS app privacy labels and Google Play’s data safety sections as a way to systematically present users with concise summaries of an app’s data practices. Mobile app nutrition labels also open the door to the development of technology that supports at-scale privacy compliance analysis and the collection of information about data practices across large collections of mobile apps (e.g., an entire app store, but also across specific categories of mobile apps). The labels also raise new research questions, such as investigating the efficacy, usefulness, and usability issues of the labels in the wild, as well as the mismatch between the labels’ disclosures and the privacy choices made available to mobile app users in permission managers. Nonetheless, little has been done to evaluate how extensively the content of current mobile app privacy labels actually addresses the privacy concerns and questions of individuals.

In this paper, we analyze a corpus of privacy questions [24] collected from mobile app users on Amazon Mechanical Turk<sup>1</sup> to determine to what extent these mobile app labels could answer users’ privacy concerns and questions. While there are differences between iOS labels and Google Play labels, our results indicate that an important percentage of people’s privacy questions are not answered or only partially addressed in today’s labels. Findings from this work not only shed light on the additional fields needed to be included in mobile app privacy labels but also provide insight into how well current privacy labels match users’ mental models. We provide recommendations to improve the scope of label content to better match users’ concerns and questions. This paper also exemplifies the misalignment between engineered privacy notices and controls and the privacy questions raised by users.

---

<sup>1</sup><https://www.mturk.com>

## II. RELATED WORK AND BACKGROUND

Mobile devices can collect a wide range of data, including location, contacts, health information, and photos, which can expose sensitive details about a user’s personal life. According to Pew Research, 54% of mobile app users have refrained from using an app, 30% have declined to install an app, and 19% have disabled location tracking on their devices due to privacy concerns [4]. These findings indicate that a significant number of individuals are conscious of and concerned about the manner in which their personal information is processed by the applications on their mobile devices. Unfortunately, current privacy notices in the form of lengthy privacy policies linked from the app are ineffective due to the small display size of mobile devices [26] and poor user comprehension [21], [27], [25]. The prevalent method of presenting privacy information and seeking consent for app permissions management systems on Android and iOS is the “ask on first use” approach. However, this approach cannot provide users with the necessary information to make informed decisions when downloading new apps.

### A. Privacy Nutrition Labels for Mobile Apps

Drawing inspiration from food nutrition labels and standardization efforts in other domains, Kelley et al. pioneered the creation and refinement of a label tuned to privacy [9]. Subsequently, Kelley et al. created short-form privacy nutrition labels for Android apps and found through a lab experiment that the content of the app privacy labels and the timing of the display could assist users in making more privacy-protecting decisions [12]. Later, the findings from the user studies conducted by Balebako et al. suggested even if an app privacy notice includes information that users are concerned about, it is unlikely to be retained if it is only presented in the app store [3].

In December 2020, Apple first introduced privacy nutrition labels for apps in the App Store. Recent research studies have, in particular, examined Apple’s privacy labels given their required use in the Apple app store since the introduction of iOS14 [18], [7], [28], [15], [14], [19]. Zhang and colleagues conducted an interview study on the usability and effectiveness of real-world mobile app privacy labels with end users [28]. They uncovered misunderstandings of and dissatisfaction with the iOS privacy labels that hinder their effectiveness, which includes a confusing structure, reliance on unfamiliar terms, and disconnect between labels and permission controls. In the meanwhile, two research teams, Kollnig et al. and Koch et al., have studied compliance issues of iOS apps not matching their respective privacy labels, including the presence of inaccurate and misleading label information [15], [14].

Another line of work focuses on the creation side of Apple’s privacy labels [7], [17], [18]. Through observing and interviewing iOS app developers, Li et al. identified common challenges for correctly and efficiently creating privacy labels and opportunities to improve their clarity, validity, and consistency [18]. Other recent work involves the development of tools aimed at assisting developers in creating more accurate labels [7] and in-app privacy notices [17].

To the best of our knowledge, the present study is the first to compare the content scope of privacy nutrition label disclosures in the Google Play and Apple’s App Store, with a particular focus on determining to what extent they address typical privacy questions people have.

### B. Privacy Questions About Mobile Apps

As part of a study of mobile app privacy question answering functionality, Ravichander et al. [24] collected a dataset (the “PRIVACYQA” dataset) of privacy questions that people had for a diverse sample of 35 mobile apps. The set of apps were selected to include well-known apps and apps with smaller install bases, also covering a broad range of application categories across the Google Play Store. The study, which involved recruiting Amazon Mechanical Turkers, asked each participant to provide five free text questions per application related to a subset of 35 mobile apps. The study was designed to elicit questions that mattered to participants as they were presented with the name, description, and navigable screenshots of the app as shown in the Google Play Store. The resulting dataset comprises 1,750 questions. Though the authors cannot make any hard claim about how representative this dataset is, it provides a sufficiently diverse collection of privacy questions to warrant comparison with the content of the mobile app labels. Our study leverages this publicly available dataset and explores to what extent these questions can be addressed by iOS and Google Play mobile app labels.

## III. METHODOLOGY

We selected this specific dataset [24] because the questions in this corpus were elicited in a context intended to mimic that of a user examining an app in an app store. Participants were presented with information about an app, including its name, description, and navigable screenshots, similar to what one would find in an app store, and were instructed to ask privacy-related questions about the app.

Our first goal of this study was to understand the nature and topics of questions asked by users in the corpus [24]. We applied thematic analysis as an organizational tool to classify and describe the questions,

as well as a process to interpret, connect, and transform the questions into themes [13]. The lead author first familiarized herself with the data by reading all 1,750 questions. Subsequently, the lead author coded all questions, generated an initial codebook, and met with the second author several times to refine the codebook. The lead author then re-coded all questions individually using the finalized codebook. Given the qualitative and exploratory nature of the study, these methods were deemed sufficient [22]. The final codebook includes 18 themes with 67 codes. The themes and example questions are shown in Table I. Most questions were labeled with one code, while 60 questions were annotated with more than one code, totaling 1,647 codes. The thematic analysis and the generated themes help us to better understand the corpus and also facilitate our next step.

The second objective of this study was to evaluate whether users' privacy questions could respectively be answered by the iOS and Google Play privacy labels. To minimize the impact of app developers' inaccuracies in specifying privacy labels or of apps that may not have been published on both platforms or lack privacy labels, we made the assumption that developers utilized the privacy labels optimally to disclose their apps' privacy practices. This means we did not examine the actual privacy labels within the app stores, but instead evaluated if the iOS and Google privacy labels have the capability to address user privacy questions.

Both authors analyzed and discussed each of the 67 codes (sub-themes) to determine if questions under each sub-theme could be answered using the labels provided by Google or Apple. We randomly sampled example questions for each sub-theme, compared them to the definitions of the Google and Apple labels [8], [2], and reached a consensus on whether those questions could be answered or addressed. We deemed a question fully addressed by the app labels if any part of the label, including definitions, contained implicit or explicit answers to that question. We provide further explanations of implicit answers in Section IV-C. We considered a question partially addressed by the app labels if the presence or absence of a label section provided relevant information but not a complete answer. Table II shows examples of answerable and partially addressed question themes and the corresponding label or definition snippets that can be used to answer these questions. Following the analysis process, the lead author conducted an evaluation of all questions in the data set to determine whether each question could be answered or partially addressed using either the Google or Apple labels.

## IV. RESULTS

### A. Question Themes

We present all 18 question themes resulting from the thematic analysis in Table I. The theme with the highest frequency, accounting for 22% of all questions, pertains to the data being collected by apps. These questions, typically phrased as "Does this app collect X?" were interpreted as "Can this app collect X?" considering that privacy labels do not necessarily indicate actual data practices but rather the potential for data collection. For instance, an app might be able to collect the user's GPS location as long as the user does not deny the app access to their GPS location. Approximately one-sixth of these *data collected* questions, totaling 60, pertain to what types of data are collected. Over 20 questions address the issue of whether the app collects data at all, including questions such as "Do you keep my data and upload to your server?" A handful of questions pertain to whether the collected data is anonymous or not. The remaining *data collected* questions are related to whether specific data types of data are being collected, such as search history, contacts, usage data, etc.

Approximately 12% of questions in the corpus are related to app security, encompassing a variety of topics such as the overall security of the app, inquiries about recent security breaches, technical questions such as whether data is encrypted or whether security protocols are being used, how the app handles passwords, or whether payment data is secure. It is worth noting that even though participants received prompts to ask privacy-related questions, they asked security-related questions as well, indicating that they view their security questions as legitimate privacy questions.

The third and fourth most frequently asked question themes were about data sharing and selling, respectively accounting for 9.2% and 8.6% of all questions. Participants wondered whether their information is shared/sold, what type of information is shared/sold, and to whom. The fifth theme resolves around the types of permissions that apps might need, such as whether a specific permission (e.g., camera, microphone) is accessed or the necessary permissions for the app to function properly. The sixth theme pertains to specific privacy questions related to the functionalities and features of the app. For instance, one question regarding the *TripAdvisor* app says, "Can I review stuff without having my name attached?" Another question about the app *Recipe, Menu & Cooking Planner* reads, "Will anyone see the recipes that I upload?"

Together, these top six themes (one-third of all themes) add up to approximately two-thirds of all questions asked. The remaining 12 themes are listed in Table I and account for just one-third of all the questions.

Question Theme	Types of Questions Under This Theme	Count	%
Data collected	Does the app collect PII, location, search history, payment, texts, health, calls, IP, calendar, other? Is any information recorded?	364	22.1%
App security	How secure is the app? Is my payment information secure with the app? How will my password be stored?	199	12.1%
Sharing	Is my data shared, with whom, and what data is shared?	151	9.2%
Selling	Is my data sold, to whom, and what data is sold?	141	8.6%
Permissions	Any permission required to run this app? Does it have access to my camera or access to my microphone?	140	8.5%
App-specific privacy	Is my status in the app visible to other users?	128	7.8%
Purpose	Will the app use my data for marketing purposes? Why do you need those permissions?	95	5.8%
Who has access	Do app company employees have access to my data? Can the government request my data?	73	4.4%
Privacy risks	Will the microphone secretly be turned on to listen to my surroundings?	64	3.9%
Retention	Will my data be saved permanently? For how long is my data kept?	56	3.4%
Privacy controls	Can I make my profile private? Is there a way to opt out of data sharing?	49	3.0%
Retained method	Will it store any information on my phone? How do you store my data and information?	40	2.4%
Account required	Do I need an account to use this app? Do I have to sign in using a social media account?	37	2.2%
External access	Does [APP] look at other stuff on my phone besides in app? Does the app have access to financial apps I use?	32	1.9%
Deletion	Do I have any rights as far as whether I want my account info deleted?	31	1.9%
Privacy protections	What safeguards does the app use to protect the privacy of my data?	29	1.8%
Privacy policy	Is there a privacy policy? Where can I read your privacy policy?	9	0.1%
Cookies policy	Do you use or collect cookies?	9	0.1%
Total:		1647	

TABLE I: Themes identified, types of questions under each theme, and the number of questions under this theme

### B. Question Themes Mostly Answered by Labels

Table III shows a summary of the question themes that can or cannot be answered by the labels. As seen in the table, both the iOS and Google labels include information on the collected data categories and can answer most of the questions. Note already that not all questions can be answered by the labels and that this varies between Apple and Google. For instance, Google labels specifically mention the collection of IP addresses and calendar information, while iOS labels do not include these data categories. Google labels also allow developers to indicate what data types are optional “where a user has control over its collection and can use the app without providing it” [8], therefore answering a few of questions related to what data are required to use the app. A handful of questions with regards to whether the collected data is anonymous can be answered using only Apple’s privacy labels, as these labels include a section on “data not linked to you.”

While Google labels contain information on security, iOS labels do not mention security at all. App developers can declare optionally in Google labels that their app “has been independently validated against a global security standard... MASA (Mobile Application Security Assessment)” [8]. This review<sup>2</sup> covers a wide range of security-related topics [1], addressing many questions in the app security theme, such as password handling and encryption. However, the review does not cover all the security questions the participants had. For instance, it does not indicate how payment information is stored or cannot help answer questions about whether an app has had a breach in the past. We considered user questions such as “Is the app secure?” or “What

protection do you offer against hackers?” answered if the Google label for an app indicates that an optional review has been conducted.

Both labels provide information on whether an app shares data and the types of data being shared, but neither label directly states with whom the data is shared, only referring to third parties in general. Apple labels only require disclosure of data sharing when it is used for advertising or “tracking,” while Google labels require developers to disclose any non-first-party sharing, unless it is for legal purposes or if the data is anonymous. No data sharing needs to be reported if the action is clearly a “user-initiated action” with clear disclosure and user consent [8] or “it is clear to the user what data is collected” [2].

Both Google and Apple labels address most questions related to the purpose of data collection. While they adopt slightly different definitions of purposes, both include categories such as app functionality, analytics, personalization, and advertising or marketing. Questions such as “Why is my data needed?,” “Will you use my data for advertising?,” and “What does the app do with my personal information?” can be answered by both labels. However, some questions about the specifics of how data is used for personalization or advertising, such as “How are features personalized?,” are not addressed by either label.

Only Google labels contain information related to data deletion [8]. However, some questions pertain to the deletion of specific data types rather than the complete removal of a user’s information. Such questions cannot be answered.

<sup>2</sup><https://github.com/appdefensealliance/ASA/blob/main/MobileAppSecurityAssessment/MobileSecurityGuide.md>

### C. Implicit Answers

Two themes contain questions that cannot be directly answered directly from the labels, but the label definitions contain implicit answers. Questions under that “data collected” theme, such as “Do you keep the data of mine and upload to your company?” can be inferred from the fact that both labels ask developers to declare user data that is transferred out of users’ devices, implying that the data listed on the labels is uploaded to servers. As per both labels, data solely residing on users’ devices are not considered to be “collected.” Similarly, while the labels do not use the term “selling,” questions about whether data is sold are addressed by information provided under Apple’s “data used to track you” section as Apple defines tracking to include user data sharing with a data broker.

### D. Question Themes Not Addressed by Labels

The lower half of Table III lists the themes identified in Table I that are not addressed by either Google’s or Apple’s labels, as evidenced by the zeroes under the Google and Apple columns. The most frequently asked theme (8.5% of all questions) pertains to permissions, with participants asking whether an app accesses specific permission(s) of the phone. It is important to note that “accessing” information in an app does not equate to collecting that information. Data collection only occurs when the information leaves the device. In other words, label entries about data collected by an app do not allow us to answer questions about permissions used by an app. This is the case for both iOS and Google.

Other questions related to permissions, including questions on the necessary permissions needed in order for the app to function (e.g., “What type of permissions does the app need to operate?”), can also not be answered by either label. About 4.4% of questions ask who has access to their information in general or specifically inquire about whether specific entities such as the government or employees of the app company may have access. 3.4% of questions pertain to data retention. 2.8% of questions are related to whether the app requests external access to other apps, accounts, or data outside the app. 2.2% of questions are about whether users are required to create an account or use a social media account to use the app. Nine questions (0.6%) are related to how the app handles cookies. These questions are relatively easy to answer not only because they request factual answers but also because they are generally app-agnostic. The remaining question themes are harder to answer in general.

Around 7.8% of questions pertain to specific app functionality. For instance, a question about the DNA genetic testing app 23andMe reads, “If my genetic data

turns out to be unexpected, can my family see it?” About 3.9% of questions address concerns about privacy violations or potential privacy risks, such as “Does having this on my device create a privacy concern?” Approximately 3% of questions are about privacy controls offered by the app, such as “Can I selectively block scripts on pages that I feel are invading my privacy?” when referring to the *Cake Web Browser* app. Another 2.8% of questions relate to how the app is protecting users’ privacy, such as “Can you guarantee my privacy while playing your game?” The questions under these themes are often app-specific and request more sophisticated answers.

### E. A Comparative Summary of iOS and Google Labels

Our evaluation found that only around 40% of the question themes could be answered by the iOS or Google Play privacy labels, as shown in the top half of Table III. Specifically, 43.2% of questions could be answered by Google Play labels, while 38.6% could be answered by iOS labels. The questions that could be answered by Google Play labels but not by iOS labels pertained to 1) additional data types (such as IP, calendar, and calls), 2) security-related questions, 3) whether the app data can be deleted, 4) optional tags for data that is not necessary for users to provide in order to use the app. In contrast, Apple’s labels provided more information related to data selling, which was not addressed by Google’s labels. Overall, Google’s labels addressed more questions than iOS labels.

## V. DISCUSSIONS

### A. Limitations

Our study investigates the crowd-sourced privacy questions in a public dataset. Even though the questions elicited are specific to the apps present in the dataset, the broad selection of apps and the questions, when analyzed as a whole, can to some extent reflect users’ questions and concerns about apps. We cannot and are not making generalizable claims about our findings since our analysis mainly serves as an exploratory starting point. This paper focuses on the scope of the label contents, and other issues, such as the usability problems or whether labels are reliable or factual, are outside of the scope of this paper. Instead, we try to shed light on the potentially missing elements in label design and the unmatched mental models of users. As privacy researchers, we can only provide an upper limit when assessing whether labels address users’ questions, assuming a complete and perfect understanding of the labels. Our results do not indicate if actual users find their questions answered. Further research is needed to evaluate the effectiveness and efficacy of privacy labels in addressing the questions of users with varying levels of technical expertise.

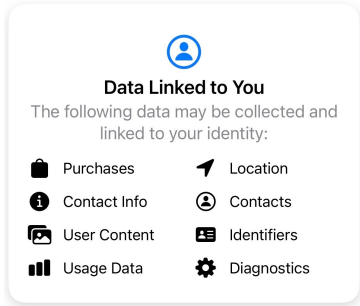
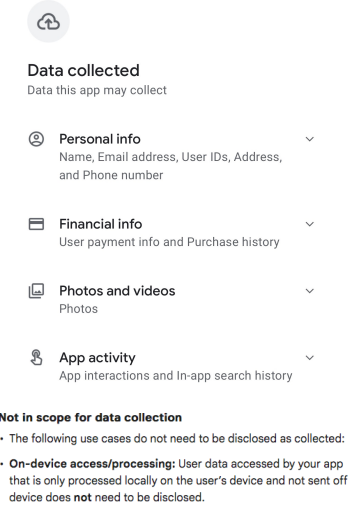
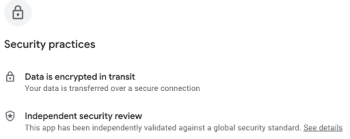
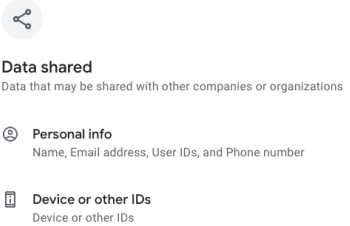
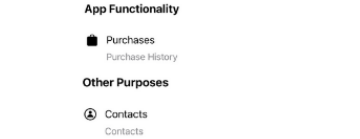
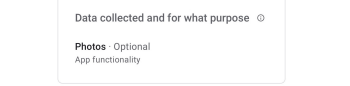
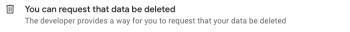
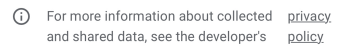
Question Theme	User Question Example	Apple Label	Google Label				
Data collected	What kind of data does [APP] collect?  Do you keep the data of mine and upload to your company?	 <p>"Collect" refers to transmitting data off the device in a way that allows you and/or your third-party partners to access it for a period longer than what is necessary to service the transmitted request in real time.</p>	 <p><b>Not in scope for data collection</b></p> <ul style="list-style-type: none"> <li>The following use cases do not need to be disclosed as collected:</li> <li><b>On-device access/processing:</b> User data accessed by your app that is only processed locally on the user's device and not sent off device does <b>not</b> need to be disclosed.</li> </ul>				
App security	Are you certified to be secure?	N/A					
Sharing	Is information shared with any third parties?	<table border="1" data-bbox="662 930 1019 1014"> <thead> <tr> <th>Purpose</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>Third-Party Advertising</td> <td>Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads</td> </tr> </tbody> </table>	Purpose	Definition	Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads	
Purpose	Definition						
Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads						
Selling	Which information, if any, does the app sell to third parties?	<table border="1" data-bbox="662 1140 1019 1224"> <thead> <tr> <th>Purpose</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>Third-Party Advertising</td> <td>Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads</td> </tr> </tbody> </table>	Purpose	Definition	Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads	N/A
Purpose	Definition						
Third-Party Advertising	Such as displaying third-party ads in your app, or sharing data with entities who display third-party ads						
Purpose	How does this app utilize my data?						
Deletion	Can I remove all my data if I choose not to use this app again?	N/A					
Privacy Policy	Where can I read your privacy policy?	The developer, Google LLC, indicated that the app's privacy practices may include handling of data as described below. For more information, see the <a href="#">developer's privacy policy</a> .					

TABLE II: Sample user questions and corresponding privacy label entry in the iOS and Google Play Stores. N/A means that the question does not have a label addressing it.

### B. Missing Key Information

Our analysis revealed that many question themes were not addressed by the iOS or Google privacy labels. This highlights the need for additional information to effectively address mobile users' privacy concerns or

questions.

1) *Recipient of Information:* Participants wanted to learn who has access to their information and also whom their information is shared with or sold to. They also asked about access to their information by the

Question Theme	# of Questions	Answered by Google	Answered by Apple	Answerable or not
Data collected	364	325	310	Could be answered or partially addressed by labels
App security	199	161	0	
Sharing	151	119	122	
Selling	141	0	125	
Purpose	95	87	87	
Deletion	31	18	0	
Privacy policy	9	7	7	
Permissions	140	0	0	Not answerable by labels
App-specific privacy	128	0	0	
Who has access	73	0	0	
Privacy risks	64	0	0	
Retention	56	0	0	
Privacy controls	49	0	0	
Retained method	40	0	0	
Account required	37	0	0	
External access	32	0	0	
Privacy protections	29	0	0	
Cookies policy	9	0	0	
Sum	1647	717 (43.6%)	651 (39.5%)	

TABLE III: Questions can be answered by Google Play or iOS privacy labels

government or the employees at the app company. This underscores the importance of disclosing the recipients of information, which aligns with the principles of Contextual Integrity [23] stating that it is imperative to disclose the recipient of the information flow. Therefore, privacy labels should include information about entities with whom user data is shared or sold. They should also address common questions, such as explaining to users whether the government or app company employees can access users’ data. For example, messaging app Signal<sup>3</sup> clearly states on its website that its end-to-end encryption keeps users’ conversations secure and that no one, including the government or Signal employees, can read their messages or listen to their calls.

2) *App Permissions*: Before the introduction of runtime Android permissions, the Google Play Store used to display a list of permissions that users needed to agree to before downloading an app. This information was no longer in the Play Store since Android 6. Google has changed its stance on including the permission list in the privacy label and currently does not include it [5]. Our analysis suggests that a good number of user questions pertain to what permissions the app needs or has access to, particularly about location, camera, microphone, and contacts permissions. Currently, users can only view the requested permissions for an app after installing it.

<sup>3</sup><https://signal.org/en/>

Users would also like to know the retention of their information, the availability of privacy controls, and the privacy risks of installing or using these apps.

### C. *Implicit Answers and Mismatching Mental Models*

Our analysis reveals that a few question themes only have implicit answers, which might not be apparent to regular users.

1) *Definition of Data Collection*: Participants used terms like “store,” “save,” and “keep” when asking about data retention and whether their information is being stored. These questions often took the form of “Are you storing any of my information?” or “Do they collect my data and upload it?” This suggests that some users may not equate data collection with storing user data on servers, and it might be beneficial to emphasize that data collection is taken outside of the users’ devices or stored on servers.

2) *Data Selling*: One area of concern among participants was whether their data could be sold or not, with 8.6% of the questions related to this issue. This specific concern is addressed in the recent consumer privacy regulations in California (CCPA/CPRA), which require data controllers to disclose whether and with whom they may be sharing users’ data and to also provide users with privacy options to opt out of such selling. Even though users seem to want to know

specifically about selling, neither the iOS nor the Google privacy labels readily use the term “selling,” making it difficult for users to find answers to these questions. For instance, Google requires disclosing what data is shared with third parties but does not require disclosing the purpose of the sharing or with whom the data is being shared. Apple’s privacy labels come closer to disclosing whether data is sold under the definition of CCPA/CPRA by introducing the concept of “tracking,” which focuses on sharing data with third parties in return for some type of consideration. However, Apple does not explicitly use the word “selling,” making it difficult for users to understand what is being disclosed and in particular whether their data is being sold [2].

3) *What about Security:* Google’s privacy labels already contain security information of an app, currently including whether “data is encrypted in transit” and “optional security review.” App developers can claim in the Google labels that for an app, “data is encrypted in transit: your data is transferred over a secure connection.” This, however, only seems to pertain to a very small number of questions—only 3 out of 199 security questions are about how secure data is during transit. Other aspects of security, such as whether user password is encrypted, are of more importance to users. Although the optional security review covers a wide range of topics, it might be unclear to users what such a security review entails. Furthermore, it is worth noting that this review is optional and not adopted by many apps.

#### D. Privacy Question Answering Functionality

Privacy labels are an important step towards the standardization of data practice disclosures. Prior work found that most users in an interview study reported that they like the concept of privacy labels in the Apple app store [28]. These labels also open the door for compliance analysis [14], [15].

1) *Decreasing User Burden:* Even though privacy labels are designed to help users quickly grasp the important data collection and usage practices without them having to read the text of privacy policies, current labels can already be overwhelming for some apps. For instance, the DoorDash iOS privacy label contains 106 entries of data types organized around 5 purposes and 2 sections. Concurrently, our results show that users have a rather diverse set of privacy questions, with more than half of these questions unlikely to be addressed in current labels. These two findings reveal a challenging tension, with labels appearing already overwhelming yet failing to address a substantial percentage of privacy questions typical users can be expected to have. The paper specifically identifies additional information that one might consider including in labels if one would like to have a better chance of answering people’s typical privacy questions in Section V-B.

Given the amount of label information for each app and the large number of apps on each user’s phone, it is unrealistic to expect users to go through the privacy labels for each app on their phone. There is a need to reduce user burden and to help users quickly locate privacy information that they care about. Future research might want to explore the use of machine learning and natural language processing techniques to automatically extract and analyze standardized notices as a way of providing users with chatbot functionality to quickly answer their questions or refer them to parts of the labels pertaining to their questions.

2) *App Specific Questions:* Our analysis also reveals that many user questions pertain to querying about available privacy controls and app-specific privacy information, which fall beyond the scope of privacy labels. However, these concerns are still relevant to users. Recent development of advanced question-answering chatbots, trained on large language models, presents new research opportunities to provide users with personalized answers to their privacy-related questions regarding specific apps. By doing so, users can make informed decisions without feeling overwhelmed by excessive privacy details. Utilizing these advanced chatbots to answer privacy questions can ease the burden on users to navigate complex privacy information. Further research is necessary to assess the feasibility, accuracy, and comprehensiveness of the answers provided by these chatbots.

## VI. CONCLUSION

We conducted thematic analysis on a dataset of privacy questions mobile app users have about a variety of apps. We evaluated whether these questions can be answered by iOS privacy labels or Google’s data safety sections. Our results indicate that an important percentage of people’s privacy questions are not answered or only partially addressed by today’s labels. We hope that the findings will help inform future refinements of existing mobile app labels as well as the design of more effective ways of communicating data practices to users.

## ACKNOWLEDGMENT

This research has been supported in part by a grant from the National Science Foundation under its Secure and Trustworthy Computing program (grant CNS-1801316), an unrestricted research grant from Google under its “privacy-related faculty award” program. The US Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notice thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official



policies or endorsements, either expressed or implied of NSF, the US Government, or Google.

## REFERENCES

- [1] A. D. Alliance. Mobile application security assessment. [Online]. Available: [\url{https://appdefensealliance.dev/masa}](https://appdefensealliance.dev/masa)
- [2] Apple, "App privacy details on the app store," <https://developer.apple.com/app-store/app-privacy-details/>, 2021.
- [3] R. Balebako, F. Schaub, I. Adjerid, A. Acquisti, and L. F. Cranor, "The impact of timing on the salience of smartphone app privacy notices," in *CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)*. Association for Computing Machinery, October 2015. [Online]. Available: <https://doi.org/10.1145/2808117.2808119>
- [4] J. L. Boyles, A. Smith, and M. Madden, "Privacy and data management on mobile devices," *Pew Internet & American Life Project*, vol. 4, pp. 1–19, 2012.
- [5] I. C. Campbell, "Google is reinstating app permissions list on play store," *The Verge*, November 5, 2020. [Online]. Available: <https://techcrunch.com/2022/07/21/google-app-permissions-play-store/>
- [6] R. Chen, F. Fang, T. Norton, A. M. McDonald, and N. Sadeh, "Fighting the fog: Evaluating the clarity of privacy disclosures in the age of ccpa," in *Proceedings of the 20th Workshop on Privacy in the Electronic Society*, 2021, pp. 73–102.
- [7] J. Gardner, Y. Feng, K. Reiman, Z. Lin, A. Jain, and N. Sadeh, "Helping mobile application developers create accurate privacy labels," *International Workshop on Privacy Engineering (IWPE'22)*, 2022. [Online]. Available: <https://privacyassistant.org/media/publications/IWPE2022.pdf>
- [8] Google, "Provide information for Google Play's Data safety section," <https://support.google.com/googleplay/android-developer/answer/10787469>, August 2022.
- [9] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder, "A "nutrition label" for privacy," in *Proceedings of the 5th Symposium on Usable Privacy and Security, SOUPS 2009, Mountain View, California, USA, July 15-17, 2009*, ser. ACM International Conference Proceeding Series, 2009. [Online]. Available: <https://doi.org/10.1145/1572532.1572538>
- [10] P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor, "Standardizing privacy notices: An online study of the nutrition label approach," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10*. New York, NY, USA: Association for Computing Machinery, 2010, p. 1573–1582. [Online]. Available: <https://doi.org/10.1145/1753326.1753561>
- [11] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall, "A conundrum of permissions: Installing applications on an android smartphone," in *Financial Cryptography and Data Security*, J. Blyth, S. Dietrich, and L. J. Camp, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 68–79.
- [12] P. G. Kelley, L. F. Cranor, and N. M. Sadeh, "Privacy as part of the app decision-making process," in *2013 ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '13, Paris, France, April 27 - May 2, 2013*, 2013, pp. 3393–3402. [Online]. Available: <https://doi.org/10.1145/2470654.2466466>
- [13] M. E. Kiger and L. Varpio, "Thematic analysis of qualitative data: A mee guide no. 131," *Medical Teacher*, vol. 42, no. 8, pp. 846–854, 2020, pMID: 32356468. [Online]. Available: <https://doi.org/10.1080/0142159X.2020.1755030>
- [14] S. Koch, M. Wessels, B. Altpeter, M. Olvermann, and M. Johns, "Keeping privacy labels honest," *Proc. Priv. Enhancing Technol.*, vol. 2022, no. 4, pp. 486–506, 2022. [Online]. Available: <https://www.petsymposium.org/2022/files/papers/issue4/popets-2022-0119.pdf>
- [15] K. Kollnig, A. Shuba, M. Van Kleek, R. Binns, and N. Shadbolt, "Goodbye tracking? impact of ios app tracking transparency and privacy labels," in *2022 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 508–520. [Online]. Available: <https://doi.org/10.1145/3531146.3533116>
- [16] J. Korunovska, B. Kamleitner, and S. Spiekermann, "The challenges and impact of privacy policy comprehension," in *28th European Conference on Information Systems - Liberty, Equality, and Fraternity in a Digitizing World, ECIS 2020, Marrakech, Morocco, June 15-17, 2020*, 2020. [Online]. Available: [https://aisel.aisnet.org/ecis2020\\_rp/51](https://aisel.aisnet.org/ecis2020_rp/51)
- [17] T. Li, E. B. Neundorfer, Y. Agarwal, and J. I. Hong, "Honeysuckle: Annotation-guided code generation of in-app privacy notices," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 3, sep 2021. [Online]. Available: <https://doi.org/10.1145/3478097>
- [18] T. Li, K. Reiman, Y. Agarwal, L. F. Cranor, and J. I. Hong, "Understanding challenges for developers to create accurate privacy nutrition labels," in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: <https://doi.org/10.1145/3491102.3502012>
- [19] Y. Li, D. Chen, T. Li, Y. Agarwal, L. F. Cranor, and J. I. Hong, "Understanding iOS privacy nutrition labels: An exploratory large-scale analysis of app store data," in *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, ser. CHI EA '22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: <https://doi.org/10.1145/3491101.3519739>
- [20] A. M. McDonald and L. F. Cranor, "The cost of reading privacy policies," *IS: A Journal of Law and Policy for the Information Society*, vol. 4, p. 543, 2008.
- [21] A. M. McDonald, R. W. Reeder, P. G. Kelley, and L. F. Cranor, "A comparative study of online privacy policies and formats," in *Privacy Enhancing Technologies*, I. Goldberg and M. J. Atallah, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 37–55.
- [22] N. McDonald, S. Schoenebeck, and A. Forte, "Reliability and inter-rater reliability in qualitative research: Norms and guidelines for cscw and hci practice," *Proc. ACM Hum.-Comput. Interact.*, vol. 3, no. CSCW, nov 2019. [Online]. Available: <https://doi.org/10.1145/3359174>
- [23] H. Nissenbaum, "Privacy as contextual integrity," *Washington Law Review*, vol. 79, no. 1, p. 119, 2004.
- [24] A. Ravichander, A. W. Black, S. Wilson, T. Norton, and N. Sadeh, "Question answering for privacy policies: Combining computational and legal perspectives," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 4949–4959. [Online]. Available: <https://www.aclweb.org/anthology/D19-1500>
- [25] J. R. Reidenberg, J. Bhatia, T. D. Breaux, and T. B. Norton, "Ambiguity in privacy policies and the impact of regulation," *The Journal of Legal Studies*, vol. 45, no. S2, pp. S163–S190, 2016.
- [26] R. I. Singh, M. Sumeeth, and J. Miller, "Evaluating the

readability of privacy policies in mobile environments,” *International Journal of Mobile Human Computer Interaction (IJMHCI)*, vol. 3, no. 1, pp. 55–78, 2011. [Online]. Available: <https://doi.org/10.4018/jmhci.2011010104>

- [27] T. Vila, R. Greenstadt, and D. Molnar, “Why we can’t be bothered to read privacy policies - models of privacy economics as a lemons market,” in *Economics of Information Security*, ser. Advances in Information Security. Springer, 2004, vol. 12, pp. 143–153. [Online]. Available: [https://doi.org/10.1007/1-4020-8090-5\\_11](https://doi.org/10.1007/1-4020-8090-5_11)
- [28] S. Zhang, Y. Feng, Y. Yao, L. F. Cranor, and N. Sadeh, “How usable are ios app privacy labels?” *Proc. Priv. Enhancing Technol.*, vol. 2022, no. 4, pp. 204–228, 2022. [Online]. Available: <https://doi.org/10.56553/popets-2022-0106>