

Transparency or Information Overload? Evaluating Users’ Comprehension and Perceptions of the iOS App Privacy Report

Xiaoyuan Wu*
Carnegie Mellon University
wxyowen@cmu.edu

Lydia Hu*
Carnegie Mellon University
ljhu@andrew.cmu.edu

Eric Zeng
Carnegie Mellon University
ericzeng@cmu.edu

Hana Habib
Carnegie Mellon University
htq@cs.cmu.edu

Lujo Bauer
Carnegie Mellon University
lbauer@cmu.edu

Abstract—Apple’s App Privacy Report (“privacy report”), released in 2021, aims to inform iOS users about apps’ access to their data and sensors (e.g., contacts, camera) and, unlike other privacy dashboards, what domains are contacted by apps and websites. To evaluate the effectiveness of the privacy report, we conducted semi-structured interviews ($n = 20$) to examine users’ reactions to the information, their understanding of relevant privacy implications, and how they might change their behavior to address privacy concerns. Participants easily understood which apps accessed data and sensors at certain times on their phones, and knew how to remove an app’s permissions in case of unexpected access. In contrast, participants had difficulty understanding apps’ and websites’ network activities. They were confused about how and why network activities occurred, overwhelmed by the number of domains their apps contacted, and uncertain about what remedial actions they could take against potential privacy threats. While the privacy report and similar tools can increase transparency by presenting users with details about how their data is handled, we recommend providing more interpretation or aggregation of technical details, such as the purpose of contacting domains, to help users make informed decisions.

I. INTRODUCTION

Digital devices and services collect user data for purposes such as providing functionality, conducting analytics, and targeting advertisements. More transparency about how user data is dealt with can be valuable for those who want to prevent unwanted data sharing or usage [5], [30], [54]. However, explaining how user data is handled can involve both high volumes and complexity of information, making “transparency” difficult to define and achieve [41], [43], [47]. Traditional privacy policies provide disclosure but are often difficult to understand [24], [43], [52], [53]. Newer tools such as privacy nutrition labels and dashboards attempt to improve

the comprehensibility of how data is collected, stored, and shared by online services, apps, and devices [33], [54], [57].

In 2021, Apple released a new iOS feature called the App Privacy Report (“privacy report”) to provide users with “a more complete picture of how the apps [they] use treat [their] data [2].” The privacy report aggregates two categories of information. The first is related to the phone’s data and sensors (e.g., location, camera, contacts), where the privacy report lists which apps are using these data and sensors and at what times. The second is related to network activity, where the privacy report lists the domains that apps or websites contact and the frequency of that contact. The content of the privacy report can help identify commonly studied privacy risks (e.g., apps that have been granted unnecessary permissions, third-party tracking, and advertising).

Users’ reactions to privacy transparency tools have been studied previously (e.g., Google My Activity Dashboard, iOS privacy labels), finding that users had positive responses to how the tools informed them about data handling processes on their digital products and devices [4], [5], [15], [56]. However, Apple’s privacy report differs from privacy dashboards that have been deployed on other popular platforms, most significantly in that it attempts to present network activity information in substantially more detail, i.e., with individual domain names. Because these aspects of the privacy report are new, there is little understanding of whether and to what extent they help users.

We investigate whether the privacy report achieves its stated goals of improving transparency and user understanding of app behavior, as well as how users react to what they learn through the privacy report. Further, we investigate whether interacting with the privacy report causes any intention in users to change how they interact with their phones or whether it motivates any privacy-protective behaviors.

More specifically, we conducted semi-structured interviews with 20 participants recruited from Prolific to examine their interactions with and perceptions of the privacy report. Through this, we aimed to answer the following research questions:

RQ1. (*Understanding*) Does the privacy report help users

*The first two authors contributed equally to the paper.

identify and understand overpermissioning, cross-app tracking, or third-party data collection occurring on their phones?

RQ2. (*Attitudes*) What are users' attitudes toward the information they learn from the privacy report?

RQ3. (*Intent to change behavior*) Does interacting with the privacy report influence users to change their app and/or privacy behaviors? Why?

We find that participants easily understood information about which apps used their phone sensors and were mostly unsurprised by the types and times of sensor access. When participants learned of unwanted access by interacting with the privacy report, they were able to map those concerns to settings where they could rescind app permissions. However, sometimes participants expected that they could control permissions in ways that were not actually available, e.g., participants expected that they could turn off specific sensor access for some built-in apps, even though iOS does not always allow this.

The sections of the privacy report that described apps' and websites' access to domains were more difficult for participants to understand. On one hand, all participants identified cases of third-party data collection and cross-app tracking by recognizing that multiple apps or websites they used had interacted in some way with the same domain. On the other hand, almost all participants expressed confusion or explicit misunderstanding in interpreting some aspect of the network activity. Participants were confused about, for instance, what it meant for an app to contact a domain, why there were so many domains contacted, and what the broader implications were. When participants were concerned about their apps and websites contacting domains, some identified potential ways they believed would address the perceived risks (e.g., deleting apps or blocking apps from sharing data with specific worrisome domains), but these options were mostly infeasible (e.g., due to unwillingness to compromise app functionality or lack of ways to block specific domains). A few participants could not identify any actions to take at all.

Overall, the privacy report effectively conveys information about how apps access phone data and sensors, in that participants more consistently felt informed by this section and identified how they could act in the cases when they found unwanted sensor access. In contrast, the sections related to network activity attempted to convey more technical information but failed to introduce, contextualize, or interpret the reasons behind or implications of apps and websites contacting third-party domains. As such, participants often did not know what to take away or how to act on the information in a way that left them feeling better off. These findings help explain why the current privacy report does not help users protect their privacy, and why similar privacy dashboards might also struggle. Based on our findings, we make several recommendations for how this privacy report, and similar privacy transparency features, can improve. For instance, we suggest including more explanations for both why phone sensors are accessed and why domains are contacted, which may help users better judge which types of activity and data sharing they should be concerned about. Some of our suggestions may be challenging to implement, as they could require changes that go beyond the user interface. We discuss both the challenges

and the importance of implementing these additional features if privacy dashboards are to more successfully help users understand and act on privacy information.

II. BACKGROUND

The privacy report became available to users in iOS version 15.2. The feature is turned off by default but can be turned on at any time in Settings. When activated, the privacy report keeps information from the past seven days. There are four sections. The first ("Data and Sensor Access") shows which data and sensors different apps are accessing and at what times. The next two sections compile the domains that apps ("App Network Activity") and websites ("Website Network Activity") contact. The last section presents the same information about domains as the previous two but is organized differently. Each section is described in more detail below.

a) Data and Sensor Access: This section, pictured in Figure 1a, displays iPhone apps that have accessed data or sensors in the past seven days, starting from the most recent access. "Data" specifically refers to things like the user's photos, media library, and contacts, while "sensor" refers to microphone, camera, and location. Each app is labeled with the data or sensor(s) that it accessed and the elapsed time since the last access. Users can select any app to see a full list of data and sensors it has accessed, then select any of the data or sensors to view timestamps or timeframes the access occurred.

b) App Network Activity: This section, pictured in Figure 1b, lists apps in order from most to least active in how many total domains the app contacts. Each app has a bar underneath labeled with the number of unique domains the app contacted.

An app can be selected to view the list of unique domains that the app contacted in the past seven days. Each domain is labeled with a bar and number showing the number of times the app contacted that domain in the past week. A domain can also be selected for a "domain summary page" (Figure 1c), with all the apps and websites that have contacted this domain, with a time and date label for the most recent access.

c) Website Network Activity: This section is organized identically to the previous section, but instead of reporting activity from apps, it reports on websites accessed from within apps and which domains those websites contacted. Selecting any of the domains leads to the same domain summary page in Figure 1c. When an app or website contacts a domain, we will refer to it as "network activity."

d) Most Contacted Domains: This section, pictured in Figure 1d, aggregates all the domains that are contacted by any apps or websites on the phone and organizes them in order of most contacted to least contacted.

III. RELATED WORK

In this section, we examine methods and challenges of communicating privacy practices with users and review how privacy dashboards approach the issue. We review findings on users' perceptions of such tools. As the privacy report studied here is unique to iPhones, we review the challenges of communicating privacy on mobile devices and previous efforts to address them.

a) Communicating privacy information to users: Presenting privacy-related information to users is challenging because users perceive privacy differently and have various levels of technical expertise [55]. Previous works have explored how best to present information about the collection and handling process of personal data. The communication is frequently done through privacy policies [6], [19], [20], [32], [36], which have been found to be challenging for users. Many users do not read or cannot comprehend privacy policies due to their length and complexity [24], [43], [47], [52], [53]. To address this issue, researchers have proposed ways to improve the readability of privacy policies, including breaking privacy policies into layers and sections, creating a software standard for browsers to automatically parse and present privacy policies, and using privacy “nutrition” labels [10], [33], [37]. Though some of the aforementioned methods can improve readability, privacy policies only state the data practices of online platforms without providing context on when or why users’ data are being collected and shared with third parties.

Privacy dashboards have gained traction in recent years as a way to communicate privacy practices while including more context about user data, with prior work identifying key components to make effective dashboards [17], [41], [57]. Examples of privacy dashboards include ones provided by Meta, Google, and Apple, which inform users about what data they collect, how they (intend to) use it, and who they share it with [2], [21], [38]. A unique aspect of such dashboards is that users can review what happened to their data and in some cases exercise a degree of control over how service providers’ handle their data [30]. For instance, Google’s Ad Settings and Meta’s Privacy Center provide control options for users to opt-out of certain data sharing practices [23], [38].

b) User perceptions of privacy dashboards: With the adoption of privacy dashboards by online platforms such as Google, Apple, and Meta, researchers have sought to understand how users perceive these tools. Overall, users found privacy dashboards helpful in how they provided previously inaccessible information of what types of user data (e.g., personally identifiable information, online activities, etc.) are collected by and shared with different parties. For instance, after reviewing Google’s My Activity dashboard, users had increased awareness about how Google collects their activity history and what inferences (e.g., topics of interest) were made from such data [14], [15]. Similar to a privacy dashboard, a browser extension designed by Weinshel et al. presented information about online trackers to users. They found users’ understanding of cross-site online tracking for targeted advertising improved [54] after interacting with the information. Our results from the Data and Sensor Access section are in line with previous studies’ findings in users find information in privacy dashboards to be helpful, in presenting information that was hard to access otherwise, and in improving their understanding of online privacy.

Despite some positive user feedback on privacy dashboards, previous work has also suggested ways to improve. Balash et al., after studying user perceptions of third-party access to Google accounts, suggested privacy dashboards require third parties to specify reasons for access to information from users’ Google accounts [5]. Other studies looked at the level of abstraction needed when informing users of privacy practices

(e.g., data handling) through privacy dashboards. These works disagree on whether raw (e.g., the webpages users visited) or inferred (e.g., potential interests of the user) information best helps users in making privacy decisions [28], [40], possibly due to the different contexts in which the experiments were conducted. We observed participants prefer raw information (e.g., app A accessed sensor B) when they have the background knowledge to understand it, but ask for more interpretation when the information (e.g., lists of domain names) is obscure to them.

Although privacy dashboards have been studied in contexts such as online accounts and browser-oriented settings, there is a lack of understanding of how users interact with privacy dashboards for mobile devices. One unique aspect of mobile devices is that many of their sensors are accessed by third-party apps or websites for data collection, complicating user perceptions of related privacy issues. Previous works have discussed that misuse, overcollection, and leakage of personal information to third parties as well as errors by data users are concerning privacy risks for mobile users [26], [50]. Further, researchers investigated ways to provide users with control over mobile data and sensors [7], [11], [16], [42], [46], [48]. While we understand which privacy risks are of concern to users and some effective ways to provide control options, less is known about how to communicate privacy information with users on mobile devices, especially in a centralized privacy dashboard. Understanding this will improve the design of such features and help increase user awareness of privacy risks when using mobile devices so that they can take actions to protect their privacy. Our study aims to address the gap and understand users’ perceptions of the privacy report, identify what existing information is helpful, and suggest what additional details are useful to add to privacy dashboards.

IV. METHODS

To assess users’ understanding of privacy risks, attitudes, and intent to change behavior (e.g., changing app settings) after interacting with the privacy report, we conducted hour-long semi-structured interviews with 20 participants recruited from Prolific. During interviews, participants first freely explored their phone’s privacy report and then answered questions about their understanding and reactions. As the privacy report is only available on iOS 15.2 or later, we screened participants based on whether they owned an iPhone for personal use, had a compatible iOS version, and were willing to participate in a virtual interview. Lastly, we conducted a follow-up survey to ask about participants’ interactions with the privacy report after the interview. In the following section, we describe in detail our screening survey, interview, and follow-up survey (Section IV-A). We then discuss the recruitment process (Section IV-B), how we analyzed responses (Section IV-C), ethical considerations (Section IV-D) and limitations (Section IV-E).

A. Study Procedure

a) Screening survey: We first conducted a screening survey through Prolific to obtain a pool of eligible participants, based on their iPhone ownership, iOS version, and willingness to participate in the interview. We balanced our sample across genders and levels of privacy awareness. Data collection for the screening survey occurred between 17th October 2023 and

9th December 2023. Below we describe the screening survey, which can be found in full in Appendix A.

The survey starts with questions that determine participants' eligibility for the interview and asks about their willingness to participate through Zoom (questions S1–S3 in Appendix A). Then, we asked about their existing privacy knowledge and practices (questions S4–S10 and S15) and demographic details (questions S11–S14) [25]. Based on participants' responses, we categorized them as either "privacy knowledgeable" or "not privacy knowledgeable" and recruited a diverse sample based on the categorization and demographics data for the interview. Participants on average took six and a half minutes to complete the screening survey and received \$1.25 as compensation.

b) Interview: We conducted semi-structured interviews to learn what privacy risks participants can identify after seeing the privacy report, how they react to the information shown, and what they intend to do to remediate any concerns they identified. The qualitative nature of the study allowed us to understand the nuances of participants' perceptions of the privacy report. Interviews took place between 7th November 2023 and 21st December 2023.

Eligible participants from the screening survey were asked to agree to the informed consent terms of the interview, prompted to install the Zoom application on their iPhones, turn on the privacy report through their settings, and schedule an interview. Interviews took place at least one week after the participant turned on their privacy report, as the privacy report does not start generating information until turned on and keeps data from the past seven days. We asked but did not require participants to export and send us their privacy reports prior to their interviews so that we could verify that their privacy reports had been generated properly. The full interview script can be found in Appendix B.

We began the interview with introductions, a review of the consent form, and a chance for participants to ask questions. We then referred to their screening survey answers to ask questions about their existing privacy preferences and tools, especially on their iPhones (question I1 in Appendix B). Since the privacy report had been available for about two years, we asked what participants' previous knowledge or interactions with the feature were. (questions I2–I6). Next, we instructed participants to share their phone screens through Zoom before starting the main part of the interview.

We started the main part of the interview by asking participants to freely explore the privacy report while thinking aloud. During this process, we noted participants' interactions with the privacy report and the comments they made.

When participants finished their own exploration, we directed them to each of the four sections of the privacy report: Data and Sensor Access, App Network Activity, Website Network Activity, and Most Contacted Domains (described in Section II). For each section, we asked participants if they could find unexpected data and sensor access or domains contacted (RQ1; questions I7–I9, I15–I17, I23–I25). We then asked if seeing the information made them feel differently about their phone or the apps and websites they used (RQ2; questions I10, I18, I26, I30). We checked how well participants understood the information by asking how they would explain

that section to someone else (questions I11, I19, I27, I31). Next, we asked about their intentions, if any, to change how they interact with their apps or phones after interacting with each section (RQ3; questions I13, I21, I29, I33). We concluded the interview by asking participants about the feedback they had for the privacy report, if they intend to use it again, and if any of their existing privacy concerns were addressed by it (questions I34–I38). As the interview was semi-structured, we asked additional follow-up questions based on participants' answers to learn more about their reasoning. The interview took about an hour and participants received \$25.00 upon completion.

c) Follow-up survey: We invited participants to complete a follow-up survey one month after their interview was conducted. In the survey, we asked whether participants continued to use the privacy report (questions Q1–Q3 in Appendix C), what motivated them to use it (question Q4), and how they responded to the privacy report when they used it after the interview (questions Q5 and Q6). We compared these responses to what each participant said during the interview. For participants who said during their interview that they would check the privacy report again and did not, or vice versa, we asked what changed their mind (question Q7). Data collection for the follow-up survey occurred between 15th December 2023 and 23rd January 2024. Participants took an average of 3 minutes and 31 seconds to complete the follow-up survey and received \$1.25.

B. Recruitment and Demographics

We recruited 190 participants over the age of 18 from the US for the screening survey through Prolific. About half (47.9%) of them were female and over half (57.9%) were under the age of 34. More demographics of screening survey participants can be found in Table VI.

Eligible screening survey participants were invited to the interview. Twenty participants joined the interview; 11 were male and nine female. Based on their responses to questions S4 to S8, we categorized six males and five females as "privacy knowledgeable" and the rest as "not privacy knowledgeable" so that we could interview users with different levels of privacy awareness. One month after their interview, we sent participants the follow-up survey, which 14 of the 20 interview participants completed. Half of these were male and half female. Detailed demographics of participants for the interview and follow-up survey can be found in Table I.

C. Data Analysis

We used thematic analysis to qualitatively analyze interview transcripts. Two researchers independently coded the first five interviews using recordings, transcripts, and notes, then met to compare, discuss, and develop a preliminary codebook. Each researcher then independently iterated on the preliminary codebook by coding all 20 interviews using notes and recordings. Both researchers then met to discuss and resolve discrepancies and finalize the codebook before applying it to each interview.

Each relevant participant comment was labeled with which of the four privacy report sections prompted the comment, and the topic (e.g., "Volume of domain activity") discussed. A list

of these topics can be found in Appendix E. Then, a code was applied based on how the comment answered a research question, (e.g., “Confused” corresponds to the second research question about user attitudes). A list of codes can be found in Tables II, III, and V, and example quotes for each code can be found in Table VIII. We reached thematic saturation with our sample size. No new themes that we coded for emerged after interview twelve of 20.

TABLE I: Demographics of interview and follow-up survey participants

	Interview		Follow-Up Survey	
	No.	%	No.	%
Gender				
Female	9	45	7	50
Male	11	55	7	50
Age				
18 - 24	5	25	3	21.4
25 - 34	4	20	3	21.4
35 - 44	4	20	2	14.3
45 - 54	3	15	2	14.3
55 - 64	2	10	2	14.3
65 - 74	2	10	2	14.3
Highest Degree of Education				
High school or equivalent	1	5	1	7.1
Some college, no degree	4	20	4	28.6
Associate degree	1	5	0	0
Bachelor’s degree	11	55	6	42.9
Master’s degree	3	15	3	21.4
Total	20	100	14	100
IUIPC	Avg.	S.D.	Avg.	S.D.
Control	6.00	1.06	5.93	1.18
Awareness	6.70	0.61	6.79	0.42
Collection	6.09	1.03	6.27	0.92

TABLE II: RQ2 interview codes and the amount of participants for whom at least one instance of that code was recorded, in the Data and Sensor Access section (out of 19 participants) and Network Activity sections (out of 20 participants)

RQ2: Attitudes			
Grouping or Description	Code	Data & Sensors	Network Activity
Positive attitudes	Informed	■■■■	■■■■
	Reassured	■■■■	■■■■
	Interested	—	■■■■
Neutral attitudes	Not surprised	■■■■	■■■■
	Not bothered	■■■■	■■■■
	Curious	■■■■	■■■■
Negative attitudes	Concerned	■■■■	■■■■
	Confused	■■■■	■■■■
	Surprised	■■■■	■■■■
	Resigned	—	■■■■
	Strange	■■■■	■■■■

■■■■ = a few; ■■■■ = some; ■■■■ = about half; ■■■■ = most; ■■■■ = almost all

TABLE III: RQ3 interview codes and the amount of participants for whom at least one instance of that code was recorded, in the Data and Sensor Access section (out of 19 participants) and Network Activity sections (out of 20 participants)

RQ3: Intent to Change Behavior			
Grouping or Description	Code	Data & Sensors	Network Activity
App & browsing related	Change app permissions	■■■■	■■■■
	Delete app	■■■■	■■■■
	Find alternative app	■■■■	■■■■
	Change browsing behavior	—	■■■■
	Decrease app/website usage	—	■■■■
	Consider additional privacy tools	—	■■■■
Information seeking	Look for more information elsewhere	■■■■	■■■■
	Check report in future	■■■■	■■■■
	Check phone settings	■■■■	■■■■
Nonspecific intentions	Desire to act but unsure of what can be done	■■■■	■■■■
	No behavior change (no reasons given)	■■■■	■■■■
	General prevention of unwanted activity	—	■■■■
Reasons for not wanting to change behavior	App/website functionality is necessary	■■■■	■■■■
	Not sufficiently concerned	■■■■	■■■■
	Accepting of privacy tradeoff	■■■■	■■■■
	Insufficient information to justify action	■■■■	■■■■
	No effective actions seem to exist	■■■■	■■■■

■■■■ = a few; ■■■■ = some; ■■■■ = about half; ■■■■ = most; ■■■■ = almost all

Due to the qualitative nature of our results, we followed previous work in the field and used the following quantifiers: almost all (>75%), most (55%–75%), about half (45%–54%), some (25%–44%), a few (<25%) when reporting results in Section V [12], [27].

For responses to the follow-up survey, we quantitatively analyzed results for questions Q1 to Q3. One researcher first independently coded responses to questions Q4 to Q7 using the open coding technique and developed a codebook for the follow-up survey. Another researcher then applied codes to each participant’s responses before discussing with the first researcher and resolving discrepancies.

D. Ethical Considerations

This study was approved by our Institutional Review Board. Prior to each stage of the study, participants were asked to consent to the terms for that part of the study. We did not collect personally identifiable information during the study, but linked data collected to Prolific IDs. During the recorded interview, participants shared their phone screens with us, which may have exposed sensitive information. To minimize such exposures, we asked participants to avoid bringing up

sensitive information before starting the recording and guided them to turn on “Do Not Disturb” mode so that phone notifications would not appear.

Sixteen participants sent us their privacy reports. The privacy reports were received through our institutional email and are stored encrypted on servers approved by our institution’s IRB for research use. In accordance with the procedure approved by our IRB, the data will be retained on our institution’s server for a minimum of three years.

Some participants had extreme concerns about what they saw in the privacy report, stemming from misunderstanding network activity. To reduce such concerns, we explained or pointed them to Apple’s documentation outlining what domains could be used for (e.g., analytics, logging into a website). We provided explanations only after recording participants’ initial understanding.

E. Limitations

Recruitment of our study was restricted to Prolific users residing in the U.S. who use an iPhone as their personal mobile device. Those from other parts of the world and those who use mobile devices other than iPhones may have different perceptions than what we observed. Nonetheless, our participants’ privacy concerns were similar to those identified in previous literature about mobile privacy issues [26], [40], [50]. We also note that previous work suggests that Prolific participants’ responses are fairly representative when it comes to users’ perceptions of privacy-related topics [49]. Additionally, as with all surveys and interviews, responses may be affected by social desirability bias, where participants responded in ways they believed to be desirable for the researchers [39]. To mitigate such bias, we informed participants before starting the interview that one of the goals of our study is to explore users’ reactions to the privacy report and we do not have predefined expectations for their responses.

V. RESULTS

Here, we present findings from the semi-structured interview, grouped by the two types of information conveyed in the privacy report. Participants found it easy to understand the Data and Sensor Access section (Section V-A), and what they learned mostly matched their expectations of how and when apps exercised permissions to access data (e.g., contacts, photos) and sensors (e.g., camera, location). We then discuss the more complex sections about network activity (Section V-B), which participants understood less well and had more negative reactions to. Finally, we discuss the follow-up survey (Section V-C) and summarize the results by answering each research question (Section V-D).

A. Data and Sensor Access Section

The Data and Sensor Access section showed information that participants were able to interpret and generally found unsurprising. Individual cases of surprise, concern, and confusion were motivated by specific apps having unexpected permissions to use the data or sensors on their phone (Section V-A1). In such cases, participants consistently related those permissions to a setting they could change, though the actual options did not always match participants’ expectations

of what they had control over (Section V-A2).¹ Our findings here echo those of previous studies about permission controls and transparency mechanisms, for mobile devices and other platforms [1], [8], [16], [31], [44], [45], [48], [51]. For instance, prior work on run-time permissions identified negative user reactions to overpermissioning [31], [45], which we also observed. Despite such findings, operating system vendors (i.e., Google, Apple) have not yet consistently included explanations or control options to address user concerns. Our results lend further support to why changes are necessary; we share more details about how our results compare to those of previous work in Sections V-A1 and V-A2.

1) Identifying overpermissioning: Participants largely understood and were unsurprised by the times and types of permissions shown in the Data and Sensor Access section of the privacy report. The privacy risk of interest was overpermissioning, which almost all (17) participants identified at least one instance of, i.e., if they pointed out a type (if they could not connect an access to a function that they understood the app performed) or time (accesses occurred when they did not recall using the app) of data or sensor access they did not expect.

a) Participants understood which types of data and sensors were accessed, when, and by which apps: All but one participant understood without any guidance that this section of the privacy report lists each app that accessed a type of data or a sensor, and the time of each access. There was still some minor misunderstanding, including a few (3) participants who thought the list of apps was what they used recently, instead of which apps were accessing the data and sensors. In contrast, some (7) participants specifically identified that some apps were performing background data or sensor accesses. The one participant (P28) who had trouble understanding was confused about the privacy report as a whole and thought that the privacy report itself was accessing apps, data, and sensors, instead of displaying when apps were using data and sensors.

Most (14) participants were unsurprised by this section. Almost all (18) participants commented that the accesses made sense to them in light of how they use a certain app, for instance, P18 and P34’s weather app using location for forecasts, or P54’s video conferencing app using the microphone and camera. When looking at the list of times that these sensors were accessed, participants were also able to connect the timestamps to when they remembered using that app.

b) Participants could not connect unexpected accesses to app functionality: Almost all (17) participants identified at least one instance of overpermissioning. The most frequent offender was Apple’s Health app: most (14) found it unexpected that Health periodically accessed their Contacts and participants could not guess why the permission was needed. Confusion was exacerbated by the fact that some rarely or never used the app, which comes preinstalled on iPhones.

From the developer’s perspective, the access to contacts may seem more reasonable. One of the Health app’s features is providing a quick way to reach emergency contacts, which access to contacts makes possible. However, the app does

¹For this section only, the numbers are out of 19 participants and not 20 participants, because P10’s privacy report was missing the Data and Sensor Access part entirely.




Fig. 1: (a) Data & Sensor Access (b) App Network Activity (c) Domain Contacted by Multiple Apps (d) Most Contacted Domains

not ask users to grant this permission and the emergency contacts feature is not immediately obvious, contributing to most (14) participants exhibiting surprise and confusion at the permission. Built-in Apple apps (Health, Podcasts, Photos) generally do not prompt users to grant specific permissions, with a few exceptions. In contrast, participants seemed to be more familiar with the permissions for apps they downloaded, and often specifically recalled granting the permissions they saw those apps exercising.

c) Participants noted unexpected timings of data and sensor accesses: As in previous studies about Android app permissions, we found participants were surprised by apps' background access to the data and sensors on their iPhones [31], [51]. A few (3) participants found apps that were accessing data or sensors when they did not recall using the app. Even if the type of permission was in line with their expectations, the access occurring in the background surprised them. For instance, P54 was surprised by Apple Maps using location when they were not actively using the app for navigation. P34 said that Instagram accessing Photos was expected, but only when they were using the app to share content, as opposed to browsing. Some of these participants were bothered and wondered if background access could occur even if their phone was turned off. Participants who mentioned these unexpected background accesses all wanted to limit those permissions and check the privacy report in the future for these kinds of incursions.

2) Addressing overpermissioning: Participants had concrete ideas about what they could do to address the overpermissioning that bothered them, though the available actions were not necessarily as straightforward.

a) Participants had ideas about how to address overpermissioning: Though overpermissioning surprised participants, not everyone was bothered by it. About half (9) were explicitly concerned to varying degrees. In such cases, participants

consistently mentioned an action they thought could address the specific apps that they felt were using data or sensors unnecessarily. Half (10) of participants, when asked if they would change anything about how they use their phone in response to what they saw in this privacy report section, described changing the permissions to limit unexpected or background accesses. Other participants said overpermissioned apps did not bother them enough to change any settings.

A few (4) participants mentioned deleting an app entirely as a possible way to address overpermissioning. However, almost all (3) of them attributed any desire to delete an app to the fact that they rarely used it: reducing unwanted access to data and sensors was a positive byproduct rather than the primary goal of deletion. The privacy report motivated them to remove apps, but the privacy information conveyed was not the main reason. Rather, the primary reason for removal was the fact that the app was not providing any utility to them that would justify the access to data and sensors.

b) There were mismatches between participants' expectations and available actions: Though participants expected the ability to restrict permissions regardless of the app in question, the control options are not consistent across all apps. Taking again the case of the Health app accessing contacts, to the best of our knowledge, it is not possible to rescind access to contacts from the iOS settings menu or anywhere else on the phone. Apple's documentation says users can "control whether third-party apps have access" to information such as contacts and Photos [3], but does not address the same for built-in apps. Therefore, though half (10) of the participants identified an action they might wish to take to address overpermissioned apps, the available control options may not satisfy their privacy needs.

Another complication is how fine-grained the available control options are, compared to what participants expressed they wanted to do. Users cannot restrict background access for

all data and sensors through iOS settings; granting permission is not specific to how the app is being used at a given moment. Users cannot, for instance, restrict Instagram’s access to photos to only when a user is sharing content through the app, as in P34’s case. The specificity and consistency of available control options did not always align with participants’ desire to restrict permissions. Prior work has recommended providing users with more fine-grained controls, including asking for permissions at run-time and restricting background access to mobile data and sensors [1], [31], [45]. However, our findings suggest that features that might provide rich controls are still inadequate or unavailable, especially for background access to data and sensors.

B. Network Activity Sections

In contrast to other widely available private transparency tools (e.g., Android’s Privacy Dashboard, Meta’s Privacy Center), which present how users’ data is shared with third parties over the network in a less specific way, or not at all [21], [22], [38], the privacy report presents network activities at the level of individual domains in the App Network Activity, Website Network Activity, and Most Contacted Domains sections [2]. As described in Section II, these sections aggregate the same type of information, and subpages in each section lead to the same domain summary page (Figure 1c). As such, we analyzed participants’ reactions and understanding of these three sections as a whole and present the results together.

Participants had an immediate, core confusion about what the domains were and why there were so many (Section V-B1). Despite this, all participants recognized instances of third-party data collection and cross-app tracking (Section V-B2), and were concerned and confused about the details and implications (Section V-B3). Participants named some types of information that would help their understanding of the network activity (Section V-B4). Regardless of their level of understanding, they did not consistently identify realistic, effective actions to address concerns that came from learning about network activity, and were often resigned (Section V-B5). While similar high-level results (i.e., negative reactions and inability to draw actionable conclusions from difficult-to-understand information) have been reported in prior studies, we observed such results in a new context, as this privacy report conveys network activity at a level of detail not present in other privacy interfaces [8], [48], [54].

1) *Lack of basic understanding about domains:* The large quantity and technical nature of the domain names led all (20) participants to express confusion or surprise. Some domains were familiar or contained words that pointed to their origin or purpose (e.g., “analytics.tiktok.com” or “maps.googleapis.com”). However, more of the domains were for technical purposes unknown to participants, prompting descriptors like “scary” (P54) or “jarring” (P46). These domain names were often not very expressive, descriptive, or human-readable (e.g., “r2—sn-q4flrn7y.c.2mdn.net” or “yt3.ggpht.com”), with one participant describing them as “mumbo jumbo” (P71). Internet Protocol (IP) addresses that apps and websites contact are also reported, under the category “Unnamed Domains,” prompting further confusion for those who did not understand IP addresses. Almost all (17) participants expressed how little they could glean from the

domain names, commenting that they were “perplexed” (P97) with one participant asking, “what is a domain?” (P106)

Not only were participants confused by the domain names, but the way domains were presented in high quantity across multiple subpages was also confusing. Individual apps often had long lists of domains associated with them (i.e., in Figure 1b), and the final privacy report section, though entitled “Most Contacted Domains,” in fact lists every single domain that any app or website contacted, producing an even longer list (Figure 1d). From the 16 reports that we obtained from participants, there was an average of 1,763 unique domains present in the “Most Contacted Domains” section. Participants felt unclear as to why there would be such a quantity of domains and what they could be doing with the data. When looking at various pages of domains, one participant said,

I don’t recognize any of these, and then scrolling down to see like, 7000 “Unnamed Domains” is really strange. —P62

Almost all (17) participants asked at least one question about what the purpose of the domain contacts could be. The nature of these questions is discussed further in Section V-B4.

2) *Identifying third-party data collection and cross-app tracking:* We examined how well participants were able to recognize third-party data collection and cross-app tracking (RQ1, outlined in Section I). All (20) participants identified some instances of these two privacy risks occurring, but had several points of confusion about the details.

a) *Participants recognized but had questions about third-party data collection:* We categorized participants as identifying third-party data collection if they pointed out that to their knowledge, a domain was not from the app or website that they were currently looking at, and that their activity was being accessed by said domain. All (20) participants identified instances of this occurring, but asked questions about the nature of the network activity and the parties involved.

In addition to more general confusion about the purpose of network activity, participants’ surprised reactions varied depending on which apps or websites they saw were contacting unrecognizable domains. Participants were less surprised to see social media apps contacting unrecognizable domains, as they expected apps such as Facebook or TikTok would share their data with third parties for analytics or advertising purposes. In contrast, some (7) participants were more surprised and concerned by apps and websites they had not suspected would track them (e.g., basic mobile games, an app used to control Christmas lights, a state government website) also having large amounts of third-party domain contacts. When noting that her state government website seemed to be tracking her, one participant made a comparison to Amazon, saying,

Amazon...is a site that people have warned about tracking...so I’m a little bit surprised to see...it’s on the lower side in comparison to a government website. That feels strange to me. —P62

b) *Participants recognized but had questions about cross-app tracking:* The second privacy risk we explore in the network activity sections is cross-app tracking. In the domain summary page (Figure 1c), the privacy report includes

a label that refers to cross-app tracking: “If multiple apps or websites contact a domain, it could indicate the domain is combining your activity into a profile.” Not every participant read the label (some were prompted to), however, all (20) were able to identify cross-app tracking by pointing out multiple apps or websites interacting with a single domain. A few (3) participants expressed unease at the prospect of a profile being built about them, whereas the rest expressed that it was in line with their expectations relating to targeted advertising. A few (2) participants even felt reassured, or had more confidence in their apps’ behaviors after seeing many apps contact a single domain, because they felt it was more likely that the domain was contacted for a benign or legitimate purpose.

Participants wondered about the nature of potential profiles being made about them, for instance,

I don’t feel more informed...just a bit of frustration, ‘cause if you’re gonna tell me that it’s building a profile, what kind of profile? Why is it building a profile? —P18

P71 wanted to know if the profile was being built by the domain owner (Google) or by the app that was contacting that domain. P97 wondered about how secure this profile was and what other parties might have access to it. Some (6) participants expressed surprise or confusion at how apps or websites that they thought had no connection were all listed together under one domain and could be contributing to one profile. P148 wondered about how long she could be tracked for after interacting with a webpage, asking if, because she had made an online purchase there, she would now be tracked by that site “for all time.”

A few participants (2) were also unclear about which parties were sharing data, and in which direction. For instance, P62 said that her banking app was either “getting tracked or doing tracking,” but could not distinguish any further. When seeing a list of apps that all contacted one domain, they thought the apps were also talking to each other. These participants were concerned that their data was not only being transmitted to external domains, but also to other unrelated apps.

3) Confusion and concern about the purpose of network activity: Though participants identified instances of the two privacy concerns investigated in these sections, there was more confusion and surprise about the purpose, quantification, and negative implications of the network activity.

a) Domains for functionality vs. advertising: All (20) participants attributed the network activity they saw to advertising purposes. Most (14) also noted that the activity could be occurring for functionality purposes, though some needed to read additional documentation to recognize this. Within the privacy report’s documentation (accessible by tapping on “Learn more...” from the domain summary page in Figure 1c), Apple provides some examples of why network activity occurs (e.g., “provide video streaming or game play content, or connect you to other devices”). In some cases (7) during interviews, after we saw that a participant had an incomplete picture of why network activity occurs, we pointed them toward the documentation. In all such cases, reading the above explanation improved participants’ understanding of why network activity occurs. But only a few (3) participants

navigated to this page on their own, and even fewer read it unprompted, even after navigating to it. Though the “Learn more” page provides helpful examples about why network activity occurs, its location is one that users may not actually go to when using the privacy report. As a result, without additional guidance, some participants would come away with the inaccurate view that network activity occurs solely for targeted advertising.

b) Volume of network activity vs. usage of app: Participants’ interpretation of the volume of network activity also varied. As described in Section II and pictured in Figure 1, the network activity sections have quasi-bar charts with the number of unique domains an app or website contacted and the number of times each domain was contacted. Participants were confused by these numbers; instead of relating the quantities to how active an app or website was in contacting domains, about half (9), at least initially, related the numbers to how much they themselves used the app or website. Accordingly, some (7) were surprised at how certain apps’ activity levels compared to how much they used that app. In some cases, participants expressed pleasant surprise if a frequently used app had lower activity than anticipated. In other cases, seldomly-used apps with high activity levels caused worry. For instance, P98 was concerned to note that a currency conversion app was the second most active in the amount of network activity, even though he rarely used it.

c) Potential misuses of personal data: As the reasons behind network activity were unclear, participants also had questions about the implications. While most (16) participants related the network activity to targeted advertising, a few questioned whether the activity could mean that their personally identifiable or otherwise sensitive information was at greater risk, or if there were other unknown consequences. Some (6) participants paid more attention to certain kinds of apps’ network activity due to the more sensitive functions done through those apps, such as banking and gambling. One participant said that she knew targeted advertising was widespread but wondered if tracking could be taking place for other unwanted purposes that “we’re less aware of” (P62).

A few (2) participants expressed outsized concern and suggested negative implications that would be difficult to infer using only the information in the privacy report. For instance, when he saw apps contacting domains but did not know exactly why those domains were contacted, P71 expressed concern about his banking apps, and wondered if this meant there was hacking occurring. P109 related that a shopping website had her payment information stored and wondered “what the connection or correlation between the [domain and the shopping website] is,” “how much access [domains] have,” and whether the domain she was looking at could, for instance, get access to her music app and then delete her music library. From looking at the privacy report, these participants had specific negative consequences in mind that to our knowledge are highly unlikely, indicating that the privacy report can cause a disproportionate level of worry.

4) Information participants wanted to see: When asked what would make the network activity sections more helpful, participants most commonly wanted more explanations, including details about individual domains and general guidance about what they should be taking away.

a) *Participants wanted descriptions for domains:* Participants wanted the privacy report to contain more easy-to-understand information about the domains, since it was often difficult for them to learn anything from the domain names. A commonly expressed sentiment was “I don’t really know what any of this stuff is” (P79). Questions included, “What is that domain for?” (P10), “How are they tracking [me]?” (P18), and “What type of data are they collecting?” (P34). A few (3) participants noted how they could not click on the domain name or copy the text to search for more information.

Participants specifically wanted to know who owned the domains. The privacy report includes labels for a few domain owners (e.g., “Google LLC” in Figure 1d), but more often, owners are unidentified. About half (9) of the participants said that knowing the domain owner was informative and could provide context or credibility to associated activity. A few (2) participants noted that activity originating from well-known companies like Amazon and Facebook made them feel more comfortable than if the origin of the domain was unknown.

b) *Participants wanted to know the purpose of domain contacts, especially which ones were necessary vs. not:* In addition to identifying the domains, participants wanted to know why apps and websites were contacting them. For certain recognizable domains, most (12) participants asked about why they were in contact with seemingly unrelated apps or websites. For instance, Panda Express (restaurant) app and Snapchat (P46), or a recipe website contacting LinkedIn (P62). In comparison, for social media apps like Instagram or TikTok, participants had ideas about the kinds of data that those apps might be giving to third-party domains. For other interactions between apps and domains that they could not intuit a reason for, some (6) participants asked about what specific types of data were being transferred.

Some (7) participants also wanted to be able to distinguish between domains contacted for advertising and tracking related reasons, versus for necessary functionality. They proposed that each domain’s purpose could be labeled (e.g., providing fonts), as well as labeling the reason for each instance of an app or website accessing a domain to better understand what kind of data might be shared.

c) *Participants wanted interpretation, e.g., “Should I be concerned?”:* More abstractly, a few (4) participants said that they did not know what conclusions they should be making. P106 noted that the privacy report’s explanation about potential profile building “makes you worry,” then asked, “how would I know that my privacy is negatively impacted?” They noted that it was difficult to know how worried to feel, which details they should pay attention to, or along what axis they should direct their concern. P54 described that the privacy report only allowed “inferences that something’s happening,” as opposed to a clear picture. Another participant suggested the privacy report could highlight certain app behaviors so that they could better tell what to be alarmed about:

I...wish [the report] would flag it, like, ‘It is alarming that Discord is contacting all these different websites,...or, ‘You should be fine with this,’ ‘You should have more questions about this.’ —P109

5) *Lack of effective means to address concerns:* Participants could not identify an effective, realistic action that would

address their privacy concerns, or were resigned and did not seek action. About half (10) of the participants said they ultimately would not change anything about how they used their phones in response to what they saw in the network activity sections. In a few cases, their inaction was because the network activity was not bothersome, or did not bother them enough to change something, for instance,

I’m not thrilled about it, but I’m not gonna go out of my way to do anything about it. —P97

For the participants who did want to act on what they learned, they either could not identify any actions to take or found actions they did know of to be undesirable in some way.

a) *Some participants could not identify any actions at all, even when they wanted to do something:* A few (4) participants expressed that though they wanted to mitigate the concerns that arose from learning about the network activity, they did not know what kinds of actions were available or possible. For instance, P46 was worried, but said,

I don’t really know how to change my behavior in a way that would make me more comfortable. —P46

In contrast to those who were not bothered by the activity that they saw, these participants wanted to change either something about the phones or something about their habits, but did not know what would be helpful.

b) *The actions participants did identify seemed unrealistic or ineffective to them:* Most (14) participants identified potential actions they could take in response to what they saw in the network activity sections. But in comparison to the Data and Sensor Access section, where participants consistently brought up changing app permissions (Section V-A2), there was a greater variety of remedial actions identified to address the perceived risks from learning about cross-app tracking and third-party data collection. Participants also often found these actions unsatisfactory in some way in having low perceived efficacy or negative impact on their ability to use the app services.

One action participants discussed was looking for more information about certain domains, through a search engine, internet forum, or by asking someone they knew. While most (13) participants either took time during the interview to lookup a domain or expressed an intention to do so later, a few (2) expressed that looking up individual domains was not a realistic course of action, for instance,

Am I supposed to search all these [domains] one by one?...it’s overwhelming. The regular person is not gonna do that. —P113

Many domains listed are not accessible if entered into the address bar of web browsers in the same way that URLs are. For instance, P98 tried to visit a domain from his report, and said, “nothing shows up.” Almost none of the participants who wanted to look up domains said anything concrete about what they would do after learning that information.

Half (10) of the participants also brought up (at least as a supposition) deleting an app or stopping the usage of a website to prevent unwanted network activity. However, for

most (6/10), deletion was infeasible because the service the app or websites provided was necessary to them, often for their jobs. Moreover, as in Section V-A2, of the participants who indicated they would delete an app, many (4) said it was because they were not using the app. There was not necessarily a specific privacy concern that motivated deletion, but rather a broader awareness of their data being shared even though they were not getting any value out of the app.

Another action a few (3) participants described was the possibility of exercising fine-grained control over the network activity, by blocking specific domains. Along with labels to distinguish between domains that were for functionality purposes and those for advertising or otherwise less desirable purposes (as discussed in Section V-B4), they wanted to prevent their data from being shared with the non-essential domains. Of the three participants who brought up this idea, most (2/3) then said this kind of control was not currently possible. In particular, when discussing the Facebook app, P34 pointed out how the app's settings did not allow blocking specific domains. Another participant noted the risk of compromised functionality if they elected to block many of the domains:

There's limited account options and privacy options to most of these apps...I can set all my privacy settings to the highest settings...but...what the hell am I gonna be able to use [Facebook] for, if I don't allow certain things? —P54

These participants' understanding was that they did not have a level of control, either through the phone or app settings, where they could isolate certain troubling domains.

A few (3) more participants suggested using privacy tools such as ad blockers or search engines and browsers they knew to be less invasive than Google, such as DuckDuckGo, to reduce unwanted network activity, but expressed hesitation to switch to tools they were not used to. P159 mentioned that they already used ad blockers, but still saw worrying network activity in their report and did not know what else to do.

c) Regardless of their level of understanding, most participants were resigned: Most (12) participants expressed some degree of resignation to the volume of domains, both known and unknown, that their apps and websites contacted, as well as the implications they understood. They acknowledged how their data was handled by apps and websites felt invasive or excessive, but that they were willing to still use the service, often because they saw no better alternative.

If I visit websites, there's nothing I can really do to prevent this from happening. —P34

Some (6) participants discussed the privacy tradeoff, describing that they expected to give up some privacy to benefit from otherwise free services. One participant was also resigned to simply not being able to understand the network activity information, even with further explanations:

It doesn't really matter how it's presented. I think it's gonna go over my head regardless. —P97

C. Follow-up Survey

In the follow-up survey, we aimed to find out if participants used the privacy report after the interview concluded, and if

so, what their interactions involved. Of the 14 participants who completed the follow-up survey, almost all (13) kept their privacy report turned on, and most (8) checked the privacy report again after the interview. Half (4/8) of those who looked at the report again said they checked it once or twice in the last month, while the rest checked more often.

The main reason that participants checked their privacy reports was curiosity about apps' access to data and sensors as well as network activities. For example, P159 said they have "downloaded new apps and wanted to see what [apps] are all using." They also found such information useful and some (5) participants took action after seeing the privacy report. Some (4) participants looked up domain names to learn more about them, while others changed their browsing habits or deleted apps, for instance,

I've deleted TikTok and I've tried to avoid certain sites that have a lot of tracking on them. —P18

The participants who said during the interview that they would use the privacy report again but said they did not end up doing so attributed this to forgetting or lacking time.

D. Answering Research Questions

Below, we summarize our findings by addressing each of the research questions.

a) RQ1. Does the privacy report help users identify and understand overpermissioning, cross-app tracking, or third-party data collection happening on their phones?: The privacy report helped participants both identify and understand overpermissioning. In the Data and Sensor Access section, participants consistently pointed out times or types of permissions that felt inconsistent with app functions. In the Network Activity sections, all participants identified instances of cross-app tracking and third-party data collection; however, they often did not understand basic concepts related to apps or websites contacting domains, such as what domains are, why apps or websites contacted them for reasons other than targeted advertising, what kind of user data the domain would receive, and what the implications were.

b) RQ2. What are users' attitudes toward the information they learn from the privacy report?: Participants were largely unsurprised by the permissions apps exercised. Individual cases of surprise, curiosity, and sometimes concern were prompted by instances of unexpected times or types of access to data or sensors on their phone.

On the other hand, participants were confused, surprised, and often overwhelmed by the amount of difficult-to-understand domains. They were concerned about what types of data was shared with third parties and who was doing what with their data. A few participants felt more reassured by seeing some domain information that lent credibility to the activity (e.g., the domain was owned by a company they trusted). Many participants were also resigned to the high amount of network activity they saw.

c) RQ3. Does interacting with the privacy report influence users to change their app and/or privacy behavior?: Interacting with the Data and Sensor Access section motivated about half of the participants to want to rescind certain app

permissions. Their expectations of having this control option did not always match the available phone settings, especially for built-in apps. A few (4) participants wanted to delete overpermissioned apps that they also rarely used.

As the Network Activity sections were more confusing overall, it was difficult for participants to identify actions that they were both willing to take and felt would be effective in reducing unwanted third-party data collection and cross-app tracking. When they imagined remediations, participants often felt they were unrealistic or only hypothetical: deleting an app, blocking individual domains, or using privacy tools such as ad blockers. The one action that was more accessible was looking up more information about individual domains, though it was unclear what participants would do with the additional information. The follow-up survey showed that few participants actually looked up more information about domains in the one month following their interview, perhaps indicating the unlikelihood of such action.

VI. DISCUSSION

In this section, we highlight key results and discuss how these findings relate to previous work on privacy dashboards. Based on these results, we make recommendations to help improve the comprehensibility of privacy dashboards for users and help them make more informed privacy decisions. In the first part, we echo previous work in recommending that providers of privacy transparency tools include explanations for why apps access users' data and sensors. Doing so could improve users' understanding of app permissions and make them more at ease with granting permissions (Section VI-A). We then discuss the presentation of network activity, which is a less common feature in existing privacy transparency tools. The privacy report's implementation of this, though informative in some ways, was challenging for users to understand. We suggest ways to improve the comprehensibility of network activity for users through labels and summaries. Finally, we suggest a way to make the previous recommendations more feasible by separating domains by their purpose (Section VI-B).

A. Explain Reasons for Data and Sensor Accesses

Our results align with some previous findings about how users react to learning about data and sensor access. The privacy report's Data and Sensor Access section conveys information about apps' access to users' data, similar to privacy dashboards like Google's "Apps with Access to Your Account" page or Android's Privacy Dashboard [21], [22]. Participants expressed a general understanding that phone apps require access to certain data and sensors to provide functionality (e.g., location for navigation), and what they saw in this section of the privacy report matched their expectations in most cases. This result confirms previous findings that privacy dashboard users are not likely to be surprised or concerned when seeing information (such as third-party access to Google account information) *that they can understand* [5], [7].

Though participants were mostly unsurprised by what they saw in the Data and Sensor Access section, almost all (17) identified specific unexpected instances of an app accessing data that did not align with how participants used that app

or thought it functioned (e.g., iPhone's built-in Health app accessing contacts data). About half (9) of the participants raised concerns about instances of background access or unexpected time of access (e.g., Instagram accessing photos when users did not post new photos). In some cases, these concerns caused participants to want to change apps' permissions or consider uninstallation.

Participants' concerns were rooted not solely in the apps' usage of data and sensors (for a specific purpose) but in the fact that the information provided in the privacy report did not allow them to determine whether such usage aligned with their privacy preferences. Providing explanations for the motivations behind access would likely have made participants more comfortable with the apps' behavior, which we believe is in the interests of both users and app developers.

Recommendation 1. *Include reasons for data and sensor access.* We suggest that providers of privacy dashboards (e.g., Apple or Google) encourage app developers to reveal the reasons that apps access data or sensors (e.g., the Health app uses contacts data to reach emergency contacts). Our suggestion is supported by previous studies that show users respond positively and are more likely to approve apps using their data when the app provides an explanation [48]. A study on Google's "Apps with Access to Your Account" page similarly recommended that privacy dashboards include contextual information (e.g., reasons for permission) [5]. Further work could explore the optimal context and methods to present the explanations we suggested to users. Besides including explanations within the privacy dashboards themselves, another possible method is through run-time notifications (similar to the current location access notifications system, which found success on mobile devices [18]).

Many others have recommended explaining the reasons behind data access [5], [7], [16]. However, there seems to be little industry incentive to implement such changes. That is, though prior work has indicated that explaining why their data is being accessed can improve user experience, privacy transparency tools do not seem to be moving in this direction. Our work further emphasizes that these changes seem necessary in order to improve users' ability to make informed privacy decisions. Furthermore, we recognize that this recommendation may not be easy to implement, as it requires actions by both platform providers and developers that they may not be sufficiently motivated to take; nonetheless, our findings demonstrate the need for such explanations in privacy dashboards and we leave it to future work to investigate ways to implement them.

B. Label and Summarize Network Activities

Presenting detailed network activity is not a common feature in privacy transparency tools. The privacy report's Network Activity sections list every domain contacted by a user's apps and websites, for purposes such as functionality, advertising, and analytics. Some prior studies have also aimed to show network activity to users to help them understand, for instance, online tracking [54] or the data handling processes of home IoT devices [29]. However, little is known about how users perceive such information. These prior attempts have not been as widely available as Apple's privacy report feature, as to the best of our knowledge, network activity has not

been included as part of comparable privacy dashboards (e.g., Meta’s Privacy Center, Android’s Privacy Dashboard) outside of research studies [21], [38].

Our work sheds light on how helpful (or unhelpful) users find the network activity information shown to them in the privacy report. Participants from our study often found information that was new to them in the network activity sections of the privacy report. Because of this, we encourage developers of other similar privacy dashboards to try to convey information about network activity to users. More work is needed to find out how best to convey that information, however, as we found that participants were often very confused by the privacy report’s presentation of network activity. Most of the domain names were uninterpretable and many participants were unable to intuit what types of data were transferred to and from the domains, for what purposes, and whether they should be concerned. Participants were also often overwhelmed by the large number of domains they encountered in their privacy reports (on average, the privacy report included 1,763 domains per participant). Therefore, we make two suggestions for presenting network activity to users, and one broader suggestion about how to make these steps more feasible for developers of privacy transparency tools.

Recommendation 2. *Label domains with their purpose.* We suggest labeling the domains in the privacy report to explain why apps and websites are in contact with the domain. Users can then better interpret the general purpose (e.g., functionality or tracking) of an app or website contacting a domain and have a better basis to decide how to interact with that app or website. While presenting network activities is novel to privacy dashboards, categorizations of domains into first- and third-party, or tracking and non-tracking exist (though are not typically exposed to end users), and are often used by studies measuring online tracking [9], [13], [34]. For example, the webXray list includes information on domain ownership, allowing distinction between first- and third-party domains [35]. Using existing domain categorization lists to help users understand online tracking has been done in prior studies [54]. Therefore, we suggest privacy dashboard providers (e.g., Apple) incorporate information provided by the aforementioned lists to help users understand domains beyond what the domain names reveal (which is often very little).

Recommendation 3. *Summarize domain information.* We further suggest that the privacy report provide options to aggregate domains or to synthesize the raw domain information that it currently provides. Prior work has provided similar suggestions when studying existing privacy tools [5], [8]. We believe such summarization could benefit users of the app privacy report (and privacy transparency tools in general). Besides lacking the right types of information about domains to allow them to reach useful conclusions, participants also reacted negatively to the long lists of domains they were shown. The privacy report could allow users to view domains grouped by owners or purposes of domain contacts, as opposed to listing all the domains on a single page. Another way to reduce the volume of information presented to users is to provide higher-level summaries. Simpler, more digestible metrics could include the total number of tracking-related domains that have been contacted by apps. Doing so could help users better gauge the extent to which tracking is occurring

through their app activity.

Recommendation 4. *Disentangle domain functionality to make labeling easier.* Categorizing domains according to why they are contacted may be difficult to implement. The same domain can be accessed for both key functionalities (e.g., providing fonts) and for other goals (e.g., tracking users’ online activity), making it sometimes impossible to delineate domains by their purpose. Future work could examine how to overcome this challenge. For example, app developers could try to design apps so that they use each domain only for a single purpose (e.g., just for functionality vs just for tracking). Platform providers could attempt to design mobile OSes and APIs that help developers (or require them to) explain the purposes of network accesses so that these could later be conveyed to users via privacy dashboards. Regardless of the implementation, the disambiguation of network accesses by purpose could help users understand the behavior of their apps and help them make informed decisions about what measures to take to control that behavior.

VII. CONCLUSION

Through an interview study of users’ perceptions of the privacy report, we found that the effectiveness of tools that aim to increase privacy transparency depends on the comprehensibility of the information presented. Participants understood how and when apps accessed their data and sensors, and identified privacy settings they wanted to change to more closely match their preferences. In contrast, participants were surprised and confused by network activities, especially about the number of domains their apps and websites contacted, the purpose, and what the implications were. In their identification of third-party data collection and cross-app tracking, participants could not consistently determine how concerned they should be or what they could do to address these privacy risks. To help users better understand and act on the information about their data and devices, we recommend the privacy report and similar privacy dashboards label, explain, and summarize information about how apps and websites access user data and interact with third-party domains. Implementing this may require creating infrastructure (e.g., in the OS or development environments) to enable such explanations or aggregation.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers and the shepherd for their feedback; Weiran Lin, Jenny Tang, Shuyang Deng, Wenhao Song, and Yingke Chen for their help in piloting the study; and Ben Weinshel for his insights about the App Privacy Report. This work is supported in part by the National Science Foundation under grant CNS-2114148.

REFERENCES

- [1] H. Almuhammedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal, “Your location has been shared 5,398 times!: A field study on mobile app privacy nudging,” in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2015, pp. 787–796.
- [2] Apple, “About app privacy report,” <https://support.apple.com/en-us/102188>, last Accessed: January 27, 2024.

- [3] —, “Control access to information in apps on iphone,” <https://support.apple.com/guide/iphone/control-access-to-information-in-apps-iph251e92810/ios>, last Accessed: February 7, 2024.
- [4] D. G. Balash, M. M. Ali, X. Wu, C. Kanich, and A. J. Aviv, “Longitudinal analysis of privacy labels in the apple app store,” <http://arxiv.org/abs/2206.02658>.
- [5] D. G. Balash, X. Wu, M. Grant, I. Reyes, and A. J. Aviv, “Security and privacy perceptions of Third-Party application access for google accounts,” in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 3397–3414.
- [6] K. Belgum, “Who leads at halftime?: Three conflicting visions of internet privacy policy,” *Richmond Journal of Law & Technology*, vol. 6, no. 1, p. 1, 1999.
- [7] F. Bemmam, M. Windl, J. Erbe, S. Mayer, and H. Hussmann, “The influence of transparency and control on the willingness of data sharing in adaptive mobile apps,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, pp. 1–26, 2022.
- [8] W. Cao, C. Xia, S. T. Peddinti, D. Lie, N. Taft, and L. M. Austin, “A large scale study of user behavior, expectations and engagement with android permissions,” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 803–820.
- [9] D. Cassel, S.-C. Lin, A. Buraggina, W. Wang, A. Zhang, L. Bauer, H.-C. Hsiao, L. Jia, and T. Libert, “OmniCrawl: Comprehensive measurement of web tracking with real desktop and mobile browsers,” *Proceedings on Privacy Enhancing Technologies*, 2022.
- [10] L. Cranor, “P3p: making privacy policies more useful,” *IEEE Security & Privacy*, vol. 1, no. 6, pp. 50–55, 2003.
- [11] S. Egelman, A. P. Felt, and D. Wagner, “Choice architecture and smartphone privacy: There’s a price for that,” in *The Economics of Information Security and Privacy*, R. Böhme, Ed. Springer, 2013, pp. 211–236.
- [12] 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*, ser. CHI ’19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 1–12.
- [13] S. Englehardt, D. Reisman, C. Eubank, P. Zimmerman, J. Mayer, A. Narayanan, and E. W. Felten, “Cookies that give you away: The surveillance implications of web tracking,” in *Proceedings of the 24th International Conference on World Wide Web*, ser. WWW ’15. International World Wide Web Conferences Steering Committee, 2015, pp. 289–299.
- [14] F. M. Farke, D. G. Balash, M. Golla, and A. J. Aviv, “How does connecting online activities to advertising inferences impact privacy perceptions?” *Proceedings on Privacy Enhancing Technologies*, vol. 2024, no. 2, pp. 372–390, 2024.
- [15] F. M. Farke, D. G. Balash, M. Golla, M. Dürmuth, and A. J. Aviv, “Are privacy dashboards good for end users? evaluating user perceptions and reactions to google’s my activity,” in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 483–500.
- [16] A. P. Felt, S. Egelman, M. Finifter, D. Akhawe, and D. Wagner, “How to ask for permission,” in *7th USENIX Workshop on Hot Topics in Security (HotSec 12)*. Bellevue, WA: USENIX Association, Aug. 2012.
- [17] S. Fischer-Hübner, J. Angulo, F. Karegar, and T. Pulls, “Transparency, privacy and trust – technology for tracking and controlling my data disclosures: Does this work?” in *Trust Management X*, ser. IFIP Advances in Information and Communication Technology, S. M. Habib, J. Vassileva, S. Mauw, and M. Mühlhäuser, Eds. Springer International Publishing, 2016, pp. 3–14.
- [18] H. Fu, Y. Yang, N. Shingte, J. Lindqvist, and M. Gruteser, “A field study of run-time location access disclosures on android smartphones,” in *Proceedings 2014 Workshop on Usable Security*. Internet Society, 2014.
- [19] S. E. Gindin, “Creating an online privacy policy implementing the privacy policy,” *Preventive Law Reporter*, vol. 18, no. 3, pp. 10–12, 1999.
- [20] J. Gladstone, “The u.s. privacy balance and the european privacy directive: Reflections on the united states privacy policy the impact of e-commerce on the laws of nations,” *Willamette Journal of International Law and Dispute Resolution*, vol. 7, pp. 10–32, 2000.
- [21] Google, “Android privacy settings and permissions | android,” <https://www.android.com/safety/privacy/>, last Accessed: January 27, 2024.
- [22] —, “Share some access to your google account with third-party apps - google account help,” https://support.google.com/accounts/answer/14012355?visit_id=638489582225474141-3518119556&rd=1, last Accessed: April 17, 2024.
- [23] —, “How personalized ads work - android - my ad center help,” <https://support.google.com/My-Ad-Center-Help/answer/12155656?hl=en>, 2024, last Accessed: January 31, 2024.
- [24] M. A. Graber, D. M. D’Alessandro, and J. Johnson-West, “Reading level of privacy policies on internet health web sites,” *The Journal of Family Practice*, 2002.
- [25] T. Groß, “Validity and reliability of the scale internet users’ information privacy concern (iupc) [extended version],” <https://doi.org/10.48550/arXiv.2011.11749>, 2020, last Accessed: January 27, 2024.
- [26] J. Gu, Y. C. Xu, H. Xu, C. Zhang, and H. Ling, “Privacy concerns for mobile app download: An elaboration likelihood model perspective,” *Decision Support Systems*, vol. 94, pp. 19–28, 2017.
- [27] 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*, ser. CHI ’20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1–12.
- [28] E. Herder and O. van Maaren, “Privacy dashboards: The impact of the type of personal data and user control on trust and perceived risk,” in *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, ser. UMAP ’20 Adjunct. New York, NY, USA: Association for Computing Machinery, 2020, pp. 169–174.
- [29] D. Y. Huang, N. Aphorpe, G. Acar, F. Li, and N. Feamster, “IoT inspector: Crowdsourcing labeled network traffic from smart home devices at scale,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 4, no. 2, pp. 1–21, 2020.
- [30] M. Janic, J. P. Wijbenga, and T. Veugen, “Transparency enhancing tools (tets): An overview,” in *2013 Third Workshop on Socio-Technical Aspects in Security and Trust*, 2013, pp. 18–25.
- [31] J. Jung, S. Han, and D. Wetherall, “Short paper: enhancing mobile application permissions with runtime feedback and constraints,” in *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*, ser. SPSM ’12. Association for Computing Machinery, 2012, pp. 45–50.
- [32] G. Karjoth and M. Schunter, “A privacy policy model for enterprises,” in *Proceedings 15th IEEE Computer Security Foundations Workshop. CSFW-15*, 2002, pp. 271–281.
- [33] 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*, ser. SOUPS ’09. New York, NY, USA: Association for Computing Machinery, 2009.
- [34] T. Libert, “Exposing the invisible web: An analysis of third-party HTTP requests on 1 million websites,” *International Journal of Communication*, vol. 9, no. 0, p. 18, 2015, number: 0.
- [35] —, “An automated approach to auditing disclosure of third-party data collection in website privacy policies,” in *Proceedings of the 2018 World Wide Web Conference*, ser. WWW ’18. International World Wide Web Conferences Steering Committee, 2018, pp. 207–216.
- [36] S. Lichtenstein, P. Swatman, and K. Babu, “Effective online privacy policies,” *ACIS 2002 Proceedings*, 2002.
- [37] A. M. McDonald, R. W. Reeder, P. G. Kelley, and L. F. Cranor, “A comparative study of online privacy policies and formats,” in *Privacy Enhancing Technologies*, I. Goldberg and M. J. Atallah, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 37–55.
- [38] Meta, “Meta privacy center,” <https://www.facebook.com/privacy/center/>, last Accessed: January 27, 2024.
- [39] J. Millham and R. W. Kellogg, “Need for social approval: impression management or self-deception?” *Journal of research in Personality*, vol. 14, no. 4, pp. 445–457, 1980.
- [40] E. Rader, “Awareness of behavioral tracking and information privacy

- concern in facebook and google,” in *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, 2014, pp. 51–67.
- [41] P. Raschke, A. Küpper, O. Drozd, and S. Kirrane, “Designing a GDPR-compliant and usable privacy dashboard,” in *Privacy and Identity Management. The Smart Revolution: 12th IFIP WG 9.2, 9.5, 9.6/11.7, 11.6/SIG 9.2.2 International Summer School, Ispra, Italy, September 4-8, 2017, Revised Selected Papers*, ser. IFIP Advances in Information and Communication Technology, M. Hansen, E. Kosta, I. Nai-Fovino, and S. Fischer-Hübner, Eds. Springer International Publishing, 2018, pp. 221–236.
- [42] I. Reyes, P. Wijesekera, J. Reardon, A. Elazari Bar On, A. Razaghpanah, N. Vallina-Rodriguez, and S. Egelman, ““won’t somebody think of the children?” examining COPPA compliance at scale,” in *The 18th Privacy Enhancing Technologies Symposium (PETS 2018)*, 2018.
- [43] W. Rodger, “Privacy isn’t public knowledge: Online policies spread confusion with legal jargon,” *USA Today*, vol. 1, p. 3D, 2003.
- [44] F. Roesner, T. Kohno, A. Moshchuk, B. Parno, H. J. Wang, and C. Cowan, “User-driven access control: Rethinking permission granting in modern operating systems,” in *2012 IEEE Symposium on Security and Privacy*, 2012, pp. 224–238, ISSN: 2375-1207.
- [45] G. L. Scoccia, S. Ruberto, I. Malavolta, M. Autili, and P. Inverardi, “An investigation into android run-time permissions from the end users’ perspective,” in *2018 IEEE/ACM 5th International Conference on Mobile Software Engineering and Systems (MOBILESoft)*, 2018, pp. 45–55.
- [46] B. Shebaro, O. Oluwatimi, and E. Bertino, “Context-based access control systems for mobile devices,” *IEEE Transactions on Dependable and Secure Computing*, vol. 12, no. 2, pp. 150–163, 2015.
- [47] L. J. Strahilevitz and M. B. Kugler, “Is privacy policy language irrelevant to consumers?” *The Journal of Legal Studies*, vol. 45, pp. S69–S95, 2016, publisher: The University of Chicago Press.
- [48] J. Tan, K. Nguyen, M. Theodorides, H. Negrón-Arroyo, C. Thompson, S. Egelman, and D. Wagner, “The effect of developer-specified explanations for permission requests on smartphone user behavior,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI ’14. Association for Computing Machinery, 2014, pp. 91–100.
- [49] J. Tang, E. Birrell, and A. Lerner, “Replication: How well do my results generalize now? the external validity of online privacy and security surveys,” in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. Boston, MA: USENIX Association, Aug. 2022, pp. 367–385.
- [50] J. Tang, U. Akram, and W. Shi, “Why people need privacy? the role of privacy fatigue in app users’ intention to disclose privacy: based on personality traits,” *Journal of Enterprise Information Management*, vol. 34, no. 4, pp. 1097–1120, 2020, publisher: Emerald Publishing Limited.
- [51] C. Thompson, M. Johnson, S. Egelman, D. Wagner, and J. King, “When it’s better to ask forgiveness than get permission: attribution mechanisms for smartphone resources,” in *Proceedings of the Ninth Symposium on Usable Privacy and Security*, ser. SOUPS ’13. Association for Computing Machinery, 2013, pp. 1–14.
- [52] J. Turow, “Americans and online privacy: The system is broken,” https://repository.upenn.edu/asc_papers/526, 2003, last Accessed: January 31, 2024.
- [53] T. Vila, R. Greenstadt, and D. Molnar, “Why we can’t be bothered to read privacy policies models of privacy economics as a lemons market,” in *Proceedings of the 5th international conference on Electronic commerce*, ser. ICEC ’03. Association for Computing Machinery, 2003, pp. 403–407.
- [54] B. Weinschel, M. Wei, M. Mondal, E. Choi, S. Shan, C. Dolin, M. L. Mazurek, and B. Ur, “Oh, the places you’ve been! user reactions to longitudinal transparency about third-party web tracking and inferencing,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’19. Association for Computing Machinery, 2019, pp. 149–166.
- [55] A. F. Westin, *Privacy and Freedom*. New York, NY, USA: Atheneum, 1967.
- [56] S. Zhang, Y. Feng, Y. Yao, L. F. Cranor, and N. Sadeh, “How usable are iOS app privacy labels?” in *Proceedings on Privacy Enhancing Technologies*, 2022.
- [57] C. Zimmermann, R. Accorsi, and G. Müller, “Privacy dashboards: Reconciling data-driven business models and privacy,” in *2014 Ninth International Conference on Availability, Reliability and Security*, 2014, pp. 152–157.

APPENDIX

A. Screening Survey

- S1 Which model of iPhone do you own?
Instructions on how to check iPhone model numbers are provided.
- o 1st generation
 - o 3G
 - o ... (list of iPhone models)
 - o 15 Pro
 - o 15 Pro Max
- S2 What iOS version does your iPhone run?
Instructions on how to find the iOS version are provided.
- o Older than iOS 15.2 (please specify) _____
 - o Newer than or iOS 15.2 (please specify) _____
- S3 Are you willing to participate in a recorded online interview study conducted via Zoom?
- o Yes, I would like to participate with my webcam turned on.
 - o Yes, I would like to participate with my webcam turned off.
 - o No
- S4 What risks are you aware of regarding apps’ access to data stored on your iPhone? [free response]
- S5 What risks are you aware of regarding apps’ access to sensors (camera, microphone, GPS, etc.) on your iPhone? [free response]
- S6 What risks are you aware of regarding apps’ on your iPhone having the ability to connect to the internet? [free response]
- S7 When you use your iPhone, what features that protect your privacy are you aware of? [free response]
- S8 When you use your iPhone, what features that protect your privacy have you used? [free response]
- S9 Have you previously heard of iPhone’s App Privacy Report?
- o Yes
 - o No
 - o Not sure
- [S10 is shown if S9 is “Yes”]
- S10 Have you used iPhone’s App Privacy Report?
- o Yes
 - o No
 - o Not sure
- S11 What is your age?
- o Under 18
 - o 18 - 24
 - o 25 - 34
 - o 35 - 44
 - o 45 - 54
 - o 55 - 64
 - o 65 - 74
 - o 75 - 84
 - o 85 or older
- S12 What is your gender?
- S13 How would you describe yourself? Please select all that apply.
- White
 - Black or African American
 - American Indian or Alaska Native
 - Asian
 - Native Hawaiian or Pacific Islander
 - Hispanic or Latino
 - Other
- S14 What is the highest degree or level of school you have completed?
- o Less than high school

- o High school degree or equivalent (e.g., GED)
- o Some college, no degree
- o Associate degree (e.g., AA, AS)
- o Bachelor's degree (e.g., BA, BS)
- o Master's degree (e.g., MA, MS, MEd)
- o Doctorate or professional degree (e.g., MD, DDS, PhD)
- o Professional degree (e.g., JD, MD)

S15 Please indicate how much you agree or disagree with each statement on the left.
See Table IV

B. Interview Script

We start the interview by giving participants a brief introduction of ourselves and the study. After the introductions, we give participants a chance to ask questions. We then ask participants if they agree to continue with the interview. If the participant agrees, we turn on the recording and proceed to the following semi-structured script.

- I1 Based on participants' response to S4 through S8, we ask follow-up questions on participants' privacy awareness as well as actions they take to protect their privacy when using their iPhones.

[I2 and I3 are asked if S9 is "Yes" and S10 is "No"]

- I2 What do you know about the report?

- I3 Why did you not use the feature?

[I4 through I6 are asked if both S9 and S10 are "Yes"]

- I4 What did you use the privacy report for?

- I5 What did you like about it?

- I6 What did you dislike about it?

Participants are given time to freely explore the entire report. They are asked to describe what they think verbally during the exploration.

Data & Sensor Access

- I7 Find apps that accessed sensors you did not expect.

[I8 is asked if the participant found unexpected access in I7]

- I8 For this access, do you think it is reasonable?

- I9 Find apps that accessed sensors at times you did not expect, or times you do not recall using the app.

- I10 Does the information from this section make you feel differently about how you interact with your phone?

- I11 How would you summarize this section to someone else?

- I12 What would you change about this section, if anything?

- I13 After seeing the information here, do you feel like you would want to do anything differently with these apps?

App Network Activity

- I14 Explain what the bars mean.

- I15 Explain what the app could do when contacting a domain.

- I16 Find an example of a domain that was contacted by multiple apps.

- I17 Find a couple of examples of first-party and third-party domains.

- I18 Does the information from this section make you feel differently about how you interact with your phone?

- I19 How would you summarize this section to someone else?

- I20 What would you change about this section, if anything?

- I21 After seeing the information here, do you feel like you would want to do anything differently with these apps?

Website Network Activity

- I22 Explain what the bars mean.

- I23 Explain what a website could do when contacting a domain.

- I24 Find an example of a domain that was contacted by multiple websites.

- I25 Find a couple of examples of first-party and third-party domains.

- I26 Does the information from this section make you feel differently about how you interact with your phone?

- I27 How would you summarize this section to someone else?

- I28 What would you change about this section, if anything?

- I29 After seeing the information here, do you feel like you would want to do anything differently with these websites?

Most Contacted Domains

- I30 Does the information from this section make you feel differently about how you interact with your phone?

- I31 How would you summarize this section to someone else?

- I32 What would you change about this section, if anything?

- I33 After seeing the information here, do you feel like you would want to do anything differently with apps or websites on your phone?

Exit Questions

- I34 Would you ever look at the App Privacy Report again?

[I35 and I36 is asked if the participant responded "Yes" to I34]

- I35 In what context would you look at the App Privacy Report again?

- I36 How often would you look at the App Privacy Report and why?

[I37 is asked if the participant responded "No" to I34]

- I37 Could you explain why you would not look at the App Privacy Report again?

- I38 Are any of the [privacy risks described in S4, S5, and S6] addressed by the App Privacy Report?

C. Follow-Up Survey

- Q1 Is your iPhone's App Privacy Report currently turned on?

- o Yes
- o No
- o I am not sure

- Q2 Have you opened and looked at the report since we interviewed you in November/December 2023?

- o Yes
- o No

[Q3 through Q6 are shown if Q2 is "Yes"]

- Q3 On average, how often did you check the report?

- o Every day
- o 2–6 times a week
- o Once a week
- o Once every two weeks
- o Once a month

- Q4 What motivated you to check the report? [free response]

- Q5 When you checked the report after the interview, what information did you find useful? [free response]

- Q6 Have you taken further actions based on information you learned in the report (e.g., changing settings, deleting apps, looking up domain names, changing browsing habits, etc.)? [free response]

[Q7 is shown if participants' answer to Q2 is inconsistent with answer to I34]

- Q7 You previously said that you [would/would not] look at the report again. What changed your mind? [free response]

D. Demographics of Screening Survey

Table VI includes demographics information of participants from the screening survey.

E. Interview Codes

Table V includes details about participants reactions to different sections of the privacy report relevant to RQ1. Table VII includes codes we used that reference specific elements of the privacy report. Table VIII includes codes we applied to answer each of the research questions and one example quote for each of the codes.

TABLE IV: Internet Users' Information Privacy Concern (IUIPC-8)

	Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly agree
Consumer online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared.	○	○	○	○	○
Consumer control of personal information lies at the heart of consumer privacy.	○	○	○	○	○
Companies seeking information online should disclose the way the data are collected, processed, and used.	○	○	○	○	○
A good consumer online privacy policy should have a clear and conspicuous disclosure.	○	○	○	○	○
It usually bothers me when online companies ask me for personal information.	○	○	○	○	○
When online companies ask me for personal information, I sometimes think twice before providing it.	○	○	○	○	○
It bothers me to give personal information to so many online companies.	○	○	○	○	○
I'm concerned that online companies are collecting too much personal information about me.	○	○	○	○	○

TABLE V: RQ1 and report feedback codes and the amount of participants for whom at least one instance of that code was recorded, in the Data and Sensor Access section (out of 19 participants) and Network Activity sections (out of 20 participants)

RQ1: Understanding			
Grouping or Description	Code	Data & Sensors	Network Activity
Accurate understanding	Identification	■■■■■	■■■■■
Asked a question	Question	■■■■■	■■■■■
Inaccurate understanding	Misunderstanding	■■■■■	■■■■■
Report Feedback			
Positive feedback	Generally helpful/informative	■■■■■	■■■■■
	Helpful to inform specific privacy action or use case	■■■■■	■■■■■
	Well organized/formatted/designed	■■■■■	■■■■■
Criticism/suggestions	Organize existing information differently	■■■■■	■■■■■
	Interpret/explain existing information	■■■■■	■■■■■
	Add new information	■■■■■	■■■■■
	Add function/feature	■■■■■	■■■■■
■■■■■ = a few; ■■■■ = some; ■■■■ = about half; ■■■■ = most; ■■■■ = almost all			

TABLE VI: Demographics of screening survey participants

	Screening Survey No.	%
Gender		
Female	91	47.9
Male	96	50.5
Prefer not to say	1	0.5
Other	2	1.1
Age		
18 - 24	39	20.5
25 - 34	71	37.4
35 - 44	39	20.5
45 - 54	21	11.1
55 - 64	14	7.4
65 - 74	6	3.1
Highest Degree of Education		
High school or equivalent	13	6.8
Some college, no degree	35	18.4
Associate degree	20	10.5
Bachelor's degree	87	45.8
Master's degree	28	14.7
Doctorate or professional degree	7	3.7
Total	190	100
IUIPC	Avg.	S.D.
Control	5.89	1.04
Awareness	6.47	0.82
Collection	5.80	1.24

TABLE VII: Codes for elements of the privacy report that participants had reactions to

Data & Sensor Access
General data/sensor access
Data/sensor access: time
Data/sensor access: type
Data/sensor access: functionality
Overpermissioning
Network Activity
Order of domains
Significance of "bar charts"
Personal app usage vs. perceived volume of domain activity
Volume of domains
Domain owner
Interpretable domain name
Mysterious domain name
"Unknown Domains" label
Domain contact: data being shared
Domain contact: implications
Domain contact: purpose
Domain contact: time of occurrence
Cross-app tracking
Functionality
Profile building
Targeted advertising
Third-party data collection
Privacy tradeoff
Resource usage
Sensitive app

TABLE VIII: Codes applied to interview responses with example quotes

<i>Description or Grouping</i>	<i>Code</i>	<i>Quotes from participants</i>
RQ1: Understanding		
Accurate understanding	Observation	"It shows all the apps that you've used, and then it shows what those apps access" —P46
Asked a question	Question	"Was it for a specific website? Was it just in the background? Was it because I used another app that somehow related to Safari?" —P159
Inaccurate understanding	Misunderstanding	"It's just collecting your most frequently used domains by app." —P52 (participant misunderstood the difference between their own use of an app and the amount of domain activity an app engages in)
RQ2: Attitudes		
Positive attitudes	Informed	"[Data and Sensor Access] allows me to understand just how many of these apps have access to your information" —P28
	Reassured	"It almost makes me feel more comfortable" —P79
Neutral attitudes	Not bothered	"Most of these things are perfectly reasonable." —P99
	Unsurprised	"Nothing stands out as glaringly incorrect or surprising." —P159
Negative attitudes	Curious	"I'm a little curious as to why it would need to access my photos." —P76
	Strange	"It's so weird for Amazon to contact my sleep app" —P113
	Surprised	"I was on [this game app] yesterday, but...I didn't know it was accessing my location." —P46
	Resigned	"I don't honestly think I'll be using the apps any differently, given that I feel that my data...is already out there." —P34
	Confused	"I have no idea what any of these [domains] do" —P46
	Concerned	"This (App Network Activities) concerns me...I can't figure out what they are." —P54
RQ3: Intent to Change Behavior		
App and browsing related	Delete app	"I would one hundred percent uninstall any app that has permissions I don't expect it to have." —P18
	Find alternative app/website	"I would take the precaution and try to avoid them, or look for other sites where I can find similar information." —P62
	Change app permissions	"I would want to see what I have [the permissions] currently set at and probably restrict them a little bit more" —P148
	Decrease app/website usage	"[I would] use [these websites] less...just across the board" —P28
	Change browsing behavior	"It helps me be more aware to not click on things that I shouldn't." —P76
Information seeking	Consider additional privacy tools	"It makes me wonder if I should be using a web browser...that does have an ad blocker." —P109
	Look for more information elsewhere	"I definitely look them up to see if they are related to the functionality, or if it's more advertising based" —P159
	Check report in future	"Anytime I download an app and use it, I will probably come in [the report] to check... what's being accessed." —P159
Nonspecific intentions	Check phone settings	"I may try to...go through some of the settings...to see if there's any sort of privacy changes...that I can make." —P34
	General prevention of unwanted activity	"I certainly wanna be more careful." —P76
	Desire to act but unsure of what can be done	"That [information] just makes you worry, but...I don't know what to do about it." —P106
Reasons for not wanting to change behavior	No behavior change (no reasons given)	"I'm concerned with it, but I'm not going to act an awful lot different." —P98
	App/website functionality is necessary	"These are [apps] that I use for work...or just things that I can't really avoid." —P46
	Not sufficiently concerned	"Since there's nothing that I'm overly concerned about, I don't think I'm...gonna go back and look at [the report] again." —P97
	Accepting of privacy tradeoff	"You're sort of trapped in a way, if you want to use the free service, that's the cost of doing business" —P54
	Insufficient information to justify action	"It's not enough information for me to decide whether I should keep the app or not." —P34
	No effective actions seem to exist	"It feels like there's very little that I can do about having these domains be contacted by the apps" —P34
Report Feedback		
Positive feedback	Generally helpful/informative	"It's a really useful feature to check out, because it displays everything in a uniform way" —P18
	Helpful to inform specific privacy action or use case	"I think this is this is even more helpful...for external apps, so you can make sure that nothing is being violated" —P52
Criticism/suggestions	Well organized/designed	"I like how it's broken down by apps and websites that connect to the same domain...I like how it's grouped." —P34
	Organize existing information differently	"This is so much information...this section could be broken up into further sections" —P46
	Interpret/explain existing information	"I just would like to see some of this explained that in a better way people like me could understand it." —P28
	Add new information	"I would want to add why exactly various apps are being accessed" —P71
	Add function/feature	"It would be cool if they had a [domain] search feature" —P113