

Your Location has been Shared 5,398 Times!

A Field Study on Mobile App Privacy Nudging

Hazim Almuhiemedi* Florian Schaub* Norman Sadeh*
Idris Adjerid• Alessandro Acquisti† Joshua Gluck*
Lorrie Cranor* Yuvraj Agarwal*

December 2014
CMU-ISR-14-116

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Note: This paper has been accepted for publication in the proceedings of CHI 2015. The present version may undergo some minor editing prior to publication in the CHI proceedings.

*School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

•Mendoza College of Business, University of Notre Dame, Notre Dame, IN, USA

†Heinz College, Carnegie Mellon University, Pittsburgh, PA, USA

This research was supported in part by the National Science Foundation under grants CNS 10-1012763 (Nudging Users Towards Privacy) and CNS 13-30596 (Towards Effective Web Privacy Notice & Choice: A Multi-Disciplinary Perspective), in part by Google, in part by Samsung, and in part by King Abdulaziz City for Science and Technology.

Keywords: Mobile, Privacy, Privacy Decision Making, Privacy Nudges

Abstract

Smartphone users are often unaware of the data collected by apps running on their devices. We report on a study that evaluates the benefits of giving users an app permission manager and of sending them nudges intended to raise their awareness of the data collected by their apps. Our study provides both qualitative and quantitative evidence that these approaches are complementary and can each play a significant role in empowering users to more effectively control their privacy. For instance, even after a week with access to the permission manager, participants benefited from nudges showing them how often some of their sensitive data was being accessed by apps, with 95% of participants reassessing their permissions, and 58% of them further restricting some of their permissions. We discuss how participants interacted both with the permission manager and the privacy nudges, analyze the effectiveness of both solutions and derive some recommendations.

1 Introduction

Previous studies have shown that smartphone users are often unaware of the data collected by their apps and express surprise and discomfort when they find out (e.g. [15, 27, 31, 12, 14]). Recently, privacy managers, such as AppOps (introduced in Android 4.3 but later removed with the introduction of Android 4.4.2), privacy controls in iOS, or ProtectMyPrivacy [6], have emerged that offer users more fine-grained control over their privacy. However, to the best of our knowledge, the effectiveness of these fine-grained controls has not yet been evaluated. Privacy decision making is known to be subject to cognitive and behavioral biases, and decision heuristics that often lead to privacy-adverse decisions in favor of short-term benefits [4]. Privacy nudges have been proposed to support users in their privacy decision making [3]. Such nudges can help make privacy risks more salient and help users move towards privacy settings that better align with their privacy expectations and concerns. Accordingly, a related question is to what extent users who have access to fine permission controls might benefit from nudges that raise their awareness of the data collected by their apps and prompts them to review their current permission settings.

In this work, we focus on two research questions: (1) Is access to a fine-grained app permission manager an effective way of helping users review and modify their app permissions? (2) Do privacy nudges regularly alerting users about sensitive data collected by their apps an effective way of possibly enhancing the effectiveness of a fine-grained app permission manager? To address these two questions, we conducted a 22-day field study in which 23 participants interacted with a permission manager – AppOps on Android – for one week, followed by an 8-day phase in which the permission manager was supplemented with privacy nudges tailored to a participant’s installed apps and their data access behavior.

Our mixed methods approach provides rich insights into (1) how and why participants review and restrict app permissions with a permission manager and (2) how they react to privacy nudges alerting them about the data collected by their apps. Our results also confirm that users are generally unaware of mobile app data collection practices. They demonstrate the benefits of fine-grained permission managers and show that the effectiveness of these managers can be significantly enhanced by the delivery of nudges intended to further increase user awareness about mobile app data collection practices.

1.1 Contributions

Our results show that permission managers – albeit not widely available yet – are an essential tool for users to manage their privacy. The privacy nudges successfully prompted almost all participants to further review their permissions, and triggered 58% of them to make restrictive app permission changes.

Our results suggest that participants gained a better understanding of ongoing data access practices. The nudges not only resulted in users restricting permissions associated with the specific apps and permissions they explicitly mentioned but also led users to review and modify settings associated with other apps and other permissions. This is in part attributed to the way the AppOps permission manager organizes information on an app-by-app basis. Based on our results, we derive recommendations for the design of effective privacy nudges on mobile devices.

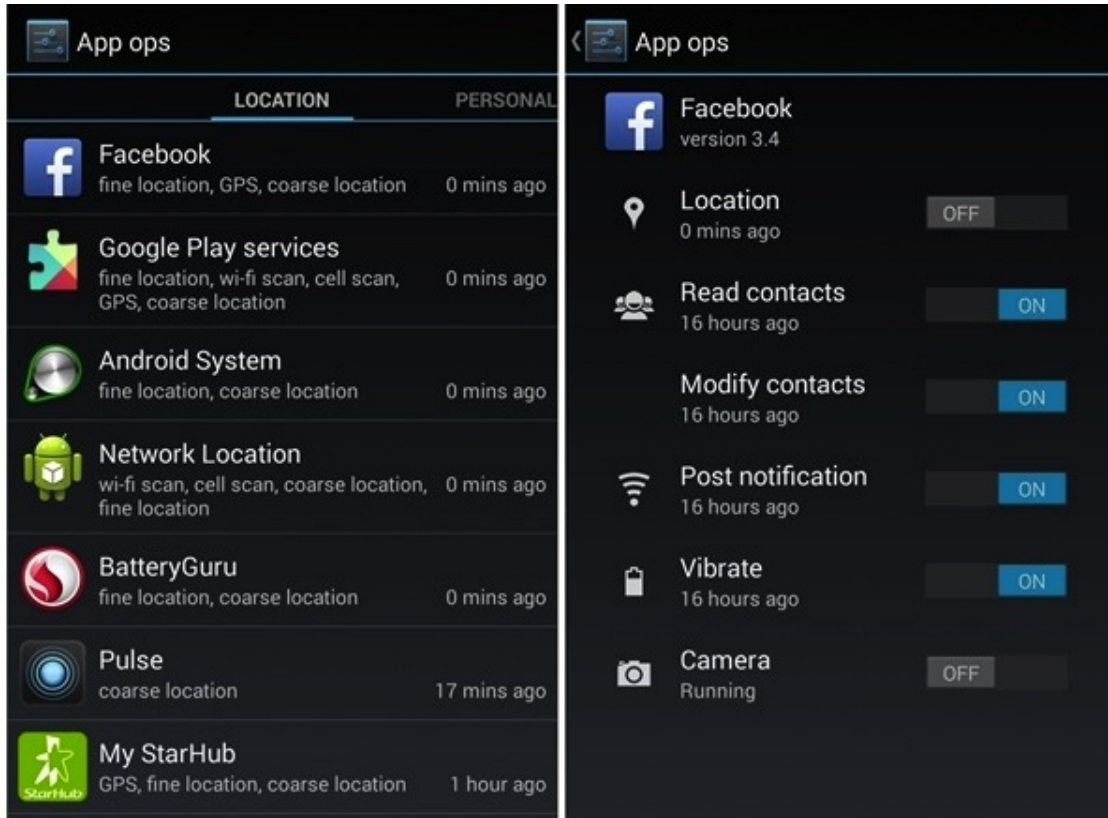


Figure 1: AppOps has different tabs that list all apps that accessed respective data, e.g., location (*left*). Selecting a specific app opens a screen showing all permissions accessed by this app (*right*).

2 Background and Related Work

2.1 Privacy Controls in Mobile Platforms

Android and iOS provide some privacy safeguards for users. Android displays an app's requested permissions at *install time*, which users have to accept to install it. In Android 4.3, Google included a permission manager, called *AppOps*, that was hidden by default and required an external app to access. AppOps allowed users to selectively grant or restrict permissions for installed apps. AppOps, shown in Figure 1, organized permissions into four categories: location, personal data (e.g. calendar, phone contacts), messaging (e.g. SMS), and device features (e.g. vibration, notification). In each tab, apps are ordered by most recent access. When selecting a specific app in an AppOps tab, users are shown all permissions for that app. As of Android 4.4.2 AppOps has been made inaccessible [10], unless the device is rooted.

Apple also introduced fine-grained *access time* permissions on a per-feature basis, such as location and contacts, in iOS 5 [34] and expanded it to other features (e.g., camera, calendar, microphone, HealthKit) in iOS 7 and iOS 8. Importantly, in iOS access is denied by default until the user explicitly grants it. Furthermore, the iOS privacy settings panel groups apps by data type

(e.g., location) rather than showing all permissions for a particular app in one place.

2.2 Enhancements to Privacy Controls in Mobile Platforms

Several privacy enhancements have been proposed for Android [20, 22, 30] and iOS [6, 11]. Apex [30] retrofits Android to enable selectively granting, denying or constraining access to specific permissions on a per-app basis. AppFence [20] uses taint tracking [13] to provide shadow data to untrusted apps and to block data from leaving the device. ProtectMyPrivacy (PMP) enables iOS jailbroken users to selectively grant or deny apps access to user information, or provide fake information to an app [6]. PMP also provides privacy setting recommendations for new apps to users via crowdsourcing. Even finer grained privacy recommendations can be provided by clustering such privacy decisions into user preference profiles [29, 28].

In contrast to the discussed approaches, we are agnostic of the mechanism for providing fine-grained privacy controls.

2.3 Asymmetric Information and Privacy Nudges

In privacy contexts, asymmetric or incomplete information refers to a disparity between users' and service providers' knowledge of collection, use, sharing practices, potential consequences, and available protections concerning users' personal information. This phenomenon has been attributed to ineffective communication of privacy risks and protections to users in privacy policies and notices [21]. More recently, researchers have identified a similar information asymmetry, in which users are unaware of the data collection performed by mobile apps, in the context of mobile privacy [15, 31]. This led to alternate proposals for presenting privacy risks to users in a manner that is more readable and salient [25]. However, research focused on privacy contexts finds that information disclosures led to fleeting and even perverse effects on behavior, casting doubts on one time information disclosures' ability to yield better user privacy decision making. For example, Adjerid et al. [5] showed that the impact of simple and readable notices can be thwarted by a mere 15 second delay between showing privacy-relevant information and privacy choices. In light of these limitations, scholars are increasingly turning to alternate methods for communicating relevant privacy information to users, such as privacy nudges [3].

Nudges are "soft-paternalistic" behavioral interventions that do not restrict choice, but attempt to account for bounded rationality in decision making [33]. Within privacy and security contexts, nudges may ameliorate some of the inconsistency in user decision making, such as the dissonance between users stated privacy concerns and actual observed behavior [3, 32]. So far, few works have evaluated how privacy nudges can differentially impact user behavior. Wang et al. found that privacy nudges can effectively influence privacy concerns and behavior on Facebook [35].

In the mobile context, the potential for nudges to support privacy decision making is appealing and may include notifications that, in contrast to traditional notices, highlight the recipients, contexts, or type of personal information being shared via a mobile device [8]. However, most related work has focused on supporting the app installation process by nudging users towards less privacy-invasive apps [9, 18, 26]. Harbach et al. [19] enrich permission dialogs with personalized examples from the user's device to make risks more salient (e.g., showing a personal photo

for gallery access). Rather than focusing on app installation, we analyze behavior of installed apps in order to increase awareness of unexpected invasive behavior and surreptitious data access. Balebako et al. [7] proposed mobile privacy nudges based on apps' access frequency to specific data and evaluated them in a lab study. Although this work extends Balebako's, it differs in three dimensions. First, this work measures the effect of privacy nudges in triggering users to review and adjust their app permissions, whereas Balebako measured perception and feeling. Second, our study evaluates privacy nudges in situ, with participants using their own devices 'in the wild', whereas Balebako's study was a lab study and participants did not use their own devices. Third, our nudges show frequency for all installed apps on participants' devices, whereas Balebako's only tested the nudges for two game apps. Fu et al. [17] proposed a per app run-time location access notification and contrasted it with the existing Android's location access disclosure method. They showed both that their method is more effective than Android's and showed anecdotal evidence of how their notification affected participants' behavior (e.g. stopped using intrusive apps). Our work extends Fu's but differs in two dimensions. First, our work measures the effect of privacy nudges in triggering users to review and adjust their app permissions, whereas Fu's focused only on transparency. Second, their work focused on location access, whereas we also examine phone contacts, calendar, and call logs. Finally, Fisher et al. [16] asked 300 iOS users to take screenshots of their location privacy settings, examined whether users permitted or restricted access to their location, and how to use these decisions to predict future privacy decisions. Our work extends this, but differs in three dimensions. First, their work focused on location access, whereas we also examine phone contacts, calendar, and call logs. Second, our work combines quantitative and qualitative data to understand why users use permission managers and for what, whereas their work only examined users' decisions at one point-of time. Additionally our work evaluates whether privacy nudges can improve the effectiveness of permission managers, whereas theirs focused on predicting future privacy decisions by users.

3 Mobile Privacy Nudge Design

We designed a mobile privacy nudge that provides *concise privacy-relevant information* and *meaningful actions* that reduce the threshold for users to act upon the nudge's content.

3.1 Nudge Content

To be effective, a nudge must garner user attention. Prior work has shown that users are unaware of, and surprised by, apps' data access practices and frequency [7, 17, 23], which suggests nudging users to review their app permissions has utility. Therefore, we designed the mobile privacy nudge shown in Figure 2 to display a succinct message describing the number of apps accessing one information type and the total number of accesses in a given period.

The nudge further lists three specific apps that accessed the information in the given period, to concretize the otherwise abstract information. In order to avoid showing only expected apps (e.g., mapping and navigation apps accessing location), the three displayed apps are selected randomly from apps that accessed that information type. The addition of "and 10 other apps" aims to pique

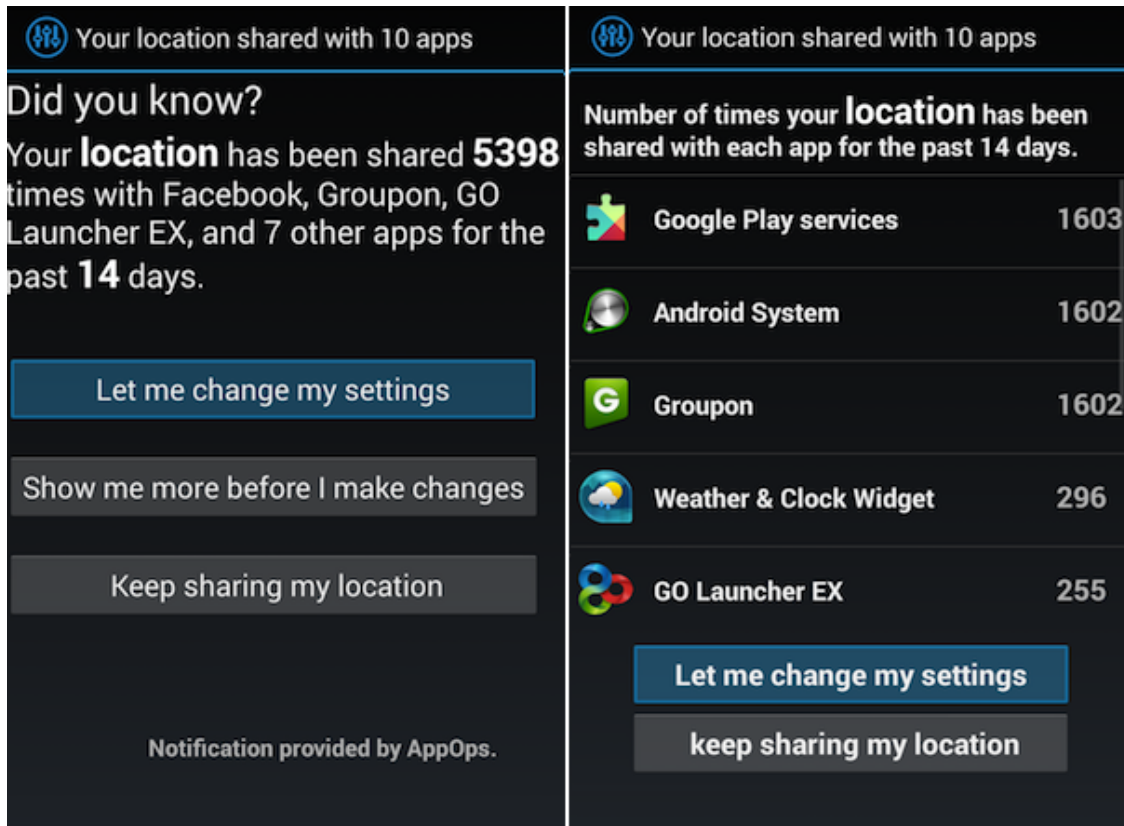


Figure 2: Screenshot of a privacy nudge for location (left) shown to a participant during our study. Nudge content is based on participant’s apps and their access to location information. “Let me change my settings” opens AppOps. “Show me more before I make changes” opens the detailed report (right). “Keep Sharing my location” closes the nudge.

the user’s interest, and trigger them to review permission settings for the particular information type.

In order to enhance the nudge’s credibility, we included cues that establish a relation between the nudge notification and the installed privacy manager (AppOps in our case), such as the AppOps icon and a tag line at the bottom (see Figure 1).

3.2 Nudge Response Options

The privacy nudge provides targeted response options to facilitate privacy management (see Figure 2). The “Let me change my settings” option opens AppOps directly. We hypothesize that facilitating access to the permission manager may lead users to review and adjust additional permissions once they switch their focus to privacy management. Since we want nudge users to select this option, it is highlighted.

The second option (“Show me more before I make changes”) opens a detailed report, shown in Figure 2, which lists each app’s access frequency for the nudge’s particular information type,

in descending order. The goal of the detailed report is to enable users to investigate which apps accessed the particular information in order to support them in comparing their expectations with apps’ data practices. Rather than just naming the option “show me more information,” we intentionally indicated that users will also be able to make changes through this option, and imply that this information may help their decision. The detailed report replicates the nudge’s other response options to make the provided information actionable. Prior work inspired the detailed report design [7].

The third option (“Keep sharing my [data]”) allows users to indicate the status quo is acceptable. Keller et al. [24] recommend employing enhanced active choice to emphasize the desired option (option 1) by “highlighting the losses incumbent in the non-preferred alternative.” Therefore, option 3 is adapted to the specific information (i.e., [data] is replaced with “location”). Finally, users can also *ignore* the nudge by pressing the “Home” button or by switching to a another app.

4 Methodology

We conducted a field study to gain insights on the effect and perceived utility of mobile privacy managers, as well as the effect and perception of privacy nudging. We implemented our privacy nudges on Android since it supported a permission manager, AppOps, which is accessible on regular non-rooted devices. Thus, our 23 participants were able to use their own phones. Our study received IRB approval.

4.1 Implementation of Study Client

We implemented a study client app and installed it on participants’ phones. The study client acted as a launcher for AppOps, which is otherwise inaccessible. Our study client collected information about app permissions accesses for specific information types, which was used to generate personalized privacy nudges for each participant’s phone. The required information was obtained by periodically recording logs created by AppOps. The AppOps logs show for each app-permission pair the last time the app tried to access the permission. The log shows when access to a permission was rejected (e.g., after the user restricted an app’s access). If the app is currently using a permission, the logs show how long the app has been accessing it. By capturing this information in five minute intervals, we gained detailed insights about apps accessing permissions, as well as the progression of permission changes made by participants via AppOps. Accessing the AppOps logs requires a one-time runtime permission (“GET_APP_OPS_STATS”), which can only be granted if the device is connected via USB after app installation.

In addition to recording access frequency and permission changes, our study client recorded participants opening AppOps, as well as their interaction with displayed privacy nudges. Permission changes had to be recorded periodically, since AppOps does not provide access to specific interaction events and modifying AppOps would have required rooting and flashing participants’ devices, which we deemed unacceptable. Hence, we used the time difference between a participant’s recorded response to a privacy nudge and an observed permission adjustment to infer whether it was triggered by the respective nudge.

4.2 Study Procedure

Our field study consisted of an entry session, three consecutive field phases lasting 22 days in total, an exit survey, and an optional exit interview. We opted for a within-subjects design as we were interested in observing phone and app usage without interventions in order to establish a baseline, as well as observing interaction with a permission manager with and without supporting privacy nudges.

Entry session: We invited participants to our lab to read and sign our consent form. Because AppOps was only available on Android versions 4.3–4.4.1, participants were required to initial that they would not update to Android 4.4.2 during the study, and could be disqualified otherwise.

Next, participants completed an online entry survey on a provided computer. The survey asked about general Android usage (e.g. frequently used apps, reasons for installing or uninstalling apps), mobile privacy and security attitudes and behaviors (e.g., screen lock use, phone encryption, awareness of apps’ permissions), and demographic questions (e.g. gender, age, phone model). While the participant completed the survey, we installed the study client on her phone with the required runtime permission to access AppOps logs, and placed it in a folder named ‘Android Apps Behaviors,’ to make it easily locatable.

Phase 1: Baseline: For the first 7 days of the study, our study client collected data about the participant’s installed apps and their data access behavior, without providing access to AppOps or showing privacy nudges. The information collected served as a baseline to better understand participants’ phone and app use, and also informed the generation of privacy nudges in phase 3.

Phase 2: AppOps Only: On the first day of the second phase, we made AppOps available through the study client and sent an Email and an SMS to participants introducing it. The message subject was “AppOps is now available to you” and it read “AppOps is an app which allows you to selectively grant/deny apps access to your personal information (e.g. location, phone contacts, calendar, SMS messages, etc) on your phone. We just made this app available to you. To use it, go to ‘Android Apps Behaviors’ folder then click on AppOps.” This notification acted as a weak privacy nudge, comparable to seeing a media article or an ad about AppOps. Participants did not receive any further interventions during phase 2, which lasted 7 days.

Phase 3: AppOps Plus Privacy Nudges: In phase 3, which lasted 8 days, participants additionally received one privacy nudge per day, sent at a random time from 11am to 8pm. In our study, we provided nudges for four information types: location, phone contacts, calendar, or call logs. They were selected both because they were shown to be the subject of mobile users’ privacy concerns [14], and because initial experiments demonstrated that these four information types constituted the most requested resources by apps, which made it likely that participants’ would have apps installed that actually accessed these information types. On the first four days of phase 3, all four nudges were shown in a random order to avoid order effect. The same nudges were then repeated in the same order on the last four days of phase 3. The first set of privacy nudges showed access statistics since the beginning of the study (i.e., 14-18 days), the second set showed access statistics for the period since the previous nudge for that data type (i.e., 4 days). If no installed apps had accessed the information type of a scheduled nudge in the respective time period, the next nudge would be shown instead.

Exit Survey and Interviews: After completing phase 3, participants were sent a link to an online

exit survey. The survey focused on the participant's experience with AppOps (e.g., prior use of AppOps, AppOps use during study, reasons for using AppOps) and the participant's understanding of, and experience with, the privacy nudges (e.g., meaning of nudge text and options, provided privacy awareness and decision support). Upon completion of the exit survey, participants were compensated with a \$30 Amazon gift card.

All participants were further invited to an optional semi-structured one-hour interview and compensated with an additional \$10 Amazon gift card. Eight participants responded and were interviewed. The interviews served to gain deeper qualitative insights to participants' experiences with AppOps and the privacy nudges. The interviews were partially tailored to a participant's behavior during the study. For example, we inquired participants' reasons for specific permission changes they made. We further presented them with specific nudges displayed to them when asking about their experiences.

4.3 Recruitment

We conducted our study from May to July 2014. Participants were recruited via Craigslist and from a city-wide participant pool maintained by our university. Ads directed prospective participants to a screening survey. Twenty-six respondents, meeting the following criteria, were invited to participate in the study: (1) Adults who have Android phones running Android version 4.3–4.4 (because AppOps is only supported by these Android versions); (2) have a mobile data plan with at least 2 GB/month (as data would have to be transferred during the study); (3) able to visit our lab for the entry session. Three were later disqualified as they upgraded to Android 4.4.2 during the study, which made AppOps inaccessible.

4.4 Limitations

Conducting a field study enabled us to evaluate our privacy nudges in-situ on participants' own devices. This increased ecological validity but introduced multiple challenges. First, we were unable to recruit a larger number of participants, because carriers (e.g. AT&T) and OEMs (e.g. Samsung) rolled out updates to Android 4.4.2 (i.e AppOps removed) around the same time, which significantly shrunk the pool of potential participants. Second, our study client required Internet connectivity. However, some participants deactivated data connection to conserve data volume or had intermittent connectivity for other reasons. This affected our data collection and caused some nudges to be lost or delivered later than scheduled. We implemented a monitoring tool to identify participants whose devices were not sending back information regularly, and then reminded them via email to remain connected. While this worked, we used this approach sparsely to avoid biasing participants' responses. Due to these technical difficulties, some participants received all eight nudges while others received fewer nudges. These limitations may have skewed our results to be more conservative, but we are confident that they did not undermine our findings in general.

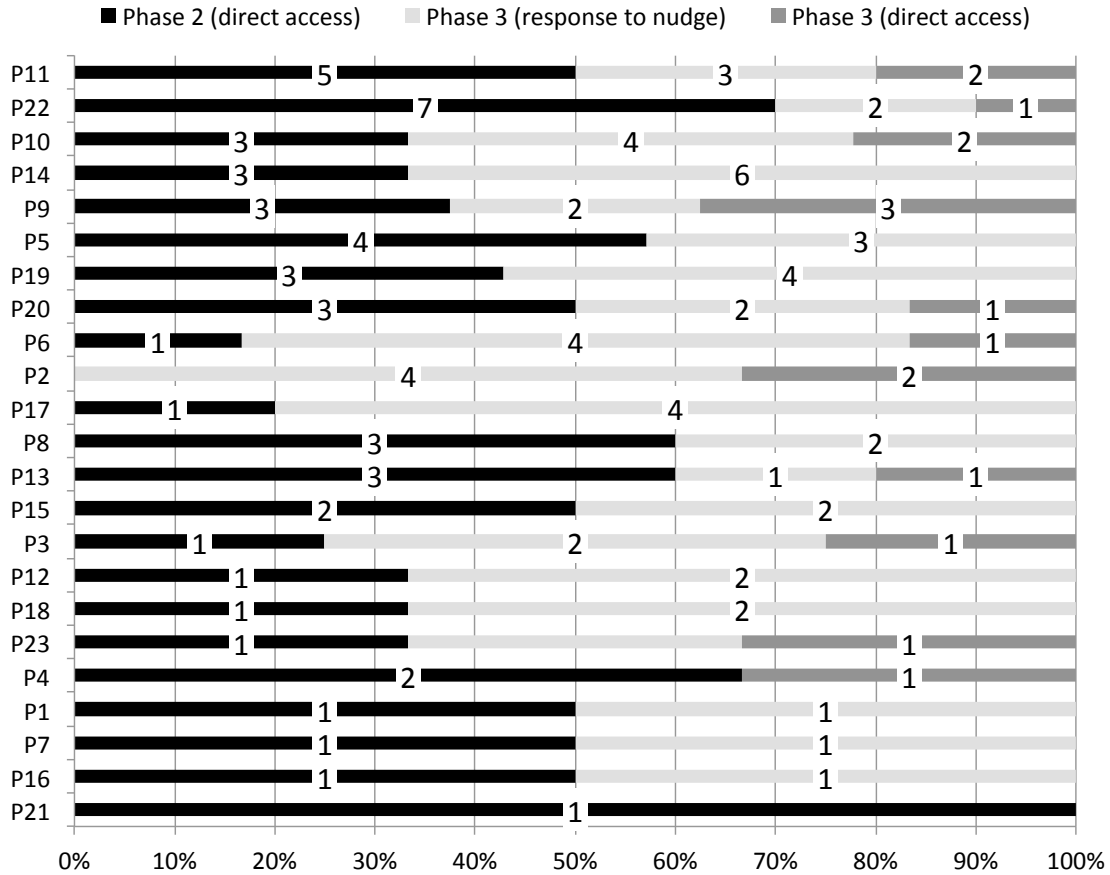


Figure 3: The number of times each participant reviewed their apps’ permissions by opening AppOps in phase 2 & 3.

5 Results

We first describe participant demographics and apps usage. Then, we report how participants interacted with the permission manager alone followed by their interaction with the permission manager with accompanying privacy nudges. Finally, we report how participants interacted with nudges and evaluate the effectiveness of the nudge’s components.

5.1 Demographics

In total, we had 23 participants (65% female; ages 18–44, median=23), of whom 21 owned Samsung devices and 2 owned an HTC One. Based on data collected in phase 1, participants had 89 apps installed on average (SD=22), including services and pre-installed apps. Twenty-one participants (91%) reported never using AppOps before; one had used AppOps, and one was unsure. Moreover, the data collected in phase 1 showed that participants could not access AppOps (e.g. not other launcher app for AppOps installed), until phase 2.

In the following, we use two main variables in our analysis:

(1) *Reviewing apps' permissions* represents how often a participant opened AppOps to review their app permissions regardless of whether they adjusted app permissions.

(2) *Adjusting apps' permissions* represents how often participants adjusted their app permissions by calculating (a) *restrictive adjustments*, i.e., how often participants restricted any app the access to a permission, and (b) *permissive adjustments*, i.e., how often participants permitted an app access to a restricted permission.

5.2 Effectiveness of AppOps Without Privacy Nudges

In phase 2, after we made AppOps available, participants reviewed their app permissions 51 times, restricted 76 distinct apps from accessing 272 permissions, and permitted access to one restricted permission.

Reviewing apps' permissions: As Figure 3 shows, 22 participants (95.6%) reviewed their permissions at least once. Of those, 12 participants reviewed their apps' permissions multiple times. P2 did not review his permissions in phase 2.

Adjusting apps' permissions: As shown in Figure 4, 15 (65%) participants restricted 272 app-permission pairs from 76 distinct apps, including both participant-installed and pre-installed apps, see Figure 5. Breaking down restrictions by information type, participants restricted apps' access to location 74 times (27%), contacts 57 times (21%), calendar 10 times (4%), and call logs 9 times (3%). Other restricted permissions included: camera 42 (9%), SMS 21 (8%), post notification 19 (7%), and recording audio 15 (6%). Only P10 made a permissive adjustment by permitting the Weather Channel app to send notifications.

In the exist survey, we asked participants if they used AppOps, what they used it for, and why. Most participants reported that they used AppOps to review their app permissions and adjust them if needed. For example, P9 responded: “[I used AppOps] to see what personal information different apps had access to and change that” because “I didn't like that too many apps could access so much information.”

In the interviews, participants further explained why they restricted apps' access to permissions. First, participants restricted unused apps, especially pre-installed apps. P10 stated: “I also blocked bunch of AT&T bloatware from accessing any information. I don't use them anyways.” Second, participants restricted permissions required for unused functionality. P13 restricted iHeartRadio access to location, explaining: “I know what stations I want listen to no matter where I'm so I turn off the location.” Third, Participants restricted apps when the purpose to access their personal information is unclear. P4 stated: “[I turned it off] because I can't think of a reason why Inkpad needs my location.”

Making the permission manager available to participants led them to actively review their app permissions and adjust them as needed. This indicates clearly that participants wanted to exercise control over apps' access to personal information.

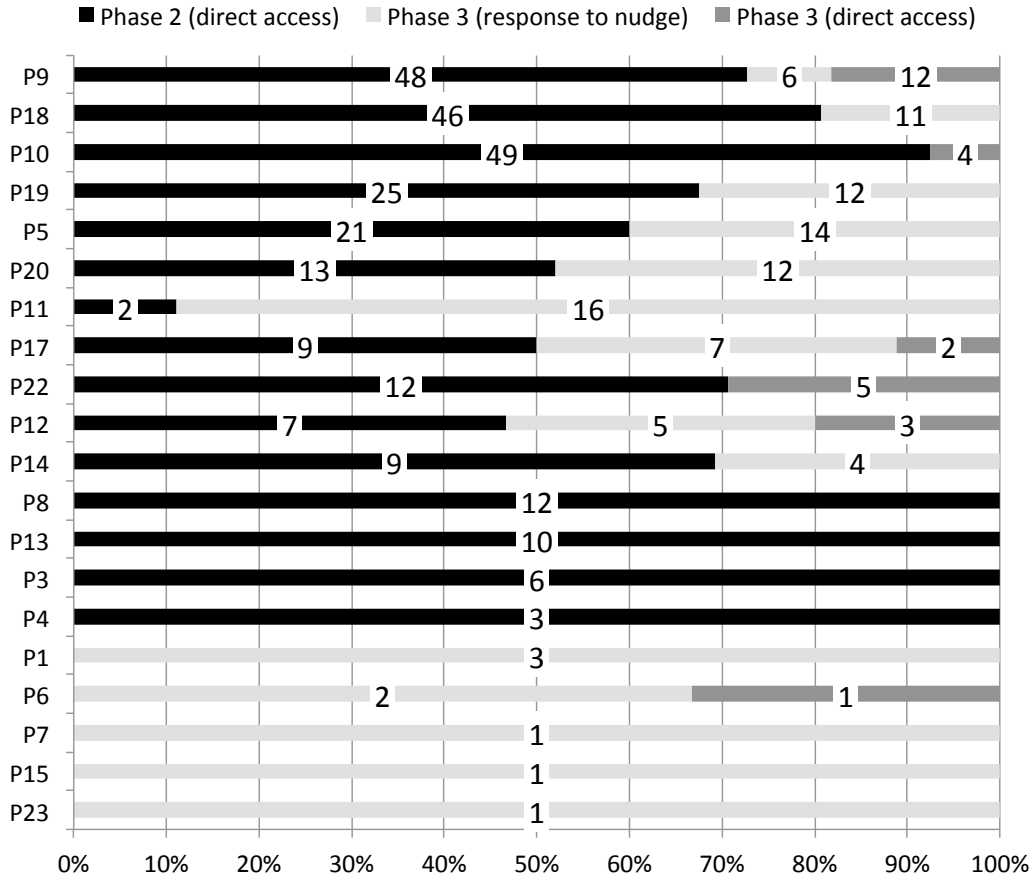


Figure 4: Number of restricted permissions for each participant in both phases.

5.3 Effectiveness of AppOps with Privacy Nudges

The goal of our nudges was to get users to review their apps permissions and adjust them as needed. To that end, we designed the nudges to complement and increase the effectiveness of AppOps. It is important to note that participants' phase 3 behavior is contingent on their phase 2 behavior. For instance, if a participant restricted access to some permissions in phase 2, these restrictions hold in phase 3, and may not require further review or adjustment. Hence, we report and analyze results from phase 3 relative to phase 2.

In phase 3, participants reviewed their apps' permissions 69 times, restricted 47 distinct apps from accessing 122 permissions, and permitted six apps access to six permissions.

Reviewing apps' permissions: As Figure 3 shows, 22 participants (95.6%), with the exception of P21, reviewed their apps' permissions at least once in phase 3. Participants could review their apps' permissions either by opening AppOps directly (same as in phase 2), or by opening AppOps in response to a nudge. Twenty-One participants reviewed their apps' permissions 53 times (78%) in response to the nudge, and 15 times (22%) by directly opening AppOps. P4 reviewed her apps' permissions only once by opening AppOps directly. Thus, the privacy nudges were the primary

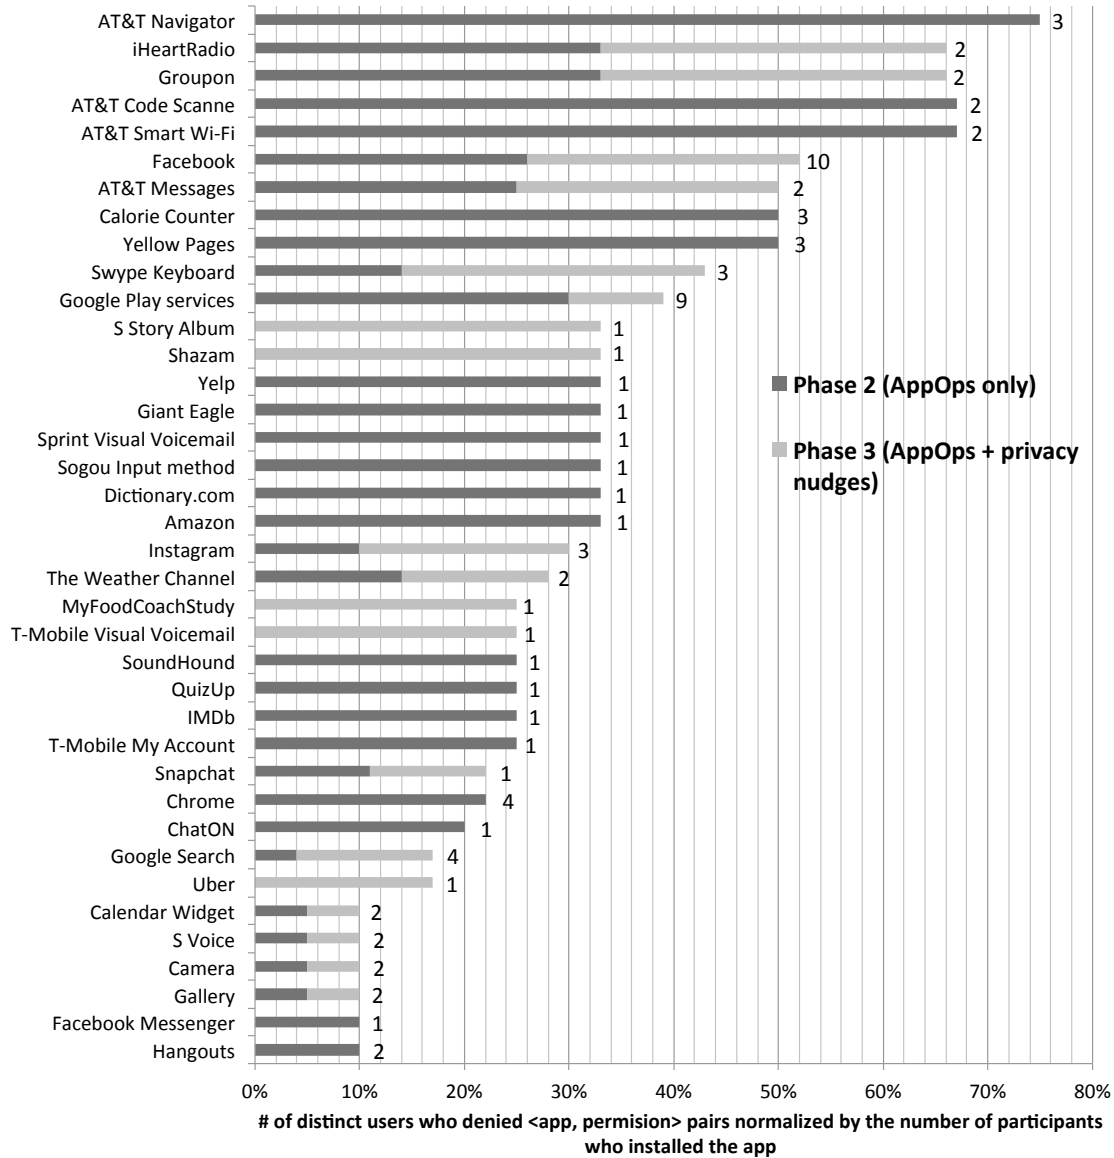


Figure 5: The apps that were revoked permissions in phase 2 & 3. The numbers to the right of the bars are the absolute number of distinct users who revoked at least one permission per app.

trigger for participants to review their apps' permissions.

Figure 6 shows that participants' interaction with AppOps declined sharply after day three: 39% (day 3), 13%, 17%, 22%, 4%, and 9%, respectively. However, the privacy nudges introduced in phase 3 positively affected participants' interest in reviewing their apps' permissions (cf. Figure 3). For instance, P2 did not review his apps' permissions in phase 2 but privacy nudges triggered him to do so six times in phase 3.

Adjusting apps' permissions: Figure 4 shows that 16 (70%) participants restricted 122 app-permission pairs, 14 in direct response to nudges. Ninety-Five (78%) of the restrictive app permis-

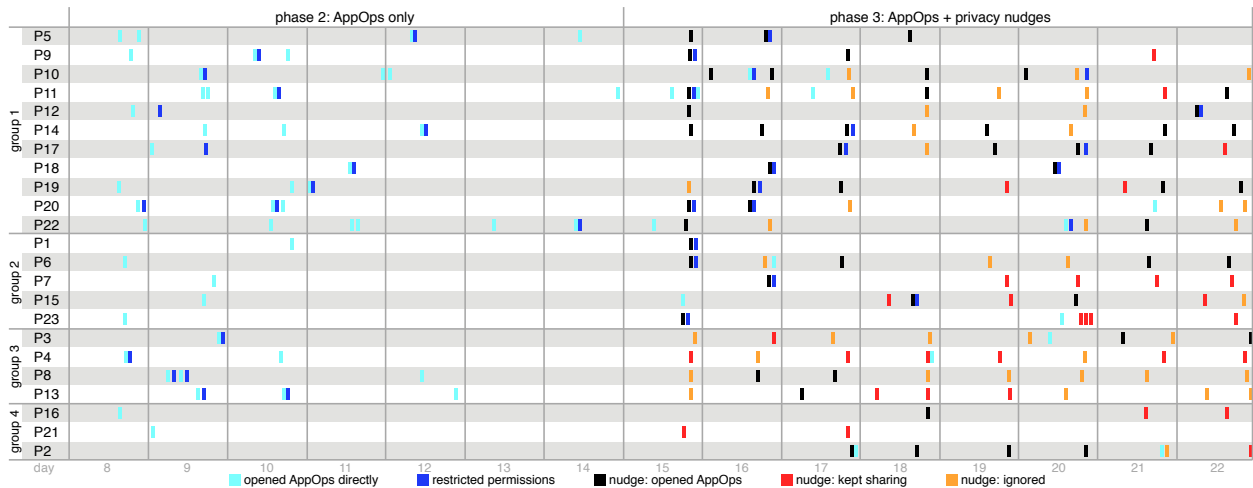


Figure 6: Timeline of participants’ interactions with AppOps and the privacy nudges during phase 2 and phase 3. In the X-axis, each column represents a day during phase 2 (7 days) & phase 3 (8 days).

sion adjustments in phase 3 were made in response to a nudge. Only three participants made permissive adjustments due to loss of app functionality. In the interview, P10 noted that he restricted and later permitted Facebook’s access to the clipboard, because he was unable to copy&paste in Facebook. Participants restricted permissions from 47 distinct apps, including both self-installed and pre-installed apps, see Figure 5. Participants restricted 122 permissions such as location 30 (25%), contacts 25 (20%), calendar 8 (7%), and call logs 6 (5%). Other restricted permissions included: post notification 10 (8%), SMS 9 (7%), camera 7 (6%), record audio 7 (6%).

To understand our data further, we analyzed participants’ restrictive adjustments in phase 2 & phase 3, and divided them into four groups based on our analysis. We additionally examined the groups’ comfort level by analyzing participants’ responses to 5-level likert scale question about location, calendar, contacts, and call logs.

Group 1: restrictive adjustments in phases 2 & 3. Although they made restrictive changes already in phase 2, the nudges led 11 participants to make additional adjustments in phase 3. For instance, P11 & P17 restricted 89% and 50% additional permissions. These participants were overall uncomfortable sharing their personal information with apps. This suggests that even active, privacy conscious participants benefited from the additional information provided by privacy nudges, which triggered them to further review and adjust their permissions to better match their privacy preferences.

Group 2: restrictive adjustments in phase 3 only. These 5 participants made no restrictive adjustments in phase 2. However, the nudges received in phase 3 triggered them to adjust their permissions. These participants were also overall uncomfortable sharing their personal information with apps. This suggests that the nudges provided additional value compared to AppOps alone, triggering them to actively review and adjust their app permissions.

Group 3: restrictive adjustments in phase 2 only. Although the privacy nudges triggered three of these 4 participants to review their permissions in phase 3, they made no restrictive adjustments.

These participants reported mixed levels of comfort sharing personal information with apps. There are multiple potential explanations for why these participants made no adjustments in phase 3. First, their phase 2 adjustments may have sufficed, particularly for P8 & P13, who adjusted 12 and 10 permissions, respectively. Second, aspects of the nudge affected participants' experience. P4 stated in the interview that she always received the nudges when she was at work and therefore never had time to interact with them. Third, participant-specific issues. In the interview, P13 reported that phase 3 was a very busy week for him, which prevented him from interacting much with the nudges.

Group 4: no restrictive adjustments in either phase. Although the privacy nudges did not trigger these 3 participants to adjust their permissions, two did review them in response to the nudges. These participants reported being comfortable or neutral sharing their personal information with apps, which may explain the lack of restrictive adjustments.

5.4 Interaction with the Nudges

In total, participants received 125 nudges. We report details of how participants interacted with them.

Did participants understand the nudge? In our survey, we presented each participant with a nudge screenshot and asked them about their understanding of the nudge, its options, the trust cues, and the detailed report. All of the participants understood the nudge, the options, and the trust cues ("Notification provided by AppOps"). Nine participants did not understand option two in the nudge ("show me more before I make changes"). Four never chose this option, possibly because they did not understand its meaning or function.

How did participants interact with the nudges? As Figure 7 shows, participants responded to 53 (42%) nudges by choosing "let me change my settings" to open AppOps and 31 (25%) nudges by choosing "keep sharing my [data]." Participants ignored 41 (33%) nudges.

Although some participants may have chosen "keep sharing my [data]" to express satisfaction with how apps were accessing their personal information, our interviews revealed they also used it to close nudges when they came at unsuitable times (e.g., busy at work or about to use another app). P4 stated: "for one [nudge] I said keep sharing because as I said earlier I didn't have time, because if I saw that normally I would have definitely changed it, if I wasn't at work." Similarly, our interviews showed that participants ignored nudges because they received them at unsuitable times. For instance, P19 also explained: "the first time [the nudge] came up I was on a run and it covered my running app. All of a sudden I couldn't hear [my running app] telling me my mileage anymore. So I opened my screen and I swiped [the lock] and that [nudge] was there and I was so confused I hit back so fast it was gone."

Did participants adjust permissions in response to nudges? We counted restrictive apps' permissions adjustments as direct response to the nudge if a participant responded to a nudge by choosing "let me change my settings" and then made these adjustments within 10 minutes. Fourteen (60%) participants made restrictive adjustments in responses to 17 (13.6%) different nudges out of 125 nudges. Of those, three participants made restrictive changes in response to two different nudges. We acknowledge that in some cases the nudge may have led participants to adjust

permissions after 10 minutes. P17 made restrictive adjustments within 22 minutes. P1 made three adjustments: one of them within an hour.

The nudges may have had indirect influence on participants' decisions to adjust their app permissions. For instance, a participant might review her app permissions in response to a nudge without adjusting them, and then later open AppOps to adjust app permissions. While these indirect effects are difficult to track, the interviews help us to identify a noteworthy occurrence. P10 responded to the first nudge by choosing "show me more before I make changes" to open the detailed report and then chose "let me change my setting". However, he never adjusted his app permissions. After a couple of hours, P10 opened AppOps directly and restricted both the Weather Channel app and HTC Location service access to location. In the interview, P10 described how the nudge helped him realize the Weather Channel app's data access practices mismatched his expectations "this weather app was the most hogging app on my cellphone. I live in a city why do you have to access my location thousands of times in [a] few days? I not only blocked this app, I removed [it]."

To explore if nudging participants about a particular data type triggered adjusting corresponding permissions, we counted the number of adjustments in which the permission matched the data type in the nudge. As reported earlier, participants restricted their permissions in response to 17 nudges. The response to 15 (88%) nudges included at least one permission that matched the data type in the nudge. Thirty (32%) out of 95 restricted permissions matched the data type featured in the nudge. In other words, 68% of restricted permissions were for data types other than the one in the nudge. A possible explanation for this behavior is how AppOps works. If a participant chooses one app, she will be redirected to a new screen listing all the personal information that the selected app has access to, as shown in Figure 1. Thus, it is possible that participants may have initially intended to only adjust the permissions matching the nudge's data type, but adjusted additional permissions as needed. This suggests that the design of the permission manager may sway participants to adjust more permissions than initially intended.

Finally, we explore the effect of the example apps included in the nudge. We checked if participants adjusted permissions for any of the randomly chosen apps in the nudge. As reported earlier, participants restricted their apps' permissions in response to 17 nudges. For nine (53%) nudges, the participants adjusted apps' permissions of at least one of the apps listed in the nudge. For 2 nudges, the participants adjusted the permissions of 2 out of 3 apps listed in the nudge. Out of all 48 apps listed in the 17 nudges, the participants adjusted permissions for 11 (23%) apps, this suggests that example apps listed in the nudge have a moderate effect on participants' decisions to adjust their apps' permissions.

How did participants respond to the 1st and 2nd nudge? In response to their first nudge, 16 (70%) participants chose "let me change my settings", three (13%) choose "Keep sharing my [data]", and four (17%) ignored it. Participants' likelihood to choose "let me change my settings" decreased when they received the second nudge (after 4 days) as Figure 8 shows. This suggests that repeating a nudge about the same data type within a short time may be ineffective, likely because participant preferences do not change within a short time. This leads us to suggest a potential improvement for our nudge. The nudge should provide a mechanism for users to identify apps that they are comfortable sharing their personal information with. Thereafter, the user could receive

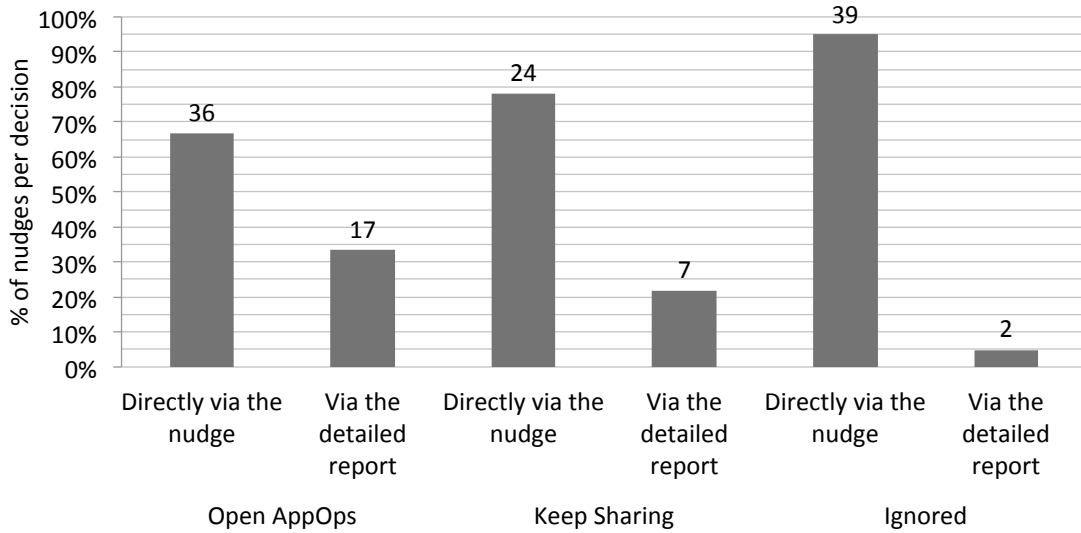


Figure 7: How participants responded to the nudges. The numbers on top of the bars are the absolute numbers.

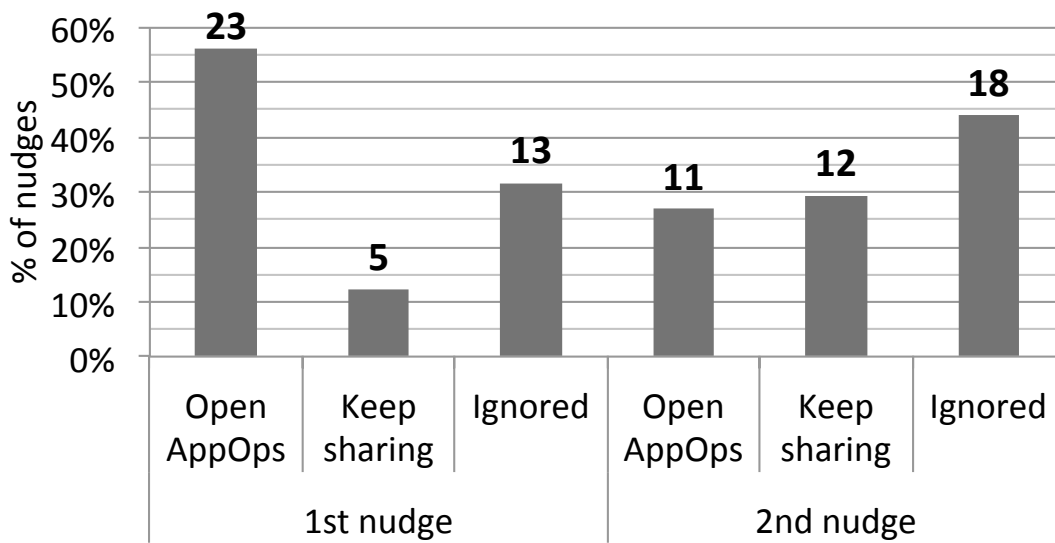


Figure 8: How participants responded to the 1st and the 2nd nudge of the same data type. The numbers on top of the bars are the absolute numbers.

more pertinent nudges including only unidentified apps.

The detailed report: The main goal of providing a detailed report, as shown in Figure 2, was to give interested participants a closer look at apps’ data access patterns. Fourteen participants (61%) chose “show me more before I make changes” to open the detailed report in response to 26 (21%) out of 125 nudges. After opening the detailed report, participants chose “let me change my settings” more often than “keep sharing my [data]” (33% vs. 22%), see Figure 7.

Participants opened the detailed report in eight nudges (47%) before making restrictive adjustments. In each case, the participant made restrictive adjustments for at least one app listed in the detailed report. This suggests that the detailed report is helpful. Though, it is more helpful when the participant intends to make adjustment as the detailed report provided a closer look at the data collection practices of individual apps.

Frequency of access: We explore if an increase in frequency personal information accesses by apps correlates with an increase in participants' likelihood to choose "let me change my settings." Using a regression analysis, we found a significant correlation ($p < .05$) between the frequency of accessing personal information by apps and participants' likelihood of choosing "let me change my settings," particularly for location ($p < 0.01$). This suggests that nudging about apps' frequency of access is effective as it triggered participants to review their apps' permissions, which was the nudge's purpose.

In the interviews, all eight participants indicated that frequencies of access to personal information by apps was the element of the nudge that caught their attention. For instance, P10 explained: "4182 [times] are you kidding me? It felt like I'm being followed by my own phone. It was scary. That number is too high." P17 stated: "the number was huge [356 times], unexpected. Again, big number a bit unexpected."

6 Discussion

In this section, we discuss the effectiveness of the permission manager and nudges, and discuss how to design a more effective nudge based on lessons learned from our study.

6.1 Permission Managers Are Essential

Access to a permission manager alone led participants to actively review and change app permissions. All but one of our participants reviewed their app permissions at least once; half of them did so multiple times. Furthermore, the permission manager led more than half of our participants to exercise control over the personal information apps were allowed to access; participants modified permissions of both popular apps and pre-installed apps. In short, our results highlight the value of using permission managers in mobile platforms, because they give users the control they may want and need. However, service providers have taken highly different paths in their handling of such tools. Violations of end users' privacy by app developers have led Apple to provide users with progressively greater privacy controls in iOS [34]. On the other hand, Google famously removed AppOps from Android phones in 2013 [10], which seemed to clash with Google's public stance of providing users with tools to exercise control over their information [2], and may exemplify ongoing tension between platform providers' (and advertisers') goal of monetizing end users' data, and end users' quest for privacy.¹

¹Analysis of Android source code shows that Google has been expanding AppOps code since then (e.g., increasing the number of permissions to control). This suggests that Google may, perhaps, provide mobile users more control in the future.

6.2 Nudges Can Increase the Utility of Permission Managers

In addition to making users aware of the permission manager, our goal was to design a nudge that assisted users in better managing their privacy. Our results show that even a simple nudge can help users utilize the permission manager to manage their privacy on mobile devices.

The privacy nudges led participants (both those who had and those who had not used the permission manager before) to review and adjust their permissions. This suggests that nudges help both active users, who may not fully utilize the permission manager alone, and users who otherwise might not have made any adjustments, to act to bring their data sharing into alignment with their privacy preferences.

Privacy nudges have been finding their way to mobile platforms. Recently, iOS 8 introduced a form of privacy nudging: if a user allows an app to access her location even when not using the app (e.g. in the background), iOS will occasionally ask if the user wants to continue allowing that [1]. This approach is consistent with our original goal in designing mobile privacy nudges: permission managers are important, but we can increase their effectiveness with periodic nudges.

6.3 How to Design a More Effective Mobile Privacy Nudge?

We learned valuable lessons from testing mobile privacy nudges with real users on their own devices. Based on those lessons, we offer some suggestions for designing and deploying a more effective mobile privacy nudge.

(1) *Personalized*: Our results show that users have different privacy preferences for different apps and data types. This suggests that a personalized nudge may be more effective. A personalized nudge could learn with which apps the user is or is not comfortable sharing her information with, and nudge accordingly. The nudge could also be customized based on the user's previous decisions. For instance, if the user specified that she is satisfied with an app's data collection practices, this app should not be included in subsequent nudges.

(2) *Salient, sticky but not annoying*: We designed our nudge to be full-screen to ensure that it was salient and hard to be ignored. Occasionally, our nudge annoyed some participants, especially when the participant was using or about to use an app, causing them to quickly dismiss the nudge. To be more effective, a privacy nudge should be salient and sticky, but not annoying. To be salient but not annoying, the nudge may cover part of the screen, like Android's special "Heads-up" notification that covers the top third of the screen. Although this format provides fewer options than a full-screen version, it may be less annoying since it is less disruptive. To be sticky, the nudge could include a "remind me later" option. After a small number of rescheduling (e.g. 2), the nudge could be transformed to a notification and sent to the notification bar.

(3) *Configurable*: Users have different preferences for receiving privacy nudges. Some participants liked the daily nudges, whereas others preferred weekly nudges. Some participants had a time preference to receive the nudge (e.g. at night, on Sunday). Moreover, some participants strongly favored receiving a nudge in form of a notification rather than a full-screen nudge. In light of this, a more effective nudge should accommodate these diverse preferences by allowing users to configure the nudge. Users should be able to easily specify the time they receive it, how often, and the form of delivery (full, heads-up, or regular notification).

7 Conclusion

In summary, results from our study indicate that Android users indeed benefit from an app permission manager such as App Ops, and do take advantage of the controls it offers. They also indicate that, even with access to such a manager, user's awareness of the data collected by their apps remains limited. Users would further benefit from receiving nudges in the form of alerts that inform them about the sensitive data collected by their apps. The nudges used in this study were fairly simplistic. Moving forward, even more effective nudges could be generated. In particular nudges would benefit from possibly being further personalized, salient, sticky, and configurable but not annoying.

References

- [1] Apple denies Chinese report that location data are a security risk. <http://on.ft.com/VXpZKR>. Published: 2014-6-12, Accessed: 2014-9-14.
- [2] Google comments on the Preliminary FTC Staff Report on "Protecting Consumer Privacy in an Era of Rapid Change". <http://www.ftc.gov/policy/public-comments/comment-00417>.
- [3] A Acquisti. Nudging privacy: The behavioral economics of personal information. *IEEE Security & Privacy*, 7(6):82–85, 2009.
- [4] A Acquisti and J Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005.
- [5] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proc. SOUPS*, 2013.
- [6] Yuvraj Agarwal and Malcolm Hall. ProtectMyPrivacy: detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proc. MobiSys*, 2013.
- [7] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. Little brothers watching you: Raising awareness of data leaks on smartphones. In *Proc. SOUPS*, 2013.
- [8] Rebecca Balebako, Pedro G Leon, Hazim Almuhammedi, Patrick Gage Kelley, Jonathan Muggan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Nudging users towards privacy on mobile devices. In *Proc. CHI-PINC*, 2011.
- [9] Eun Kyoung Choe, Jaeyeon Jung, Bongshin Lee, and Kristie Fisher. Nudging people away from privacy-invasive mobile apps through visual framing. In *Proc. INTERACT*, 2013.
- [10] EFF. Google Removes Vital Privacy Feature From Android, Claiming Its Release Was Accidental. <http://goo.gl/emMQPa>. Published: 2013-12-12, Accessed: 2014-9-14.

- [11] Manuel Egele, Christopher Kruegely, Engin Kirdaz, and Giovanni Vigna. PiOS: Detecting privacy leaks in iOS applications. In *Proc. NDSS*, 2011.
- [12] Serge Egelman, Adrienne Porter Felt, and David Wagner. Choice architecture and smartphone privacy: There’s a price for that. In *Economics of Info. Sec. & Priv.* Springer, 2013.
- [13] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. Taintdroid: An information flow tracking system for real-time privacy monitoring on smartphones. *Commun. ACM*, 57(3):99–106, 2014.
- [14] Adrienne Porter Felt, Serge Egelman, and David Wagner. I’ve got 99 problems, but vibration ain’t one: a survey of smartphone users’ concerns. In *Proc. SPSM*, 2012.
- [15] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android permissions: User attention, comprehension, and behavior. In *Proc. SOUPS*, 2012.
- [16] Drew Fisher, Leah Dorner, and David Wagner. Location privacy: user behavior in the field. In *Proc. SPSM*.
- [17] Huiqing Fu, Yulong Yang, Nileema Shingte, Janne Lindqvist, and Marco Gruteser. A field study of run-time location access disclosures on android smartphones. *Proc. USEC*, 2014.
- [18] C.S. Gates, Jing Chen, Ninghui Li, and R.W. Proctor. Effective risk communication for android apps. *IEEE Trans. Depend. Secure Comp.*, 11(3):252–265, May 2014.
- [19] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proc. CHI*, 2014.
- [20] Peter Hornyack, Seungyeop Han, Jaeyeon Jung, Stuart Schechter, and David Wetherall. These aren’t the droids you’re looking for: Retrofitting android to protect data from imperious applications. In *Proc. CCS*, 2011.
- [21] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proc. CHI*, 2004.
- [22] J. Jeon, K.K. Micinski, J.A. Vaughan, N. Reddy, Y. Zhu, J.S. Foster, and T. Millstein. Dr. Android and Mr. Hide: Fine-grained Security Policies on Unmodified Android. Technical report, University of Maryland, 2011.
- [23] Jaeyeon Jung, Seungyeop Han, and David Wetherall. Enhancing mobile application permissions with runtime feedback and constraints. In *Proc. SPSM*, 2012.
- [24] Punam Anand Keller, Bari Harlam, George Loewenstein, and Kevin G Volpp. Enhanced active choice: A new method to motivate behavior change. *J. Consum. Psychol.*, 21(4):376–383, 2011.

- [25] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A nutrition label for privacy. In *Proc. SOUPS*, 2009.
- [26] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Privacy as part of the app decision-making process. In *Proc. CHI*, 2013.
- [27] Jialiu Lin, Shahriyar Amini, Jason I Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proc. UbiComp*, 2012.
- [28] Jialiu Lin, Bin Liu, Norman Sadeh, and Jason I Hong. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Proc. SOUPS*, 2014.
- [29] Bin Liu, Jialiu Lin, and Norman Sadeh. Reconciling mobile app privacy and usability on smartphones: could user privacy profiles help? In *Proc. WWW*, pages 201–212, 2014.
- [30] Mohammad Nauman, Sohail Khan, and Xinwen Zhang. Apex: Extending Android Permission Model and Enforcement with User-defined Runtime Constraints. In *Proc. CCS*, 2010.
- [31] Irina Shklovski, Scott D. Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proc. CHI*, 2014.
- [32] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proc. EC*, 2001.
- [33] Richard H Thaler and Cass R Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, 2008.
- [34] Wall Street Journal. Apple Bows to iPhone Privacy Pressures. <http://on.wsj.com/160kjhv>. Published: 2012-2-16, Accessed: 2014-9-14.
- [35] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. A field trial of privacy nudges for facebook. In *Proc. CHI*, 2014.