Shikun Zhang, Yuanyuan Feng, Yaxing Yao, Lorrie Faith Cranor, and Norman Sadeh

How Usable Are iOS App Privacy Labels?

Abstract: Standardized privacy labels that 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. With their introduction by Apple in iOS 14 and Google's recent adoption in its Play Store, mobile app privacy labels are for the first time available at scale to users. We report the first in-depth interview study with 24 lay iPhone users to investigate their experiences, understanding, and perceptions of Apple's privacy labels. We uncovered misunderstandings of and dissatisfaction with the iOS privacy labels that hinder their effectiveness, including confusing structure, unfamiliar terms, and disconnection from permission settings and controls. We identify areas where app privacy labels might be improved and propose suggestions to address shortcomings to make them more understandable, usable, and useful.

Keywords: usable privacy and security, interview study

DOI Editor to enter DOI
Received ..; revised ..; accepted ...

1 Introduction

The prevailing legal framework for privacy in the United States revolves around the concept of "Notice and Choice," which is based on the assumption that if consumers are provided with sufficient information about the collection and use of their data, and if they are given meaningful choices, they will be able to manage their privacy adequately.

Notice is typically addressed through the publication of a privacy policy. In practice, users seldom read these privacy policies and those who do often struggle to understand what the text really means. Privacy policies are not just long and difficult to read but also tend to be ambiguous or silent about important issues. To combat those usability issues, privacy researchers have advocated for the adoption of privacy "nutrition labels" as a more succinct and effective way of informing people about data collection and use practices, focusing on a small set of issues that particularly matter to most users [39].

In Fall 2020, Apple announced that with the release of iOS 14, it would require privacy labels for mobile apps [11], reminiscent of ideas introduced and evaluated in research over a decade earlier [39–41]. Google took a similar approach and began rolling out privacy labels to the Play Store in 2022 [53].

As the Apple privacy labels began appearing in the app store, the new labels got mixed reviews in popular media. Reviewers praised the labels for making it much easier to compare privacy practices across apps but pointed out confusing language and jargon in the privacy labels [13] as well as examples of apps that appeared to have inaccurate or misleading labels [31]. Members of the United States Congress wrote to Apple CEO, Tim Cook, to request that labels be improved and reviewed for accuracy [17].

The availability of privacy labels for over 600,000 iOS apps [44] opened the door for researchers to study their effectiveness in the wild. Here we describe an interview study with 24 iPhone users that investigates lay users' experience with, understanding of, and perceptions of iOS privacy labels.

This paper makes the following contributions:

- Based on an in-depth qualitative interview study with lay users, we show that the iOS privacy nutrition labels currently play a limited role in informing or empowering participants to manage their mobile app privacy.
- We identify areas of misunderstanding of and dissatisfaction with iOS privacy nutrition labels that hinder their usability and effectiveness as privacy notices.
- We identify areas where app privacy labels might be improved and propose recommendations to address the shortcomings we observed.

Shikun Zhang: Carnegie Mellon University, E-mail: shikunz@cs.cmu.edu
Yuanyuan Feng: University of Vermont, E-mail: yuanyuan.feng@uvm.edu
Yaxing Yao: University of Maryland, Baltimore County, E-mail: yaxingyao@umbc.edu
Lorrie Faith Cranor: Carnegie Mellon University, E-mail: lorrie@cmu.edu
Norman Sadeh: Carnegie Mellon University, E-mail: sadeh@cs.cmu.edu

2 Related Work

We review related work on usable and effective privacy notices, notices beyond privacy policies, and privacy labels for mobile apps.

2.1 Usable and Effective Privacy Notices

The prevailing legal framework for privacy in the U.S. is built upon the concept of "Notice and Choice" derived from the Fair Information Practices Principles (FIPPs) [63]. Privacy notices are declarations of how entities collect, process, retain, and share personal data. Privacy policies, which tend to be lengthy legal documents that people seldom read and barely understand [34, 46, 52], are the dominant form of privacy notices. However, they are neither usable nor effective [18].

We summarize below the privacy literature on five key criteria for usable and effective privacy notices.

First, the **readability** of privacy notices is crucial for conveying information. Research has repeatedly shown privacy policies are too long and often require unrealistic education levels to read [27, 50, 65], discouraging people from reading them [27, 46, 50, 65, 73]. Research also indicates that concise privacy notices written in plain language tend to be more effective than lengthy privacy policies [22, 35].

Second, effective privacy notices should promote **comprehension** by the intended audience. Privacy policies often use legal jargon and vague language to allow potential future uses of collected data [57], making it difficult for an average person to comprehend the disclosed data practices [1, 14, 58, 72]. Vu and colleagues' eye-tracking study found that participants poorly comprehended privacy policies even if they were written at their level of education [74]. Researchers have proposed non-textual privacy notices in addition to privacy policies to convey privacy concepts, such as various indicators [59] and icons [51], but user comprehension of these notices remains a challenge [37].

Third, **salience** determines the likelihood that people will actually find and pay attention to privacy notices. Effective privacy notices should be prominently displayed and easy to access both initially and when users want to revisit them. An eye-tracking experiment found that participants were more likely to read and understand privacy policy information when it was displayed by default rather than accessible only by following a link [69]. Another study found that a prototype Android app privacy label was more likely to be noticed and remembered by users when displayed after they downloaded an app than when displayed only in the app store [6]. A recent study also indicates that concise privacy notices displayed in a salient way significantly increased user awareness of potentially risky data practices [22].

Forth, **relevance** also impacts the effectiveness of privacy notices. Frequent exposure to lengthy privacy policies containing too much irrelevant information may cause privacy fatigue [15]. Therefore, privacy notices should highlight the most relevant information to their audience, particularly about unexpected, risky data practices [29, 54]. Also, contextually relevant privacy notices tend to more effective [23, 62]. "Just-in-time" notices like mobile app permissions can provide users contextual information when a specific data practice is about to happen, allowing them to make informed privacy decisions when choices are also provided [28, 62].

Finally, actionable information about **control** makes privacy notices more useful. This typically means integrating privacy notices with privacy choices (e.g., consent, control options), allowing users to take actions about their privacy based on the disclosures in the notices [18, 30, 62]. However, caution should be taken when integrating notice and choice, since it may increase users' cognitive burden in privacy decision making [15, 16].

This study examines the usability and effectiveness of Apple's privacy labels considering the above criteria.

2.2 Notices Beyond Privacy Policies

Legal and privacy researchers have long criticized the poor usability and ineffectiveness of privacy policies as privacy notices and have explored alternative approaches [12, 18, 66]. Schaub and colleagues [61] proposed a design space including four design dimensions (i.e., timing, channel, modality, control) that should be considered to design more usable and effective notices.

Some research has shown that shorter privacy notices tend to be more effective in terms of readability and comprehension [49] and standardized privacy notices are more readable and better facilitate comparisons of privacy practices [2]. Gluck et al.'s study, however, found that extreme shortening of privacy notices did not improve readers' awareness of disclosed privacy practices, suggesting a potential length limit for short notices to be effective [35]. Layered privacy notices, which typically include a standardized short notice and a full privacy

policy [18, 47] have been proposed as a way to combine the strengths of different approaches: the readability of short notices, the comparability of standardized notices, and the full details of traditional privacy policies.

"Privacy nutrition labels" or "privacy labels" offer a standardized and succinct way of informing people about data collection and use practices, focusing on a small set of issues that particularly matter to most users. Kelley and colleagues' early studies suggested that privacy labels allow readers to find privacy-related information quicker and more accurately [39, 40]. More recently, privacy labels for Internet of Things (IoT) devices have been proposed to assist consumers in considering privacy and security when purchasing such devices [24, 25]. However, the effectiveness of privacy labels in real-world situations has not yet been investigated.

To that end, this study aims to explore the usability and effectiveness of Apple's privacy labels for apps, one of the first real-world implementations of the privacy nutrition label approach.

Rather than focusing on notice presentation, some research leverages technical approaches, including machine learning, to help users navigate the complex privacy notice and choice landscape. For example, natural language processing techniques can be used to extract key information from privacy policies [8, 55, 70, 76] to help users understand notices and configure choices. Also, privacy assistants have been proposed to reduce users' burden in managing data privacy according to their diverse individual privacy preferences [19, 20, 38, 45].

2.3 Privacy Labels for Mobile Apps

Today's mobile devices are often equipped with sensors and services capable of collecting various data (e.g., location, contacts, photos, health), which may reveal many sensitive aspects of mobile users' lives. A Pew Research survey reported that 54% of mobile app users have avoided an app and 30% have decided not to install an app due to personal data privacy concerns, and 19% of cellphone users have turned off location tracking on their devices [9]. These survey results suggest that many people are at least somewhat aware of and care about how their personal data are handled by apps on their mobile phones. However, the extremely complex mobile ecosystem makes it difficult for people to protect their data privacy on their mobile devices [68]. Even learning about how mobile apps handle data collected via their devices can be difficult. Not too long ago, the privacy notices for many mobile apps were merely links to their privacy policies hosted on websites. Not surprisingly, a study showed that small displays on mobile devices worsen the already poor user comprehension of privacy policies [64].

Currently, app permissions management systems on Android and iOS mobile platforms often serve as a de facto privacy notice mechanism. Under the dominant "ask on first use" (AOFU) model of app permissions, when a mobile app requests to access certain data or resources (e.g., location, microphone) on a mobile device for the first time, the mobile platform will prompt the user with a pop-up permission choice, which can include a concise privacy notice provided by the app developers. AOFU, as an integrated privacy notice and choice, allows mobile app users to make privacy decisions in the actual context of use. It offloads the initial privacy decision making at app installation but can significantly increase user burden later [45, 67, 75]. Also, AOFU does not assist users in determining whether or not to download a new app based on its privacy risks.

Apple's app privacy labels are standardized short-form privacy notices. Kelley et al. designed short-form privacy nutrition labels for Android apps and demonstrated in a lab experiment that participants who were shown the labels in a modified version of the app store would often take label information into account when choosing which app to download [41]. On the other hand, a multi-stakeholder effort to standardize terminology for mobile app labels resulted in a set of terms with definitions that were not precise enough to enable consistent application [7]. Furthermore, Balebako et al. observed that a prototype Android app privacy label based on these standard terms was rarely noticed by users when placed as an image in the app store [6], indicating that labels may inherit the same discoverability problem as traditional privacy notices.

In a recent study examining the creation side of Apple's privacy labels, Li and colleagues found that creating labels was nontrivial and error-prone for app developers [43], indicating room for improved clarity, validity, and consistency of labels. Other recent work proposes tools designed to help developers create more accurate labels [32] and in-app privacy notices [42].

To our best knowledge, our study is the first to investigate the usability and effectiveness of real-world mobile app privacy labels with end users.

3 Method

We conducted 24 semi-structured interviewers over Zoom between mid January and early March of 2022 to explore participants' experiences, perceptions, and understanding of Apple's privacy labels inside the iOS App Store. This study was approved by our institutional review board.

3.1 Recruitment and Screening

We recruited participants online through postings on about two dozen local Craigslists and sub-Reddit forums for geographic locations throughout the U.S. The recruitment posts did not mention or refer to privacy. Potential participants responded to our recruitment posts by filling out a short screening survey to confirm their eligibility (age 18 or older, able to speak English, located in the United States). Participants who had downloaded one or more apps from Apple's App Store in the past three months were qualified, and some were invited to participate in the follow-up interview. We employed a purposive sampling method [71] to ensure a diverse sample of participants based on age, gender, and occupation. Among the 148 people who completed the screening survey, only one person was disqualified due to not having downloaded an app in the past three months.

3.2 Interview Protocol

Participants completed the online consent form prior to the scheduled Zoom meeting. At the start of the Zoom session, participants were informed that they could stop the interview at any time or decline to answer a question, and then they were given the opportunity to ask the researcher any questions. They were instructed not to disclose any personally identifiable information.

The lead author conducted all 24 interviews, which on average lasted 64 minutes. At the end of the interview, each participant filled out a brief post-survey with additional demographic questions. All interviews were recorded via Zoom and transcribed by a commercial transcription service with participants' consent. We sent each participant a $25 Amazon.com gift card via email. The screening survey, interview script, and post-survey can be found in the Appendix.

We first asked participants about their experience using an iPhone and downloading apps, and asked them to walk us through a recent experience downloading an app. Then, we asked participants whether they had ever wondered about or investigated what information apps collected about them and whether privacy was an important factor when downloading apps.

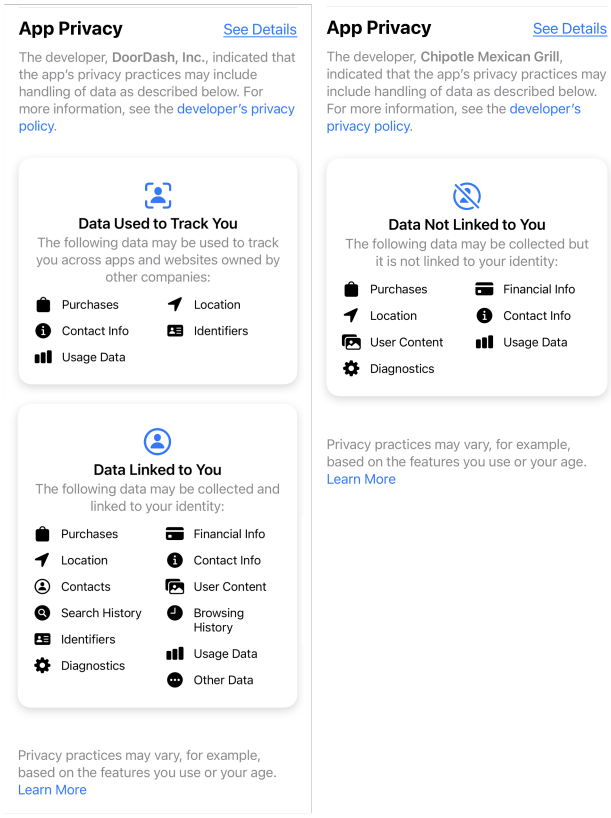

Fig. 1. Screenshots of compact privacy labels from DoorDash (left) and Chipotle (right) in the iOS App Store

Later, we asked participants to share their iPhone screen through Zoom for an interactive session. In this activity, we asked them to visit the App Store and read the compact (Figure 1) and detailed (Figures 2 and 3) privacy labels of two apps in a randomized order. For each app, we asked some specific questions related to the privacy labels, such as their understanding of terms in the labels (e.g., "Data Used to Track You," "Identifiers," "Product Interaction," "Product Personalization"), and their interpretations of the data practices disclosed (e.g., whether the app might share their data with third-party companies for advertising purposes). We also asked them to compare several similar terms (e.g., "data linked to you" and "data used to track you") and explain the differences, if there were any.

After participants completed all the questions for both apps, we asked about their general perceptions of the privacy labels, including whether they found these labels to be useful or not, what they liked or disliked

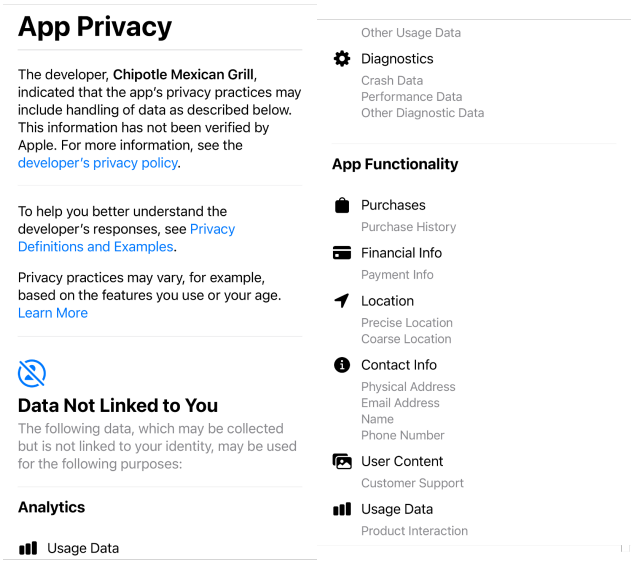

Fig. 2. Chipotle's privacy label in the iOS App Store

about these labels, and whether they would pay attention to these labels in the future. In addition, we also asked whom they considered to be the source of information presented in the labels (e.g., the app developer or Apple). We finished the interview by asking participants about their general privacy concerns and behaviors (e.g., whether they had read a privacy policy or not, whether they had experienced any of their data being misused).

3.3 Interview Design and Piloting

We carefully designed our interviews and set the question order so as to minimize any priming of participants. We iteratively piloted and refined the interview protocol with 5 volunteer participants and 1 recruited participant. Our interview protocol is designed to learn about participant awareness of privacy labels before we mention them, and then to learn about participant understanding of the labels.

In choosing which privacy labels to show, we originally designed the protocol to let participants view the privacy labels of an app they recently downloaded and an app they use frequently. After piloting, we decided to fix the apps that participants reviewed due to the unpredictability of apps and large variances of these apps' privacy labels. We considered one pair of most downloaded apps for each of four app categories (Shopping, Social Networking, Finance, and Food&Drink). Two pilot participants expressed a strong preference for a particular social networking app for privacy reasons. Finance apps also elicited greater privacy concerns from partici-

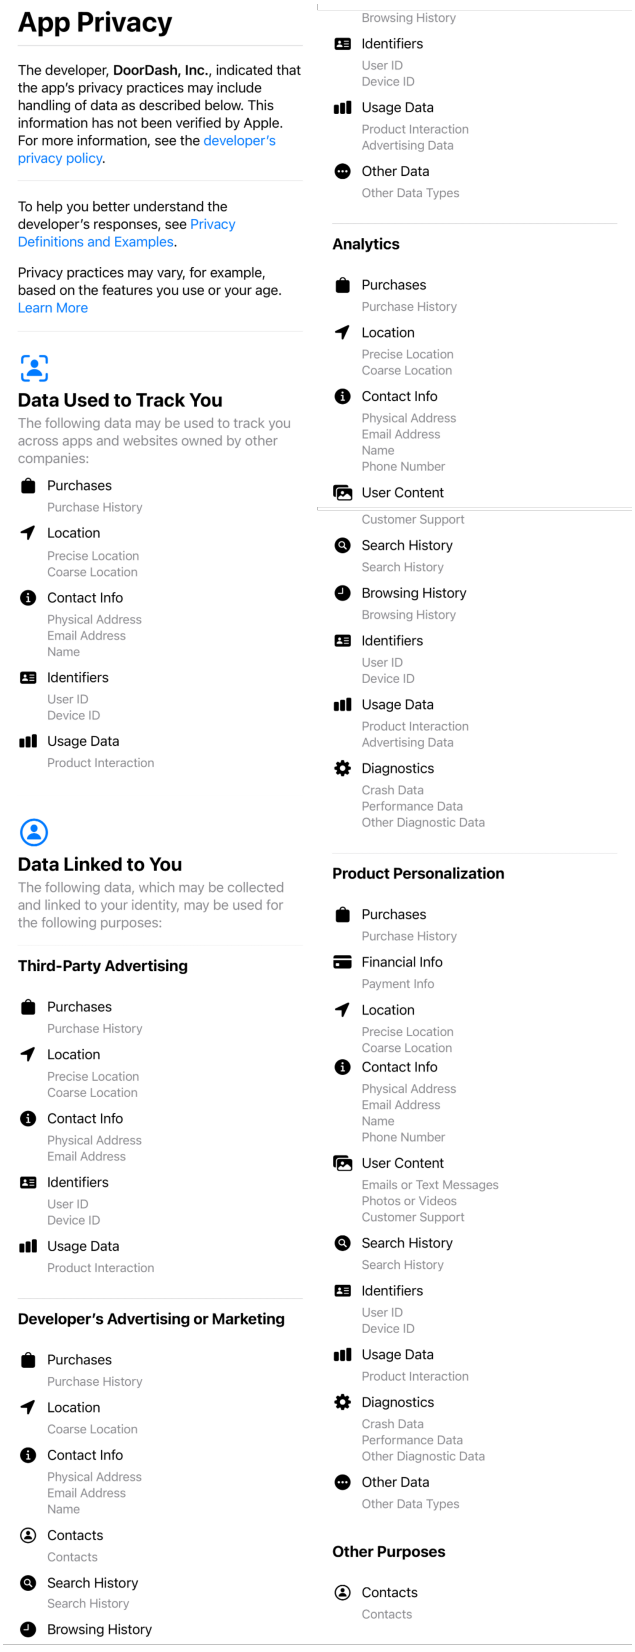

Fig. 3. DoorDash's privacy label in the iOS App Store

pants, potentially making the results more app-specific. We decided to use Chipotle (Figure 2) and DoorDash (Figure 3) because the two apps are similar in nature, but their privacy labels are very different and together cover a variety of terms and topics introduced in Apple's privacy labels.

3.4 Data Analysis

All transcribed interviews were cleaned up by the lead author and analyzed using inductive coding [10]. The lead author met with the research team several times to discuss the first few interview transcripts and generate an initial codebook. The lead author then coded the rest of the data individually using the codebook. During the process, the research team met as needed to improve the codebook and to discuss any perceived ambiguities. Given the qualitative and exploratory nature of the study, these methods were deemed sufficient [48]. The final codebook includes 204 codes across 64 categories. We released the codebook and the redacted interview transcripts with codes via the Open Science Framework.[1]

Due to the qualitative nature of this work, we try to avoid using exact numbers but adopt a consistent terminology to convey the relative sense of the frequency of major themes, similar to prior work [26, 36]. We use the terminology shown in Figure 4 to characterize the frequency of participant responses.

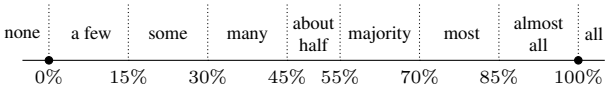

Fig. 4. Terminology used to present relative frequency of themes.

3.5 Demographics

Among all participants, 11 were male, and 13 were female. They represent a diverse background in terms of their age, education, technology experience (whether they said "yes" to the question "Have you ever held a job or received a degree in computer science or any related technology field?"), employment status, and their general understanding of what companies are doing with their data. We present participants' demographic information and some descriptive statistics about their iPhone and app usage in Table 1.

Our study participants were experienced iPhone users with a median of 10 years of usage. Ten of them downloaded one or more new apps within a day of the interview, 10 within the previous week, and 4 within the previous month. Participants estimated having on average 59 apps on their phone, far below the actual average number of 124 apps, which was obtained from their iOS Settings during the interview.

3.6 Limitations

Our study focuses on iOS privacy labels viewed on iPhones and is qualitative in nature based on a purposive sample recruited from location-specific online fora. Our sample skews more educated and younger than the general U.S. population. We aim to describe some of the challenges lay users might encounter as they interact with Apple's privacy labels in the App Store without making quantitative or generalizable claims. We recruited iPhone users with iOS 14 or above so that participants would have all potentially been exposed to the iOS privacy labels, but there could be a sampling bias resulting from targeting such a population.

We did not explore other devices in Apple's ecosystem (e.g., iPad, Mac). However, since Apple uses the same privacy label system with identical terminology and structures but slightly different layouts, many of our results could be reasonably extended to labels on other types of devices. In addition, we chose only two apps (DoorDash and Chipotle) for our study. Future work could investigate whether different apps might have induced different perceptions.

Finally, our study focused on the U.S. population and did not consider other cultural backgrounds. Future work could explore the potential impact of different languages used in the labels or expand the population coverage to account for other cultural factors.

4 Results

We report participants' perceptions about app privacy and Apple's privacy labels, their misunderstandings about the labels, and their suggestions for improvements.

1 https://osf.io/47kzt/

#	Age	Gender	Education	Tech Exp	Occupation	Employment Status	iPhone Usage	# of Apps	Recent App Download	Chipotle	DoorDash	Order
N1	60-69	F	Bachelor's	No	Event planner	Retired	4	82	<1 week	No	No	D
N2	40-49	F	Bachelor's	No	Administrator	Full-time	10	79	<1 week	No	Installed	C
N3	30-39	F	Bachelor's	No	Administrator	Full-time	10	103	<1 day	No	Installed	D
N4	18-29	M	>Master's	No	Neuroscientist	Full-time	14	115	<1 day	No	Installed	D
N5	50-59	F	>Master's	No	Administrator	Full-time	10	71	<1 month	No	Installed	C
N6	18-29	M	Bachelor's	No	Hair Stylist	Full-time	6	36	<1 week	No	No	D
N7	30-39	F	Bachelor's	No	Contracting	Full-time	7	135	<1 week	No	Installed	C
N8	40-49	F	Some college	No	Homemaker	Homemaker	4	102	<1 day	No	Installed	D
N9	40-49	F	Master's	No	Project mngr	Full-time	7	75	<1 day	Installed	No	C
N10	30-39	F	Master's	No	Operation mngr	Full-time	14	126	<1 day	No	Installed	D
N11	30-39	F	Bachelor's	No	Hair Stylist	Full-time	5	180	<1 day	No	No	C
N12	40-49	F	Master's	Yes	Tech	Full-time	10	67	<1 day	Installed	No	D
N13	18-29	F	Associate's	No	Head of HR	Full-time	12	107	<1 week	No	Installed	C
N14	40-49	M	Bachelor's	Yes	Software trainer	Full-time	6	56	<1 week	No	Installed	D
N15	40-49	M	Bachelor's	No	Office mngr	Full-time	12	249	<1 week	Installed	No	C
N16	50-59	F	Bachelor's	No	Art advisor	Full-time	10	73	<1 week	No	No	D
N17	50-59	M	Some college	No	Banquet server	Full-time	6	161	<1 week	Installed	Installed	C
N18	18-29	M	Bachelor's	Yes	Urban planner	Part-time	6.5	57	<1 month	Installed	Installed	D
N19	40-49	M	Master's	Yes	Options trader	Full-time	12	97	<1 day	Installed	Installed	C
N20	40-49	M	Master's	No	HR Director	Full-time	15	319	<1 month	No	No	C
N21	30-39	F	Some college	No	Dental asstnt	Part-time	8	231	<1 day	Installed	Installed	D
N22	18-29	M	Bachelor's	No	CWO	Part-time	8	19	<1 month	No	No	C
N23	30-39	M	Bachelor's	Yes	Reseller	Full-time	12	324	<1 day	Installed	Installed	C
N24	18-29	M	Bachelor's	No	Waiter	Full-time	10	121	<1 week	Installed	Installed	D

Table 1. Participant demographics. "Tech Exp" refers to whether they have held a job or degree in computer science or a related field; "iPhone Usage" refers to the number of years using an iPhone; "# of Apps" refers to the number of apps installed on their phone as shown in iPhone Settings; "Recent App Download" refers to when participants reported having last downloaded an app on to their phone; "Order" refers to the app that participants visited first in the App Store with "D" representing DoorDash and "C" representing Chipotle.

4.1 Perceptions about App Privacy

We present three insights from the first part of our interview prior to introducing participants to the app labels.

4.1.1 Privacy is rarely considered prior to installation

We asked participants to describe their recent or typical process of downloading apps from the App Store. Most of our participants said they already know what app to download when they visit the app store, either through recommendations from friends, articles, or ads. Some participants reported trying several apps and deleting unwanted ones. Some participants also reported searching for keywords in the app store and selecting an app. Privacy was rarely the reason to (not) choose an app during the downloading phase except for a few participants; in comparison, utility, reviews, and cost are the top factors that participants considered before their download. As described by N19 when asked about whether privacy was a reason to (not) choose an app: "Not at that stage. No. I might download the app and then decide, oh, this is too much. Maybe I might delete it. But when I'm looking for an app, no."

Even though privacy was rarely considered when participants were downloading apps, many participants did report having privacy concerns regarding specific apps. Some described removing newly downloaded apps because specific personal information was requested during sign-up. Some indicated that they had deleted apps, such as WhatsApp, Facebook, and T-mobile, after a data breach or learning about privacy concerns regarding these apps.

About half of the participants were not concerned about app privacy, as N10 acknowledged: "I try to ignore that and push that outta my head."

4.1.2 Most have questions about app privacy but lack usable sources for answers

Most participants reported having questions related to app privacy. N23 provided examples of the types of questions he had in mind: "I do think about like, what information are they taking from me? How does it affect me in my life?"

About half of the participants said they would use Google to find answers. For instance, N21 explained try-

ing to use Google: "Google searching, I have tried, but... it wasn't like a professional opinion. There were all other consumers who had wondered that same thing that I was wondering and what they think it does."

N23 mentioned using Google to find answers:

> If there's an app that might be asking for permission for something and I'm like, wait, does this seem right? I will go to a third party, Google or Twitter, so like that, and just do a little bit of research on my own... make sure that what they're asking for is actually legitimate.

Many participants considered privacy policies or terms of services as places to look for answers but reported frustration with them. For instance, N21 voiced her resignation:

> I mean, they do have what they use it for in the terms and conditions. But that's something for a lawyer to look at. I need layman's terms. I need people's terms. So I really don't understand the terms and the conditions. I just hope that they use my information for the best.

Some participants reported looking into apps' data collection practices in iPhone's privacy settings or being informed by iPhone privacy prompts, as detailed by N20:

> The only time I get concerned... when the app has a little pop-up, you know, when I'm using it, and it says this app will collect personal information about you and when would you like it to do so. And it's like, all the time, when you're using it, or never.

Some participants considered the app store as a place to find privacy-related information, as N14 described: "I would hope that in the app store that the part of the description of the app itself would have that type of stuff."

4.1.3 Most unaware of privacy labels

Even though the app privacy labels had been in the App Store for more than a year at the time of our interviews, most of our participants had not seen or read them. Among our 24 participants, a few participants said they had previously read an app privacy label in the iOS App Store. One of them was likely mistaken based on his description of what he had read. Some participants said they became aware of the existence of app privacy labels while scrolling past them in the App Store but did not stop and read them. Most were unaware of the labels. Many participants reported not scrolling down on app pages to see the labels, as N21 commented, "Don't think I go all the way down there." Others simply did not see them. For instance, N20 acknowledged, "No. If I did, then I glazed over it. This is the first time I've ever consciously seen it."

4.2 Perceptions of Privacy Labels

This section reports how participants perceived the app privacy labels after they examined the App Privacy section for both apps.

4.2.1 Most found labels useful

Most participants reported finding the labels useful. For example, N4 said the labels compared favorably to other types of privacy notices:

> I think it is useful because as society at its whole and people individually are caring more about their privacy. So it makes sense that companies and app providers are forced to actually display this stuff in a way to the customer that is not completely incredibly difficult to understand like in a 50-page ToS [terms of service], for example.

N7 considered the label useful but also noted how inconspicuous it is:

> I mean it's useful, I think, if you specifically know what you're looking for. I would think for most people they don't know that this exists, you know, like they might just scroll through and again, see the comment, like this app's amazing, or how many stars, or when it was last updated, but I don't know that people really know that this much information is provided in that section.

On the other hand, some participants did not consider the labels useful. For example, N14 was not satisfied with the labels: "It just definitely feels like it's like a company fulfilling a requirement and not necessarily like trying to tell the consumer what's happening." Similarly, N13 noted, "It's so vague that... it's like a 'there there, don't worry, we're telling you exactly what we're doing', but in reality, you're not telling me nothing."

N2 acknowledged that the labels would be unlikely to impact her use of apps, "I guess it's kind of like a nice to know, but at the end of the day... I probably wouldn't run and delete the app."

4.2.2 Most stated intention to use labels

Most participants reported that they would refer to the labels in the future. N15 described:

> It makes me alarmed the amount of information that's being collected about me through apps. And I'm gonna take a serious look at it after this. I'm gonna look at it and I'll take action. I'll make sure that if there's something easy I can do to set what I'm allowing on my all apps on my phone, if I can do it in one easy step, that's what I'll do. But if not, if I have to go through each app, then I'll have to do that.

Some participants reported that they would look at the privacy labels for apps that they are not familiar with or to compare apps, as described by N18, "If I'm looking for a certain app... if I had choices among other apps that were of similar caliber... this [privacy label] might influence my decision somewhat."

A few participants even deleted apps during the interview or planned to delete apps upon completing the interview.

Those participants who said they did not find the labels useful were also less likely to say they would use them in the future.

4.2.3 Some mistakenly assumed Apple's role in producing the labels

A few participants thought Apple and app developers together produced the labels, while some participants thought the labels were provided by Apple, as N8 commented, "It says it hasn't been verified by Apple. So for me, that's confusing because I would assume that... it was Apple all this time." Most assumed the labels were provided by the app developers.

Many participants correctly deduced that Apple did not review the information submitted by app developers; some learned this from the disclaimer at the top of the labels (Figures 2 and 3). For instance, N23 said, if he hadn't read the disclaimer:

> Given that it's on the app store description page, I would assume it would be information either provided by or directly vetted by Apple because they seem to exert a pretty tight control over things that go on the app store itself at least. And I know that they also do some amount of reviewing of each version of each app that gets up.

The majority believed (wrongly) that Apple had reviewed or verified the information in the labels. N7 explained:

> [Apple is] allowing that app, that product on their system.... So I think Apple, if they're approving that app and they're behind it, then I would think they should be checking [the privacy label].

4.2.4 Many participants would not trust the labels

The majority of participants reported trusting the labels, because some trust Apple. N24 explained, "I trust that their App Store wouldn't be showing me misinformation." Some believed that they had no reason not to trust the labels or had to trust it, as N20 explained, "Why not? I'm going to trust that. It is because I have no evidence that it's not [trustworthy]."

On the other hand, many participants did not trust the labels. N10 explained her reasoning:

> Because of all these things listed that I don't even understand but know that they're wrong, I understand enough to know that it's wrong and it's not okay. I just know that they put it on here because it's mandatory that they do, but they're just doing it because they have to, not because they want me to know. There's nothing with good intentions.

Some participants were concerned about the vagueness of the labels, as N13 explained, "I trust it as far as I could throw it. I mean, it's reliable but very vague. So it's just like... I don't know to what extent you're really doing this."

Both labels included an "other" category, which concerned some participants and contributed to their concerns about vagueness. N4 noted, "I would not trust it a hundred percent. So just because of wild cards, like other data." We discuss this further in Section 4.3.3.

4.3 Misunderstandings of Privacy Labels

In this section, we report the misunderstandings that our participants expressed as they systematically looked at two privacy labels following the interview protocol in the Appendix. We focus on overall understanding, confusing terminology, and vague language.

4.3.1 Overall understanding

At a high level, about half of our participants had misconceptions about the requirements for what is disclosed in an app privacy level, and about half were confused about the label structure itself.

What needs to be disclosed

About half our participants mistakenly assumed the label included all app data collection and usage. However, Apple does not require disclosure of all app data collection: disclosure is optional for data that are not used for tracking or adverting purposes, collected infrequently, and in the app's user interface [3].

On the other hand, about half of our participants said they did not believe or were not sure that the privacy labels show all of the app data collection and handling. They were concerned that companies were only disclosing what they had to disclose. N14 explained:

> It almost seems like the developer, the companies only show you what they're forced to show you. And it's possible that the development of data capturing features and functionality is faster than the regulation to regulate it. And they only do so most of the time if they're in fear of getting in trouble for it.

A few participants were puzzled by why only some of the three boxes were displayed. N18 noted:

> The DoorDash app, it had "data linked to you", you know, "data used to track you"... whereas for Chipotle, they're just saying what's not linked to you.... I understand what they're not using to link to me, but then what are they using to link to me? Right. You can't just assume that because it's not on here, that means they aren't doing it.

Confusing label structure

The compact privacy label for each app, as shown in Figure 1, shows up to three boxes for three categories of data: data used to track you, data linked to you, and data not linked to you. Within each box is a list of specific data types, each accompanied by an icon. If there is no data for a particular category, that box is omitted. Users who click on one of the boxes or the "See Details" link at the top of the label are taken to a more detailed view that shows the same three categories. The "data used to track you" category shows specific data types, with an even more detailed list under each type. However, the other two categories are presented differently, with lists of purposes of data use under each category, and then specific types of data under each purpose, as shown in Figures 2 and 3. This structure was not readily apparent to about half of our participants, who were confused about how to find the purposes and did not understand why specific data types were shown multiple times. For example, after viewing the DoorDash label, which discloses data being used for multiple purposes, some participants said they thought the DoorDash label contained redundant information. For example, N21 noted, "A lot of the information was repeated.... I don't see the need to keep repeating the same that it's gonna collect my purchase, my location, stuff like that."

When asked about what the purpose heading "app functionality" means on the DoorDash label, N10 said:

> It's a word that I figure out the meaning to, but I don't know how it applies to all these things listed under it. So I really can't even guess.

Similarly, N3 expressed confusion when examining the Chipotle label, and noted that she had similar concerns about the DoorDash label, "There's not really a purpose honestly... because they all kind of say the same thing. So like with the last one too, like they were all like basically a lot of repetitive information."

4.3.2 Terminology caused confusion

All participants expressed confusion about one or more terms used in the labels or gave interpretations of terms that did not match Apple's definitions. Only a few participants noticed the link on the label to review the actual privacy term definitions (the second paragraph of Figures 2 and 3), and even fewer of those clicked to view those definitions.

"Tracking" is overloaded

Apple's definition of tracking indicates that data can be used to link with third-party data for targeted advertising or to share data with data brokers. However, this was not always clear to participants. For instance, when looking at the tracking section of the DoorDash label (Fig. 3) N19 could not tell whether data would be used for advertising purposes: "It doesn't say anything about ads? It doesn't say what exactly they will use this data for." Some participants did not associate tracking with data brokers aggregating their information across other websites and companies.

Some participants who glanced at the definition of "data used to track you" were surprised to find that tracking involves sharing information with other parties rather than just collecting data on either the user's location or website usage. For example, N1 said tracking was "tracking me like where I go," while N5 described tracking as, "I think my usage, the frequency of ordering or using the site and identifying where I am. I think that's all tracking. I think tracking is what are my habits." After being told that the definition included sharing, she said, "I don't like that at all. Like now that you say it, I'm thinking I'm deleting this app [DoorDash]. I think it's too intrusive."

N7 was also surprised and concerned by the definition of "data used to track you":

> It's concerning. I will say that I didn't realize. I kind of believed that they would track like, oh, she looks up dogs. So then just random dog things would pop up like ads or like, you know, so pick up kind of like what I was searching. But seeing this, it goes way deeper than that, you know, they're actually collecting the information, which to me is kind of scary.

Confusion around "data (not) linked to you"

The three boxes in the compact label seemed to confuse participants, despite the explanations included in each box.

Partially due to not understanding tracking, most participants could not explain the difference between "data used to track you" and "data linked to you," or their explanation was inconsistent with Apple's definitions, shown in Table 2. N2 admitted, "Well I don't know what the difference is.... Maybe they're just required to have both sections."

N4 was confused about the same data categories shown under these two headings in the DoorDash label, and he wrongly assumed "data linked to you" implied data not being shared:

> I was under the perception that data that is linked to me has more identifiers that make me non-anonymous, for example, my physical address and my name. But then I read that this is also in "data used to track me." So out from the descriptions "the following data might be used to track you across apps" and "the following data, which may be collected and linked to your identity may be used for purposes." Difficult for me to understand. I can imagine that this data is perhaps just put into a database to create a profile of me, but I don't a hundred percent understand what it is."

Term	Definition
Data linked to you	Data that is linked to your identity (via your account, device, or other details). "Personal information" and "personal data" as defined under relevant privacy laws are considered linked to you. In order for data to not be linked to you, the developer must avoid certain activities after they collect it: 1) They must make no attempt to link the data back to your identity. 2) They must not tie the data to other data sets that enable it to be linked to your identity.
Data not linked to you	Data that is not linked to your identity (via your account, device, or other details).

Table 2. Apple's definitions of data linked to you and data not linked to you [4]

"Data not linked to you" on the Chipotle label also confused about half of the participants; most participants were confused after they saw contact information shown under this heading as in Figure 2. The inclusion of contact information in this category is possibly an error, as N13 observed:

> It states it's not linked to you, but obviously it is linked to you because it's your personal information, like your address, your email address, your phone number, and your name. These are all very personal things that are specifically you, as opposed to them stating this is like completely anonymous. So I don't know if they're like lying, if that's like a fine line as in, "we collect this information for our app purposes only, and then the information that we share with other companies that's not linked to you is you know what we are sharing."

When users observe potentially erroneous information in the labels it may undermine trust in the labels, as noted by N21, "Now I feel like you're lying to me because I don't see how you're collecting my name but then telling me it's not linked to me."

Entangled and overlapping definitions

Most participants were particularly confused about the terms, app functionality, analytics, and product personalization, shown in Table 3. They either admitted that they did not know the definitions of these terms or gave their own definitions in which they mixed up the terms.

For example, N19 gave a definition for app functionality that confuses it with analytics: "By collecting this data about me and gazillion other people, they analyze it and they might change the functionality based on the patterns that they see."

Term	Definition
Analytics	Using data to evaluate your behavior, including to understand the effectiveness of existing product features, plan new features, or measure audience size or characteristics.
App Functionality	Such as to authenticate you in the app, enable features, prevent fraud, implement security measures, ensure server uptime, minimize app crashes, improve scalability and performance, or perform customer support.
Product Personalization	Customizing what you see, such as a list of recommended products, posts, or suggestions.

Table 3. Apple's definitions of analytics, app functionality, and product personalization purposes [4]

Term	Definition
Third-party Advertising	Such as displaying third-party ads in the app or sharing data with entities who display third-party ads.
Developer's Advertising or Marketing	Such as displaying the developer's own ads in the app, sending marketing communications directly to you, or sharing data with entities who will display ads to you.

Table 4. Apple's definitions of advertising-related purposes [4]

Term	Definition
Product Interaction	Such as app launches, taps, clicks, scrolling information, music-listening data, video views, saved place in a game, video, or song, or other information about how you interact with the app.

Table 5. Apple's definition of product interaction [4]

N1 mixed up app functionality with product personalization, describing app functionality as "customizing my user content." N3 tried to define analytics as "for them to see how like the app is working, I think," admitting, "I don't really know what analytics is. And I don't really even know what the difference between a lot of these."

A majority of participants could not tell the difference between third-party advertising and developer's advertising, shown in Table 4. N8 described her confusion:

> The advertising and the third party? Yeah, that's a good question because it kind of feels like it means the same thing. I feel like the third-party advertising is somebody other than DoorDash. And then the developer's one is, I don't know, to tell you the truth. I don't know the difference.

Some participants mistakenly thought that developer's advertising did not involve third parties, as N4 suggested:

> Theoretically, the difference should be that third-party advertisement has the purpose of monetizing me and making money out of me by showing me ads, whereas developer's advertising or marketing is for developing a better product in the end. So it's not shared with a third party, but it's shared with DoorDash itself and when DoorDash wants to create a better product out of it.

Unfamiliar terms for frequently used data types

Three frequently used terms describing data types often confused participants.

The data category "user content," for which Apple does not provide a definition, was called out as confusing by about half of the participants. Some were also confused by the term "customer support" below the bold heading for user content in both labels. For example, not knowing what the heading meant, N10 took a guess: "User content and customers support. Maybe they're like blocking me from getting in touch with customer support.... I don't know. That's ridiculous. I really dunno what it is."

Some participants did not understand the term "identifiers," for which Apple does not provide a definition, and a few looked to the icon next to the term for answers, noting it looked like some sort of a photo ID. N4 explained, "I don't know what identifier is, but based on the icon... I'm guessing like identity-related information could be phone, name, email, etc." On the other hand, N10 was confused by the icon, "What are identifiers? That looks like a license."

"Product interaction" was also confusing to many participants, and a few of them erroneously believed that product refers to items they viewed or bought, as N18 explained:

> I think the product refers to like an item, like anything tangible, right? Cause you know, for DoorDash, you're purchasing items through them. Typically it's like food or drink. So I guess to that end, what specific restaurants am I ordering from frequently? What foods am I buying frequently if I am ordering online, stuff like that.

Other terms for data types that confused participants included "usage data," "diagnostics," and "coarse location."

Term	Definition
Browsing history	Information about the content you have viewed that is not part of the app, such as websites.
Search history	Information about searches performed in the app.

Table 6. Apple's definitions of browsing and search history [4]

4.3.3 Vague language on labels

The "other" category is alarming

Almost all participants responded with confusion when they saw terms in the label related to other data types or other purposes since they did not know what these other categories could entail. About half of the participants said these terms made them anxious or decreased the utility of the label. N3 asked, "Why even write all the details, if you're just gonna say other?" Similarly, N22 commented:

> I don't like it. This section seems to be like an extra section for people who want to be clear on what's being used, and then they give vague answers. Anyway, it doesn't make it pointless, but it makes it less useful.

N4 considered this a way for companies to make sure they covered everything that was required:

> I don't think that a company would blatantly lie to me here. This, in my opinion, would be bad intent of the company. And I don't think that the company would do this nowadays because it would be very, very stupid if this comes out like Dieselgate, you know, but as long as they can put something in there, which is other data, I can imagine that they are covered anyways, you know, because what is other data it can put in anything.

The scope of browsing and search history is unclear

Most participants made assumptions about browsing and search history that did not align with Apple's definitions, shown in Table 6, or became confused or concerned when they learned about the definitions.

N19 looked at the DoorDash label and tried to interpret the meaning of browsing history: "Is it the DoorDash website's browsing or any browsing? Is it the DoorDash app or any app? I don't know." The majority of participants were concerned when they learned that browsing history includes "content you have viewed that is not part of the app." N4 commented, "So then I think this should be a bigger topic that is communicated overall, that just by having food delivered to you, you are showing the world what you browse on your smartphone."

Even though search history is defined to be only within the app, some participants assumed it referred to other searches as well. For instance, N17 commented:

> I would hope that it's only keeping track of searches on my apps. In other words, DoorDash, I go on there and type in "tacos." It will come up with tacos. But it doesn't actually say that. So it says "search history." It could also be keeping track of all the searches I do on Google.

User content: emails, texts, photos, or videos

Many participants were concerned or not sure about to what extent the data listed under user content could be used in the DoorDash label. For instance, N22 commented, "I'd hope that this user content section is just stuff in the app for customer support purposes. But if it's looking at your emails or texts or photos or videos outside of the app, that's very disconcerting."

Many participants were disturbed by this kind of data collection, as N8 explained:

> User content, yeah. They're gonna check how I use the app, how many purchases I'm making, where I am at the time that I'm using it. I don't know what they would need emails or text messages or photos or videos. I don't know why they would need that. I didn't know they would have access to my photos. That's kind weird. That's kind of creepy. I don't know what food has to do with my photos.

Most participants were concerned to see data collection that they perceived as unrelated to the purpose of the app, such as their contacts being used by DoorDash.

4.4 Suggested Improvements

Although most participants liked the label concept, they had a number of concerns, as discussed in the previous

Term	Definition
Emails or text messages	Including subject line, sender, recipients, and contents of the email or message.
Photos or videos	Your photos or videos.
Audio data	Your voice or sound recordings.
Customer support	Data you generate during a customer support request.
Other user content	Any other content you generate.

Table 7. Apple's definitions of terms under "User Content" [4]

sections, as well as ideas for improvements. In this section, we report on improvements to the labels suggested by study participants.

4.4.1 Better structure for data and purposes

As discussed in Section 4.3.1, participants found the structure of the label confusing, and many had the misconception that some sections were redundant. Underlying this confusion is that the label designers mapped a multi-dimensional space of data types and purposes onto a list-based label representation. A matrix or tabular approach might offer a more compact and intuitive representation, as N5 suggested: "I do think those last two sections [app functionality and product personalization] were very redundant, and so I think they could do it by like a table with a bunch of checkboxes...."

4.4.2 Easy access to definitions and contextualized examples of data collected

The majority of participants suggested making the definitions of label terms more accessible. Currently, the definitions of many of the terms are available through a link from the detailed view to a web page with the definitions all in one place. However, many participants wanted to see each term linked to its definition, perhaps appearing through a hover. For example, N4 explained:

> It would be cool if this information would be hot linked there, you know, like there is this symbol with an "i" in it, which means information, you know, that would be cool. So like browsing history, for example, click, you know, I'm like, "Ooh, what does this mean?"

Others, including N8, suggested using "terminology that's just easier to understand." Many participants suggested providing concrete examples of the data being collected. N20 suggested:

> If I click user ID... I know what a user ID is, but tell me which user IDs are you tracking. So Is it me, my wife, and my kids? ... What exactly are you tracking? You know, is it my phone? Is it my watch? Is it my iPad? You know what's all really linked to user ID?

N17 also suggested that specific examples should be listed for "other" categories to ease concerns about what might be included.

4.4.3 Embedded actions and controls

Many participants voiced their frustration with the labels due to their lack of controls and suggested controls to turn off some of the data collection. N19 explained:

> I don't like it that it's too long and you can't really take any action on it. It's really just informational and you can't really turn it on off, etc. at this level.... I'd rather just be able to turn off whatever is not required.... If any of this tracking is optional, I wanna be able to turn it off.

N2 described in detail what she envisioned:

> Maybe by letting you check off things that you don't want included or make it easy to like opt out of all of this. Well, every app asks you when you install it if they can like track and share your stuff, so make it easier. So I don't have to learn or go hunting.... It could be like a checkbox or a radio slider or something....

4.4.4 Access to more information

Participants suggested other topics related to app privacy that they would like to see added to the label.

Some participants asked for information about data retention, as N18 articulated, "How long is this information stored for, right. If there was a clear understanding that your information's gonna be stored for 30 days, that would probably give me a lot more solace than not knowing right now."

Some participants were interested in knowing to whom the apps are sending their information. N7 said she would like to know "where exactly is it being [sent]? Is it being sold? Is it being just shared back and forth so that there's this hub that everybody uses?"

A few participants wanted to know where they could have questions answered or read more detailed information. N3 wondered:

> So where would I go to like ask somebody or chat with somebody or like, there should be like another link that takes you to dive in deeper if you wanted to know, because I don't know what I would even like, there's no contact information. And who would I even ask about this? It's kinda useless if I do have a question and I don't know who to ask them, it kind of seems a little useless.

Finally, a few users wanted to better understand the privacy-related implications of using the app and any measures the app was taking to protect their privacy. N14 asked, "How would you see this occur or affect you? Like because of having this app, these are the things

that are happening to you, like you're seeing targeted ads, you know. It relates to the user more."

5 Discussion

With hundreds of thousands of mobile apps now featuring privacy labels in iOS 14, these labels are for the first time available at scale to mobile app users. The introduction of privacy labels is an important step towards empowering users to better understand mobile app data practices that matter to them. At the same time, in their current deployment and configuration, these labels are not as usable or effective as they could be.

5.1 Helping Users Comprehend Complex App Privacy Practices

After examining the labels in our study, almost all participants learned new things that they did not know before and appreciated the existence of the labels. Also, about half of the participants regarded the privacy labels as useful and most reported being likely to use them in the future.

However, the labels suffer from confusing terms and definitions (see Section 4.3.2), which led to a range of misunderstandings. In addition, vague language (e.g.,"other" category, user content) impede participants' understanding of the actual data practices (see Section 4.3.3). These findings clearly demonstrate that Apple's privacy labels still fail to fully support user comprehension of the disclosed app privacy practices.

The linear structure of the labels, which is presented differently in the compact and detailed views, seems to do a poor job of communicating the multi-dimensional space of data practices where multiple categories of data (each of which is represented in a multi-level hierarchy) are used for multiple purposes. A tabular representation may be more compact and intuitive [40, 60], although the small form-factor of mobile devices may present design challenges. Additional work is needed to better understand which label elements are most important to users so that the compact version might focus on those elements.

Our findings on end-user misunderstandings extend recent studies that showed how app developers often struggle with privacy label definitions (e.g., interpretation of terms such as "tracking") and how this hampers their ability to create accurate labels [32, 43]. Our study focused on lay users who lack technical expertise and experienced a high level of confusion.

The addition of links or hover text to provide more ready access to definitions of terms and examples might aid comprehension.

With the recent rollout of Android privacy labels, we have observed that definitions of terms such as "tracking" are not completely consistent on the Android and iOS platforms. It would be helpful if the industry were to adopt standard terms and definitions for privacy labels, empirically tested with both developers and lay users. An earlier multi-stakeholder effort led by the U.S. Department of Commerce resulted in a standard set of terms for app transparency, but did not include terms describing purposes of use and the terms were not updated after user testing found them to be confusing to both experts and lay users [7].

5.2 Improving Privacy Labels' Salience

Even though the privacy labels were introduced in Apple's App Store over a year before our study, the majority of our participants were still unaware of them. Our finding shows that the discoverability of privacy labels on each individual app's page in the App Store is low, even for participants who said they were concerned about mobile app privacy. As currently deployed, users have to scroll past several sections, including images, Ratings & Reviews, and What's New, before finding the App Privacy section. Our findings on discoverability are corroborated by prior research that has shown that the location and timing of privacy labels and indicators can have a large impact on whether users pay attention to them [6, 23, 37].

Our results suggest the need for more prominent placement of privacy labels, consistent with recommendations to display concise privacy notices in salient ways [22]. Alternatively, it would be beneficial to add standardized indicators (e.g., links, icons) to signal the existence of these labels during users' app installation decision-making process.

There is also a need for additional mechanisms to bring users' attention to privacy labels for apps that users already have on their phones. For example, iOS privacy nudges [5] about background app data collection, just-in-time app permission requests, and iPhone permissions setting interfaces are potential places to include links to the privacy labels that would increase both awareness and convenience.

5.3 Promoting Privacy Labels' Role in App Privacy Management

Another key complaint from participants is that the privacy labels do not offer control options (see Section 4.4.3). Some participants reported being disappointed that even after learning the information presented by the labels, they were not provided with any actionable steps they could take. Information on the labels is not readily accessible in the permission settings (i.e., the permission manager) where users decide which permissions to grant to each app. It would be helpful if the permissions manager included the relevant information for each app that appears in the app store label. For instance, the iOS app tracking permission could be incorporated into each label with a toggle control.

In addition, the controls offered to users in the current permission manager are not aligned with the information conveyed in the privacy labels. For example, while a privacy label might inform users that their location information might be used by the app for multiple purposes such as for the app's core functionality as well as for advertising purposes, users do not have the option to grant an app access to their location for one purpose and not for another (e.g., granting access for the core functionality but not for advertising purposes). A few participants were puzzled by DoorDash never requesting the Contacts permission despite listing it in the privacy label. Even worse, when users deny a particular permission, for example, location access, some apps might still be able to extract location-related information from IP addresses, metadata associated with uploaded user photos, WiFi connections, etc [21, 33]. Such misalignments between the disclosures made in privacy nutrition labels and privacy controls made available to users create another potential source of confusion.

Furthermore, if users have already selected the global setting to turn off app requests to track, it is unclear whether any of the tracking indicated in the "Data Used to Track You" section could happen or not. It might be helpful to include a toggle to allow users to turn off tracking directly in the label and indicate appropriately whether the user has previously configured that setting.

Privacy labels are shown within the descriptions of individual apps in the App Store, but no functionality is provided to enable users to compare apps or look for equivalent apps with less invasive or more desirable data practices. The App Store should enable users to search for apps that meet certain privacy criteria, for example, filtering similar game apps that do not collect any location information or picture editing apps that do not involve sharing user information with data brokers.

5.4 Reducing User Burden in App Privacy Management

Ultimately, privacy labels are designed to empower users to quickly find answers to some of their most common questions and save them the time and effort that would be required if they had to read the text of privacy policies. Even though privacy labels offer the promise of providing users with more succinct and more effective notifications, given 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 one at a time. Prior work using machine learning and natural language processing techniques to automatically extract and analyze disclosure statements from the text of privacy policies [55, 70, 76], including privacy question answering functionality [56], has been technically challenging. With the help of these standardized notices, it will be more feasible to automatically extract relevant privacy disclosures, which in turn can support chatbot functionality to quickly address users' questions or refer them to parts of the labels pertaining to their questions.

Another way to decrease user burden is to leverage the operating system or a personal privacy assistant to act on behalf of users instead of relying on users to manually configure every app setting. Users could be selectively notified about the types of data collection disclosures that they would like to be reminded about and only show users relevant disclosure information that they personally care about [45, 67, 75].

6 Conclusion

While iOS app privacy nutrition labels offer the first wide-scale deployment of standardized short-form privacy notices, our qualitative interview study highlights the barriers that prevent these labels from achieving their desired impact when it comes to actually helping users. Findings from this work provide the basis for concrete recommendations to refine existing labels, potentially delivering benefits to millions of smartphone users, as well as informing the design and effective deployment of similar privacy labels on other platforms (e.g., Android) and in other domains (e.g., websites, Internet of Things).

Acknowledgements

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, and a gift from Innovators Network Foundation. 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] I. Adjerid, A. Acquisti, L. Brandimarte, and G. Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, SOUPS '13, New York, NY, USA, 2013. Association for Computing Machinery.

[2] A. I. Antón, J. B. Earp, Q. He, W. Stufflebeam, D. Bolchini, and C. Jensen. Financial privacy policies and the need for standardization. *IEEE Security & privacy*, 2(2):36–45, 2004.

[3] Apple. App privacy details on the app store. https://developer.apple.com/app-store/app-privacy-details/, 2021.

[4] Apple. Privacy definitions and examples. https://apps.apple.com/story/id1539235847, 2021.

[5] Apple. About privacy and location services in iOS and iPadOS. https://support.apple.com/en-gb/HT203033, February 3, 2022.

[6] 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.

[7] R. Balebako, R. Shay, and L. F. Cranor. Is your inseam a biometric? a case study on the role of usability studies in developing public policy. In *Workshop on Usable Security*, volume 23, 2014.

[8] V. Bannihatti Kumar, R. Iyengar, N. Nisal, Y. Feng, H. Habib, P. Story, S. Cherivirala, M. Hagan, L. Cranor, S. Wilson, F. Schaub, and N. Sadeh. Finding a choice in a haystack: Automatic extraction of opt-out statements from privacy policy text. In *Proceedings of The Web Conference 2020*, WWW '20, page 1943–1954, New York, NY, USA, 2020. Association for Computing Machinery.

[9] J. L. Boyles, A. Smith, and M. Madden. Privacy and data management on mobile devices. *Pew Internet & American Life Project*, 4:1–19, 2012.

[10] V. Braun and V. Clarke. Thematic analysis. *APA handbook of research methods in psychology, Vol 2: Research designs: Quantitative, qualitative, neuropsychological, and biological.*, pages 57–71, 2012.

[11] I. C. Campbell. Apple will require apps to add privacy 'nutrition labels' starting december 8th. *The Verge*, November 5, 2020.

[12] F. H. Cate. The limits of notice and choice. *IEEE Security & Privacy*, 8(2):59–62, 2010.

[13] B. X. Chen. What we learned from apple's new privacy labels. *The New York Times*, Jan January 27, 2021.

[14] 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 Workshop on Privacy in the Electronic Society*, pages 73–102, 2021.

[15] H. Choi, J. Park, and Y. Jung. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior*, 81:42–51, 2018.

[16] J. Colnago, Y. Feng, T. Palanivel, S. Pearman, M. Ung, A. Acquisti, L. F. Cranor, and N. Sadeh. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2020.

[17] Congress of the United States. House of Representatives, Committee on Energy and Commerce, letter to Tim Cook. https://energycommerce.house.gov/newsroom/press-releases/ec-chairs-question-accuracy-of-apple-s-new-app-privacy-labels, Feb 9, 2021.

[18] L. F. Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.*, 10:273, 2012.

[19] L. F. Cranor, P. Guduru, and M. Arjula. User interfaces for privacy agents. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 13(2):135–178, 2006.

[20] A. Das, M. Degeling, D. Smullen, and N. Sadeh. Personalized privacy assistants for the internet of things: providing users with notice and choice. *IEEE Pervasive Computing*, 17(3):35–46, 2018.

[21] Z. Doffman. Facebook Tracks Your iPhone Location—This Is How To Stop It. *Forbes*, May May 22, 2021.

[22] N. Ebert, K. Alexander Ackermann, and B. Scheppler. Bolder is better: Raising user awareness through salient and concise privacy notices. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2021.

[23] S. Egelman, J. Tsai, L. F. Cranor, and A. Acquisti. Timing is everything? the effects of timing and placement of online privacy indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 319–328, 2009.

[24] P. Emami-Naeini. *Informing Privacy and Security Decision Making in an IoT World*. PhD thesis, Carnegie Mellon University, 2020.

[25] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi. Ask the experts: What should be on an IoT privacy and security label? In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 447–464. IEEE, 2020.

[26] P. Emami-Naeini, H. Dixon, Y. Agarwal, and L. F. Cranor. Exploring how privacy and security factor into IoT device purchase behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19,

page 1–12, New York, NY, USA, 2019. Association for Computing Machinery.

[27] B. Fabian, T. Ermakova, and T. Lentz. Large-scale readability analysis of privacy policies. In *Proceedings of the International Conference on Web Intelligence*, WI '17, page 18–25, New York, NY, USA, 2017. Association for Computing Machinery.

[28] Federal Trade Commission. Mobile privacy disclosures: Building trust through transparency (FTC staff report). https://www.ftc.gov/reports/mobile-privacy-disclosures-building-trust-through-transparency-federal-trade-commission, February 2013.

[29] Federal Trade Commission. Internet of things: Privacy & security in a connected world (FTC staff report). https://www.ftc.gov/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things, January 2015.

[30] Y. Feng, Y. Yao, and N. Sadeh. A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021.

[31] G. A. Fowler. I checked apple's new privacy 'nutrition labels.' many were false. *The Washington Post*, January, 29 2021.

[32] 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.

[33] S. Gibbs. Google has been tracking Android users even with location services turned off. *The Guardian*, Nov 22 2017.

[34] S. E. Gindin. Nobody reads your privacy policy or online contract: Lessons learned and questions raised by the FTC's action against sears. *Nw. J. Tech. & Intell. Prop.*, 8:1, 2009.

[35] J. Gluck, F. Schaub, A. Friedman, H. Habib, N. Sadeh, L. F. Cranor, and Y. Agarwal. How short is too short? implications of length and framing on the effectiveness of privacy notices. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 321–340, 2016.

[36] H. Habib, S. Pearman, J. Wang, Y. Zou, A. Acquisti, L. F. Cranor, N. Sadeh, and F. Schaub. "It's a scavenger hunt": Usability of websites' opt-out and data deletion choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020.

[37] H. Habib, Y. Zou, Y. Yao, A. Acquisti, L. Cranor, J. Reidenberg, N. Sadeh, and F. Schaub. Toggles, dollar signs, and triangles: How to (in)effectively convey privacy choices with icons and link texts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery.

[38] C. B. Jackson and Y. Wang. Addressing the privacy paradox through personalized privacy notifications. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, 2(2):1–25, 2018.

[39] 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*, pages 1–12, 2009.

[40] 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, page 1573–1582, New York, NY, USA, 2010. Association for Computing Machinery.

[41] P. G. Kelley, L. F. Cranor, and N. Sadeh. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 3393–3402, 2013.

[42] 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.*, 5(3), sep 2021.

[43] 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*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery.

[44] 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*, CHI EA '22, New York, NY, USA, 2022. Association for Computing Machinery.

[45] B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. Zhang, N. Sadeh, A. Acquisti, and Y. Agarwal. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security*, SOUPS '16, page 27–41, USA, 2016. USENIX Association.

[46] A. M. McDonald and L. F. Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4:543, 2008.

[47] A. M. McDonald, R. W. Reeder, P. G. Kelley, and L. F. Cranor. A comparative study of online privacy policies and formats. In I. Goldberg and M. J. Atallah, editors, *Privacy Enhancing Technologies*, pages 37–55, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.

[48] 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.*, 3(CSCW), nov 2019.

[49] Y. Meier, J. Schäwel, and N. C. Krämer. The shorter the better? effects of privacy policy length on online privacy decision-making. *Media and Communication*, 8(2):291–301, 2020.

[50] G. R. Milne, M. J. Culnan, and H. Greene. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing*, 25(2):238–249, 2006.

[51] Mozilla Wiki. Privacy icons. https://wiki.mozilla.org/Privacy_Icons, June 2011.

[52] A. Oeldorf-Hirsch and J. A. Obar. Overwhelming, important, irrelevant: Terms of service and privacy policy reading among older adults. In *Proceedings of the 10th International Conference on Social Media and Society*, SMSociety '19, page 166–173, New York, NY, USA, 2019. Association for Computing Machinery.

[53] J. Porter. Google play store's app privacy labels start appearing. *The Verge*, April 26, 2022.

[54] A. Rao, F. Schaub, N. Sadeh, A. Acquisti, and R. Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Pri-*

vacy and Security (SOUPS 2016), pages 77–96, Denver, CO, June 2016. USENIX Association.

[55] A. Ravichander, A. W. Black, T. Norton, S. Wilson, and N. Sadeh. Breaking down walls of text: How can nlp benefit consumer privacy? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4125–4140, 2021.

[56] 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)*, pages 4947–4958. Association for Computational Linguistics, Nov. 2019.

[57] 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*, 45(S2):S163–S190, 2016.

[58] J. R. Reidenberg, T. Breaux, L. F. Cranor, B. French, A. Grannis, J. T. Graves, F. Liu, A. McDonald, T. B. Norton, and R. Ramanath. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ*, 30:39, 2015.

[59] J. R. Reidenberg, N. C. Russell, V. Herta, W. Sierra-Rocafort, and T. B. Norton. Trustworthy privacy indicators: Grades, labels, certifications, and dashboards. *Washington University Law Review*, 96:1409, 2018.

[60] D. Reinhardt, J. Borchard, and J. Hurtienne. Visual interactive privacy policy: The better choice? In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery.

[61] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17, Ottawa, July 2015. USENIX Association.

[62] F. Schaub, B. Könings, and M. Weber. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *IEEE Pervasive Computing*, 14(1):34–43, 2015.

[63] Secretary's Advisory Committee on Automated Personal Data Systems. *Records, Computers, and the Rights of Citizens: Report*. US Department of Health, Education & Welfare, 1973.

[64] 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)*, 3(1):55–78, 2011.

[65] R. I. Singh, M. Sumeeth, and J. Miller. A user-centric evaluation of the readability of privacy policies in popular web sites. *Information Systems Frontiers*, 13(4):501–514, 2011.

[66] R. H. Sloan and R. Warner. Beyond notice and choice: Privacy, norms, and consent. *The Journal of High Technology Law*, 14:370, 2014.

[67] D. Smullen, Y. Feng, S. A. Zhang, and N. Sadeh. The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences. *Proceedings on Privacy Enhancing Technologies*, 2020(1):195–215, 2020.

[68] C. Spensky, J. Stewart, A. Yerukhimovich, R. Shay, A. Trachtenberg, R. Housley, and R. K. Cunningham. SoK: Privacy on mobile devices–it's complicated. *Proceedings on Privacy Enhancing Technologies*, 2016(3):96–116, 2016.

[69] N. Steinfeld. "I agree to the terms and conditions": (How) do users read privacy policies online? An eye-tracking experiment. *Computers in human behavior*, 55:992–1000, 2016.

[70] W. B. Tesfay, P. Hofmann, T. Nakamura, S. Kiyomoto, and J. Serna. Privacyguide: towards an implementation of the EU GDPR on internet privacy policy evaluation. In *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics*, pages 15–21, 2018.

[71] M. D. C. Tongco. Purposive sampling as a tool for informant selection. *Ethnobotany Research and Applications*, 5:147–158, 2007.

[72] M. W. Vail, J. B. Earp, and A. I. Antón. An empirical study of consumer perceptions and comprehension of web site privacy policies. *IEEE Transactions on Engineering Management*, 55(3):442–454, 2008.

[73] T. Vila, R. Greenstadt, and D. Molnar. Why we can't be bothered to read privacy policies. In *Economics of information security*, pages 143–153. Springer, 2004.

[74] K.-P. L. Vu, V. Chambers, F. P. Garcia, B. Creekmur, J. Sulaitis, D. Nelson, R. Pierce, and R. W. Proctor. How users read and comprehend privacy policies. In *Symposium on Human Interface and the Management of Information*, pages 802–811. Springer, 2007.

[75] P. Wijesekera, J. Reardon, I. Reyes, L. Tsai, J.-W. Chen, N. Good, D. Wagner, K. Beznosov, and S. Egelman. Contextualizing privacy decisions for better prediction (and protection). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–13, New York, NY, USA, 2018. Association for Computing Machinery.

[76] S. Zimmeck, P. Story, D. Smullen, A. Ravichander, Z. Wang, J. Reidenberg, N. C. Russell, and N. Sadeh. Maps: Scaling privacy compliance analysis to a million apps. *Proceedings on Privacy Enhancing Technologies*, 2019(3):66–86, 2019.

Appendix

Screening Questionnaire

- You can check your iOS version by going to Settings>General>About, and looking at "Software Version". Please type in the version number exactly as it appears in the "Software Version" section. [free-text]
- I am at least 18 years old, reside in the US, and am a regular user of an iPhone with iOS 14 or above.
- Please select the most applicable answer. I downloaded one or more apps from the app store ___.
  – In the past week – In the past month – In the past 3 months – More than 3 months ago [disqualified]
- Roughly how many new apps have you downloaded from the app store yourself over the past 3 months?
  – None [disqualified] – Somewhere between 1 and 10 – Likely more than 10
- Have you ever done any of the following in the past? Please select all that apply.
  - Uninstalled or stopped using an app or service because of the types of data the app collects about you or how that data is used
  - Reviewed an app's privacy settings, namely what data it requests access to
  - Read (partially or fully) an app's privacy policy or end-user license agreement
  - Used a VPN or Tor for non-work-related reasons on your phone or other device
  - Decided not to download an app after looking at its privacy information in the app store
  - None of the above
- What is your age? [free-text]
- What is your gender? – Male – Female – Non-binary – Prefer to self disclose
- What is your occupation? [free-text]
- Please enter your email. Your email will only be used to contact you to set up a time for the study and to pay you, if you are selected to participate in this study. If you are not selected, your responses to the survey (including your email address) will be deleted within 3 days of the completion of recruitment for the study. Your email address will not be shared with anyone and will be stored separately from your other study data.

Interview Scripts

- Introduction: Thank you for meeting with me today. This interview is being conducted for research at Carnegie Mellon University to better understand how people interact with mobile apps in the Apple App Store. We will ask you to answer some questions and view some information in the App Store. This session should take no more than 1 hour to complete, and will be recorded via Zoom. Upon completion of the study, you will receive $25 in the form of an Amazon Gift Card that will be sent to you via email. You will be asked to share your iPhone's screen via Zoom at some point during the session to enable us to follow what you are doing on your phone as you visit the App Store.
- Please answer our questions truthfully and as thoroughly as possible. If in doubt, feel free to ask me for clarification at any point during the interview. I want to emphasize that there are no right or wrong answers. Our goal is simply to understand your opinions and thought processes. You may stop the interview at any point, or choose to not answer a question, or take a break if you wish. Please do not reveal any private or personally-identifiable information about yourself or others during the interview. If you accidentally reveal any personal information, please let me know so that I can remove it from the recording. Do you have any questions at this time?
- Part 1: General Questions about App Usage
  - For how long have you been using an iPhone? (Prompt: Any particular reason why you chose an iPhone?)
  - To the best of your knowledge, approximately how many apps do you have on your iPhone? (estimates are expected) [After getting the estimate, give them instructions to look up the actual number Settings>General>About>Applications]
  - When was the last time that you downloaded a new app on your phone?
  - Could you describe a recent experience when you decided to download an app on your phone, starting from how you discovered the new app all the way to what happened when you used it for the first time? (to the extent they went all the way - some people can stop halfway and decide not to set up an account or may even change their mind and remove the app)
  - What are some of the typical factors that influence which apps you download on your phone? (Prompt: app reviews, brand, ratings, security, ranking, data

privacy, your friends) Have you ever compared different apps before deciding which one to download (Prompt: what types of things have you compared?)? Was data privacy ever a reason that you chose or did not choose an app?

- Part 2 : Information Seeking
  - Have you ever wondered what information apps collect about you?
  - [If they say they have] How would you go about finding out what information an app collects about you and what the app does with the information?
  - [follow-up] Have you ever actually done that?
  - [If they say they have not] If you were to look for this, how could you possibly find out what information an app collects about you and what the app does with that information? Prompt: media? friends/family? experts? privacy policies/EULA? permission settings? Other? Do you think those sources are reliable? App store? Have you ever looked
- Part 3: Label Comprehension — 2 Scenarios
  - [Instructions] This section requires you to share your screen with us via zoom. Which device are you on, your iPhone or your desktop/laptop? Please open the App Store app on your phone before you start sharing your screen. Please silence your notifications and remove anything confidential from your screen. [If participant is not on iOS 15, show instructions] Please enable Do Not Disturb on your device to prevent unexpected notifications by going to "Settings">"Focus" and then "Do Not Disturb" and turn on the top toggle. [Remind if on their phone using the following sentence] Please note that we will be able see snapshots of your apps when you switch from Zoom to the App Store app. So please make sure that there is nothing sensitive displayed on your screen before you start sharing your screen.
  - [Show instructions on screen if needed] Great! Now that you are sharing your screen.
  - Could you search for the [Doordash or Chipotle] app? Are you familiar with this app? Could you describe what this app does? Have you used [APP] before? [If yes] how often do you use it?
  - Please scroll down to the "App Privacy" section. Do you remember ever seeing or reading an "App Privacy" section like this one before? (follow-up: if yes, ask about their experience with privacy labels; when did you last see one? For which app? What was the context? Did you find the information useful? Did you end up downloading the app? If not, why not? Do you typically look for this information before downloading a new app?)
  - Please take some time to read this "App Privacy" section. What do you think of the section you are seeing? What do you think this section is for? (impression testing)
  - Please click the "See Details" at the top right corner and take some time to read this as well. Let's go through the app privacy information section you just looked at systematically. Please answer the following questions based on what you see in this "App Privacy" section:
  - Starting at the top where it says "Data used to Track you", what do you think "data used to track you" means?
  - Do you know what (online) tracking is?
  - What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
  - In this section there is information about "identifiers." Do you know what an identifier is? If yes, what are (other) examples of identifiers?
  - What do you think of the fact that this app may use your identifiers to track you across apps and websites owned by other companies? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
  - In this section there is information about "product interaction" under "usage data." Do you know what product interaction entails? How do you think Doordash tracks you via product interaction(s)?
  - [If the user is looking at Doordash]
  - What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
  - [Tracking] Do you think Doordash is allowed to share your location with a third-party company that would combine your location obtained from Doordash and location data from other apps and websites to build a history of your whereabouts?
  - What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
  - [Tracking] Do you think Doordash uses data collected from other companies (including websites, apps, and offline services) to decide what ads to show you? If so, what data do you think the app uses? How can you tell? What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.

- If you scroll down a bit you should see a heading for "data linked to you," previously there was another heading for "data used to track you". What is the difference between "data used to track you" and "data linked to you"?
- The next heading is Third-party advertising. If you scroll down a bit you should see another heading Developer's Advertising or Marketing. What do you think those headings denote? Do you think that these headings refer to different practices and, if so, what are the differences?
- What do you think "third-party advertising" means?
- Do you know what targeted advertising is?
- Do you find it to be a useful practice or are you possibly concerned about it? Please explain. (There is no right or wrong answer. We're just curious to understand how you feel about these practices)
- Do you know what Developer's Advertising or Marketing means? What do you think is the difference, if there is, between "third-party advertising" and "developer's advertising or marketing"?
- Below that you will see Browsing History. What do you think "browsing history" covers?
- Do you think browsing history includes content the user has viewed that is not part of the app, such as websites?
- Are you concerned or not concerned about this data being collected?
- If you keep scrolling you will see "other data", what do you think "other data" include?
- If you scroll down you will see the Analytics heading. What do you think your data being used for "Analytics" purposes means?
- What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- [Analytics] Do you think Doordash uses data to understand or analyze your behavior (e.g., to develop new features, to measure audience characteristics)? If so, what data do you think the app uses? What do you think of it?/How do you feel about it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- If you scroll down you will see the Product Personalization heading. Do you know what that is? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- If you keep scrolling down you will see the heading "app functionality," What do you think your data being used for "App Functionality" purposes means? What do you think of all the other purposes beyond app functionality?
- What do you think of it? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- If you keep scrolling you will see "other purposes," what do you think about your contacts being used for "other purposes"? Do you find it to be a useful practice or are you possibly concerned about it? Please explain.
- Is there any information here that you do not understand?
- Is there any information that is not present but you would like to know about?
- [For participants assigned to the Chipotle]
- Let's go through the app privacy information section you just looked at systematically. Please answer the following questions based on what you see in this "App Privacy" section: Starting at the top where it says "Data not linked to you", what do you think "data not linked to you" means? The next heading is Analytics. If you scroll down a bit you should see another heading App Functionality. What do you think those headings denote?
- [Analytics] Do you think Chipotle uses data to understand or analyze your behavior (e.g., to develop new features, to measure audience characteristics)? If so, what data do you think the app uses?
- How do you think Chipotle collects "contact info" such as "Name" and "Email address" without the data being linked to you?
- Do you think Chipotle collects any data that is linked to you? For example, when you place an order?
- Do you think this App Privacy section includes all the data and all the usage of your data that Chipotle collects about you? Please explain.
- Is there any information here that you do not understand?
- Is there any information that is not present but you would like to know about?

- Perceptions of Privacy Labels
  - In general, how do you feel about the information provided here (i.e., the information under the "App Privacy" section)?
  - Do you find this type of information useful? Why/Why not? Do you find it easy to understand? Do you find it well organized?
  - Do you feel you understand most of the information provided in that section?

  - Did you learn about things you didn't know or you pretty much knew everything in that section? Or somewhere in-between?
  - What do you like or dislike about that section? How can this section be improved?
  - Do you think you might pay attention to this information later or you will probably not be looking at this?
  - [If yes] Would this information influence your decision to download the app one way or another?
  - [If no] Why do you think you will not be looking at this?
  - [if participant saw labels before] After downloading an app, did you ever find yourself going back to this section in the app store?
  - Under what circumstances? For what reason? (Prompt: to answer privacy questions you might have had)
  - If you had to guess, who do you think provided the information under "App Privacy"?
  - [Follow-up] Do you believe that this information is provided directly by Apple or vetted by Apple? Or by the app developers? Why/why not?
  - If it's the latter, do you believe that it has been reviewed by Apple?
  - [Follow-up] Do you believe it is done manually or the result of some automated processing?
  - [if participant saw labels before] If you ever looked at this type of information in the past, did you ever have that question in mind? Who did you assume provided the information in the labels? Did you think the information could be trusted?
  - Do you think the information provided in this section is reliable? How likely are you to trust this information? Why/why not?

  [Remind participants to stop sharing their screen]
- Part 5: General Privacy Concerns / Behaviors
  - Have you ever read the privacy policy (partially or fully) of a mobile app? What do you think of them? Prompt: When you read a privacy policy, what do you typically do? How much do you typically understand the privacy policies you read?
  - Have you ever regretted downloading or using an app because of data privacy issues? [Follow-up]: Did you take any further actions because of these regrets such as changing your privacy settings, uninstalling the app or limiting your use of it?
  - In the past, have you ever had your personal data misused or compromised in general? If so, what happened? What about data related to apps or web services?
- Wrap-up Alright, I have asked all the interview questions. Is there anything else you would like to say? Any questions at this point? Or any comments? Now I will be asking you to fill out a short survey. The survey link is pasted inside the chat. You will receive the e-gift card via email soon after you complete the survey. Feel free to disconnect now. Thank you so much for your participation!

Post-Interview Questions

- How much control do you think you have over the data that companies collect about you?
- How concerned are you, if at all, about how companies are using the data they collect about you?
- How much do you feel you understand what companies are doing with the data they collected about you? Prompt: a great deal, some, very little, nothing
- What is the highest level of education you have completed?
- Have you ever held a job or received a degree in computer science or any related technology field?
- Which of the following best describes your employment status?
- Privacy behavior: Have you used the following tools in the past year? Please select all that apply.

Codebook

Code categories are shown in bold type, with the list of codes in that category following. For a more detailed version with code descriptions, see https://osf.io/47kzt/.

App Privacy

- **iPhone usage length:** less than 5 years, at least 5 and less than 10 years, at least 10 years
- **Why use iPhone**
- **iPhone # of apps estimate:** $\leq$40, 50–100, >300
- **iPhone actual # of apps:** <50, 50–100, 101–200, >200
- **Recent app download time:** within 1 day, within a week, within a month
- **App download process:** search keywords in App Store, search in Google, download multiple apps and then delete unwanted, learn specific app from ads, learn specific app from recommendations, generally know the app the download when in the App Store
- **Factors considered when downloading apps:** cost, reviews, ratings, utility, descriptions,

brand/trust, rewards program, space/battery, number of downloads or reviews, bugs
- **Concerns about app privacy:** yes, no, consider privacy before downloading, remove app out of privacy concern, privacy concern for newly downloaded apps, privacy concern about Facebook related apps, privacy concern only for important apps
- **Whether have questions about app data collection** yes, no, specific question
- **Where to learn about app data collection** Google, terms of services/privacy policies, iPhone privacy settings, iPhone privacy prompts, in-app privacy settings, the App Store, media, downloaded data
- **Looking for app privacy in the App Store:** no
- **Seen app privacy section in the App Store before:** no, yes, yes but only aware of its existence
- **DoorDash app used before:** yes, no
- **Chipotle app used before:** yes, no

Label Understanding & Perception

- **Tracking:** understanding, confusion, concerned, not concerned, useful, not useful, mixed feelings
- **Tracking by identifiers:** concerned, not concerned, mixed feelings
- **Tracking implies aggregating location data:** yes, clear from the label, not clear from the label, concerned
- **Tracking used for advertising:** yes, clear from the label, not clear from the label
- **Data linked to you:** understanding, confusion, useful
- **Difference between data used to track you and data linked to you:** understanding, confusion
- **Data not linked to you:** understanding, confusion
- **Contact under data not linked to you:** confusion
- **Targeted advertising:** understanding, concerned, not concerned, mixed feelings, useful
- **Third party or developer advertising:** understanding, confusion
- **App functionality:** understanding, confusion, useful
- **Analytics:** understanding, confusion, useful, not concerned
- **Product Personalization:** understanding, confusion, useful, concerned
- **Identifiers:** understanding, confusion, concerned
- **Device id:** understanding, confusion
- **Product interaction:** understanding, confusion, mixed feelings
- **Browsing or search history:** understanding, confusion, concerned, not concerned
- **Other category:** understanding, confusion, concerned, not concerned
- **User content:** concern, confusion
- **Contacts used by DoorDash for other purposes:** concerned, mixed
- **Jargon confusion:** usage data, crash data, diagnostics, coarse location, purchase
- **App has no need for listed data on label**
- **Label first impression**
- **Confusion about label structure**
- **Confusion about label sections**
- **Label useful:** yes, no
- **Understood most of the labels?:** yes, no, in-between
- **Learned new things from labels?:** yes, no, in-between
- **Future use of labels:** yes, no, depends
- **Labels impacting later decision to download apps:** yes, no, depends
- **Labels include all app data collection practices:** yes, no, not sure
- **What participants like about the labels:** existence, increased transparency, other
- **What participants dislike about the labels:** vagueness, long and/or repetitive, use of jargon
- **How to improve the labels or what participants would want the labels to include:** add accessible definitions, add specific and contextualized examples of data collected, add privacy controls, add whether the data is being shared or sold and/or with whom, add data retention, explain in details what they do with the data and/or justification, add how data privacy is protected, arrange purposes in a table format, add contact info for further questions, other
- **Do participants trust the labels?** yes, no, depends, reason for yes, reason for no
- **Labels provided by:** app developers, Apple, both Apple and app developers, not sure
- **Labels reviewed or verified by Apple?** neither, only reviewed, verified, not sure
- **If labels are reviewed, how?** automated processing, manual review, both
- **Like Compact label**
- **Participants expect labels to be interactive**
- **Participants think labels are required**
- **Labels lack oversight or guarantee**

Privacy Attitudes and Experiences

- **Resignation**
- **Trade-off**

- **Not concerned about privacy:** generally unconcerned, nothing to hide, low perceived risk, privacy as a secondary task
- **Privacy concern:** generic, desire to remain personal autonomy, feeling watched, risks
- **Usable privacy:** user burden high, frustration with privacy policies
- **Privacy protection behavior**
- **More concerned after reading labels**
- **Past experience with personal data being misused or compromised?** no, data breaches, fraudulent activity on bank accounts, identity theft, accounts hacked
- **Aware of turning off tracking on iPhone:** yes, no, confusion