Helping Smartphone Users Manage their Privacy through Nudges

Hazim Almuhimedi

CMU-ISR-17-111
December 2017

Institute for Software Research
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:
Norman Sadeh (Chair, Institute for Software Research)
Anind K. Dey (Human-Computer Interaction Institute)
Alessandro Acquisti (Heinz College)
Adrienne Porter Felt (Google, Inc.)

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.

I would like to acknowledge the funding I received to conduct the research reported herein under a fellowship from King Abdulaziz City for Science and Technology (KACST) as well as from the National Science Foundation under grants CNS 10-1012763 (Nudging Users Towards Privacy), CNS 13-30596 (Towards Effective Web Privacy Notice & Choice: A Multi-Disciplinary Perspective) and SBE-1513957 (”Towards Personalized Privacy Assistants”), as well as from the DARPA Brandeis program and the Air Force Research Laboratory under agreement number FA8750-15-2-0277 (”Personalized Privacy Assistants for the Internet of Things and Big Data”). The US Government and King Abdulaziz City for Science and Technology (KACST) are authorized to reproduce and distribute reprints for governmental purposes, not withstanding any copyright notation. Additional funding for the research reported herein was also provided to the Mobile Commerce Lab at Carnegie Mellon University by Google and Samsung. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, the Air Force Research Laboratory, the National Science Foundation, the U.S. Government, King Abdulaziz City for Science and Technology (KACST), Google, or Samsung.

Keywords: Privacy, Mobile, Permission, App, Nudge, Privacy Decision Making.

To my parents (Saleh & Fatimah), my wife (Nada), and my kids (Lana & Saleh)

“Privacy means people know what they are signing up for in plain English and repeatedly.”
–Steve Jobs, All Things Digital (D8) Conference, 2010[1].

[1] https://www.youtube.com/watch?v=39iKLwlUqBo

Abstract

The two major smartphone platforms (Android and iOS) have more than two million mobile applications (apps) available from their respective app stores, and each store has seen more than 50 billion apps downloaded. Although apps provide desired functionality by accessing users' personal information, they also access personal information for other purposes (e.g., advertising or profiling) that users may or may not desire. Users can exercise control over how apps access their personal information through permission managers. However, a permission manager alone might not be sufficient to help users manage their app privacy because: (1) privacy is typically a secondary task and thus users might not be motivated enough to take advantage of the permission manager's functionality, and (2) even when using the permission manager, users often make suboptimal privacy decisions due to hurdles in decision making such as incomplete information, bounded rationality, and cognitive and behavioral biases. To address these two challenges, the theoretical framework of this dissertation is the concept of nudges: "soft paternalistic" behavioral interventions that do not restrict choice but account for decision making hurdles. Specifically, I designed app privacy nudges that primarily address the incomplete information hurdle. The nudges aim to help users make better privacy decisions by (1) increasing users' awareness of privacy risks associated with apps, and (2) temporarily making privacy the primary task to motivate users to review and adjust their app settings.

I evaluated app privacy nudges in three user studies. All three studies showed that app privacy nudges are indeed a promising approach to help users manage their privacy. App privacy nudges increased users' awareness of privacy risks associated with apps on their phones, switched users' attention to privacy management, and motivated users to review their app privacy settings. Additionally, the second study suggested that not all app privacy nudge contents equally help users manage their privacy. Rather, more effective nudge contents informed users of: (1) contexts in which their personal information has been accessed, (2) purposes for apps' accessing their personal information, and (3) potential implications of secondary usage of users' personal information. The third study showed that user engagement with nudges decreases as users receive repeated nudges. Nonetheless, the results of the third experiment also showed that users are more likely to engage with repeated nudges (1) if users have engaged with previous nudges, (2) if repeated nudges contain new information (e.g., additional apps, not shown in earlier nudges, that accessed sensitive resources), or (3) if the nudge contents of repeated nudges resonate with users.

The results of this dissertation suggest that mobile operating system providers should enrich their systems with app privacy nudges to assist users in managing their privacy. Additionally, the lessons learned in this dissertation may inform designing privacy nudges in emerging areas such as the Internet of Things.

Acknowledgments

I would like to express my deepest gratitude to my advisor, Dr. Norman Sadeh, for his excellent guidance, unconditional support, and providing me with the best atmosphere for doing research. I would also like to thank my other thesis committee members (Dr. Anind K. Dey, Dr. Alessandro Acquisti, and Dr. Adrienne Porter Felt) for providing me with valuable feedback over the last few years.

Over the years, I have been fortunate to work with great colleagues, friends, and collaborators during my time at Carnegie Mellon University: Shikun Aerin Zhang, Bin Liu, Dr. Florian Schaub, Dr. Shomir Wilson, Dr. Mads Schaarup Andersen, Dr. Lorrie Cranor, Dr. Yuvraj Agarwal, Dr. Idris Adjerid, Dr. Martin Degeling, Dr. Anupam Das, Dr. Sebastian Zimmeck, Josh Gluck, Dr. Mary Glavan, Peter Story, Daniel Smullen. I would also like to thank Linda Moreci (our lab manager), members of the Mobile Commerce lab, and members of the Privacy Nudge group. It has been a great pleasure and a rewarding experience working with you and knowing you all.

I would never have been able to finish my dissertation without the continuous support and encouragement from my parents (Saleh and Fatimah) and siblings (Ibrahim, Suliman, Hazar, Arwa, Abdulmalek, and Faisal). Thank you all for supporting me to complete this great journey.

I am forever grateful to my wife, Nada. I have been through ups and (many) downs and she has always been there and cheering me up. My kids (Lana and Saleh) have always been the light at the end of the tunnel especially when I almost gave up.

Finally, all praise is due to Allah (“Alhamdulillah”).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Smartphones show significant adoption by mobile users in the US and worldwide. According to recent market reports, smartphone penetration in the mobile market is about 75% in the US and 40% worldwide [44]. This increased adoption of smartphones has been attributed primarily to the availability of mobile apps. Both Android and iOS, the two most popular smartphone platforms, have more than two million apps available to download from their respective app stores [2, 50], and each store has seen more than 50 billion app downloads [1, 49].

Apps provide desirable functionality by accessing sensors integrated with mobile devices (e.g., GPS, camera, accelerometer) and by accessing users' personal data (e.g., phone contacts list, calendar, call logs)[1]. For example, the Google Maps app provides directions by accessing the user's location information. However, apps also access users' personal information for other purposes beyond functionality that users may or may not desire [55, 99, 142, 185]. The top downloaded apps both in iOS and Android collect a wide range of users' personal information such as device unique identifiers (e.g., IMEI), location information, and phone contacts list, and may utilize users' personal information for advertising, profiling, and other purposes without users' knowledge or consent [19, 33, 55, 99, 185].

In light of these perceived privacy violations, studies have shown that many users desire to control what types of personal information apps can and cannot access [22, 30, 65, 107]. In apparent response to such demands, platforms have provided their users with various techniques to manage their privacy and amongst these techniques are permission managers [4, 190]. Permission managers enable users to selectively allow/deny apps access to different types of personal infor-

[1] Going forward, we refer to both sensor data and personal data as users' personal information or simply personal information

mation. For instance, the permission manager enables a user to allow the Facebook app to access her location information but to deny the app access to her photos.

Although permission managers are intended to help users protect their privacy, users often underutilize them or do not use them at all [22, 65, 107]. This underutilization is likely the result of two main challenges. First, privacy is often not the user's primary task [79, 148, 198], and therefore the user might not be motivated enough to use the permission manager to control her app privacy settings. Second, privacy decision making is subject to decision making hurdles (e.g., incomplete information or behavioral and cognitive biases) that often lead to suboptimal privacy decisions [58, 59]. Therefore, users might not be able to make optimal privacy decisions even when using permission managers. For example, users might not deny apps access to their personal information (through permission managers) because users are unaware of how much of their personal information is being collected by apps [65, 73, 118].

Privacy nudges are a promising approach to overcome these two challenges. Nudges are "soft paternalistic" behavioral interventions that do not restrict choice, but attempt to account for decision making hurdles [58, 59, 61, 181]. As a subcategory of nudges, privacy nudges specifically focus on helping users make better privacy decisions [58]. In the mobile context, app privacy nudges aim to temporarily make privacy the primary task for users by informing them of potential privacy risks associated with their apps. The goal is to switch users' focus to privacy management and motivate users to review their app privacy settings [65]. In addition, mobile app privacy nudges also aim to overcome privacy decision making hurdles, especially the incomplete information hurdle: the disparity between users' and service providers' knowledge of collection, use, sharing practices, potential consequences, and available protections concerning users' personal information [59, 61]. To address this disparity, app privacy nudges may, for example, inform users about how much of their personal information is being collected, which brings users' attention to apps' behaviors that they may have been unaware of and helps users make better privacy decisions [65, 73, 110, 118].

As a promising approach to help users better manage their privacy, nudges are finding their way into various commercial platforms. For instance, Facebook recently introduced the "Privacy Checkup" to ensure that the user shares her personal information with the intended audience [100, 159]. The "Privacy Checkup" interrupts users who are about to make a public post, verifies with whom the user wants to share the post, and provides options to share the post with a limited audience (e.g., friends only). Similarly, in iOS 8 Apple 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 inform the user when the app accesses her location in the background and ask if the user wants to continue allowing the app to do so [48, 52]. The implementations of privacy nudges by these commercial platforms shows that nudges are gaining traction as a tool assisting users in managing their complex privacy preferences.

In the mobile app privacy context, prior work has primarily focused on designing privacy nudges to help users make better privacy decisions when choosing new apps to install (i.e. install-time app privacy nudges) [85, 118, 134, 142, 162]. The install-time nudge works by informing users of privacy risks associated with the app that the user chooses to install and recommending a more privacy-friendly app than the one chosen. However, an install-time nudge can only be effective (a) if there is an alternative app that is more privacy-friendly than the one chosen by the user, or (b) if the nudge can persuade the user to refrain from installing the privacy-invasive app, even when no alternative app exists. To the best of our knowledge, and assuming alternative apps exist at all, there is no accessible way (outside the experimental settings) to find alternatives to privacy-invasive apps. In addition, prior work has suggested that privacy notices are less likely to garner users' attention at install time in comparison to notices at run-time (e.g., during or after using an app) [74], and that users are less likely to choose another app once they have made up their minds to install a given app [174]. Thus, if the user chooses to proceed and install the app, the install-time app privacy nudges lose their effect.

To extend previous work on mobile app privacy nudges, this dissertation focuses on designing and evaluating runtime app privacy nudges[2] to help users better manage their app privacy settings. Specifically, runtime app privacy nudges aim to (1) increase users' awareness about privacy risks associated with apps that users have previously installed, and (2) motivate users to review their app privacy settings and adjust them as needed. This thesis will attempt to answer the following research questions:

- Is access to an app permission manager an effective way of helping users review and modify their app privacy settings?
- Can app privacy nudges enhance the effectiveness of an app permission manager?
- How effective is the "recent access" privacy nudge content in motivating users to review their app privacy settings in comparison to other nudge contents? The "recent access" nudge content informs users that their sensitive information (e.g., location information) has *recently* been accessed by apps installed on their phones. Both Android and iOS (the most popular mobile operating systems) use recent access nudge contents (see Chapter 4

[2] Going forward, we refer to runtime mobile app privacy nudges as app privacy nudges or simply privacy nudges unless we need to distinguish between install-time and runtime app privacy nudges.

for more details). In Android, a screen in the Settings app show the list of apps that made recent location requests (i.e. within 15 minutes). Similarly, a screen in the Settings app in iOS shows a list of apps that request access to the user's location information. Next to each app, a purple arrow refers to an app that has recently accessed the user's location information, and a gray arrow refers to an app that had accessed the user's location information within the past day.

- Can we design nudge contents that are more effective than the "recent access" nudge content in motivating users to review their app privacy settings?
- How do mobile users respond to repeated app privacy nudges?
- Can we keep users engaged with repeated nudges by updating or varying nudge contents? Which is more effective in keeping users engaged with repeated nudges, updating or varying nudge contents?
- What factors may motivate users to engage with repeated nudges?

The thesis of this dissertation can be summarized as follows.

> *It is possible to design runtime mobile app privacy nudges that increase users' awareness of privacy risks, motivate them to review privacy settings, and cause them to adjust their app privacy settings to reflect their privacy preferences. Effective nudge contents highlight unexpected contexts, purposes, and potential implications of accessing sensitive information by apps. User engagement with repeated nudges can be maintained by limiting them to users who engaged with earlier nudges, to repeated nudges that contain new information, or to repeated nudge contents that resonate with users.*

The results of the first field study (see Chapter 3) indicate that users benefit from an app permission manager, with a majority of our participants taking advantage of the controls it offers. The results, however, indicate that, even with access to a permission manager, users' awareness of the privacy risks associated with apps already on their devices remains limited (e.g., users are unaware of how frequently apps collect their personal information). Therefore, users are unable to fully utilize the permission manager to protect their privacy, and they would further benefit from receiving nudges that effectively: (1) increase awareness about privacy risks associated with their apps, and (2) motivate them to review and adjust their app privacy settings as needed. In addition, the results of the first study show that privacy nudges can indeed help users utilize the permission manager to better manage their app privacy settings. Specifically, the privacy nudges led all but two of our participants to review their app permissions at least once within the eight days after having access to the permission manager for a week. More importantly, the privacy nudges led more than half of our participants to exercise control over the personal information that apps were allowed to access. In sum, privacy nudges were shown to be a promising approach in helping users better manage their app privacy settings.

The results of the second online experiment (see Chapter 4) suggest that not all app privacy nudge contents are equally effective in motivating mobile users to review and possibly adjust their app privacy settings. Although the "recent access" nudge content is used by current mobile operating systems to inform users that their location information has been recently used by apps on their phones, our results show that it is one of the least effective nudge contents. Only 59% of respondents in our experiment indicated their willingness to review and possibly adjust their app privacy settings in response to the "recent access" nudge content. In contrast, other nudge contents were more effective in motivating mobile users to review their app privacy settings. In particular, significantly more respondents indicated their willingness to review their app privacy settings in response to five different nudge contents: informing users of the context in which their personal information has been accessed (76%), informing users of the purpose for accessing their personal information (two nudge contents: 76% and 83%), and informing users of potential implications of accessing and then reusing personal information for secondary usages (two nudge contents: 72% and 75%). In summary, app privacy nudges can help mobile users make informed privacy decisions through nudge contents designed to inform users of: unexpected app behavior, purposes unrelated to the apps' main functionalities, or potential implications of secondary usage of users' personal information.

The results of the third field study (see Chapter 5) demonstrates that user engagement with nudges decreases as users receive them repeatedly. In particular, the number of participants (51.7%) who indicated their willingness to review and adjust their app settings in response to the first nudge is more than those who did so in response to the repeated nudge (39.7%). In addition, the number of app setting adjustments (67%) made in response to the first nudge is more than adjustments made in response to the repeated nudge (33%). Nonetheless, our results identified factors that may motivate users to engage with repeated nudges. Our results show that users were more likely to engage with repeated nudges (1) if users had engaged with previous nudges, (2) if repeated nudges contained new information (e.g., additional apps, not shown in earlier nudges, that accessed sensitive resources), or (3) if the nudge contents of repeated nudges resonated with users. In summary, although user engagement with nudges decreases as users receive them repeatedly, we identified a number of factors that nudge designers should take into account when sending users repeated nudges.

In light of these studies, this dissertation makes the following contributions to the research of mobile privacy:

- We show that app permission managers are useful tools that help users to manage their privacy on mobile devices. We explored the utility of permission managers at the time when Android (the most popular mobile operating system) had not provided users with privacy tools to exercise control over how individual apps access various types of users' sensitive information[3]. Furthermore, we show that although app permission managers are useful, they cannot alone help users achieve the level of privacy protection that they want.

- We show that mobile app privacy nudges can be effective interventions that make users aware of privacy risks associated with their apps, motivate them to review their app privacy settings, and increase the utility of app permission managers. To the best of our knowledge, we are first to deploy app privacy nudges to users in situ and enable users to act upon the nudges by adjusting their app privacy settings.

- We show that not all privacy nudge contents equally motivate mobile users to review and adjust their app privacy settings. In particular, we show that the "recent access" nudge content (which informs users that their location information has been recently accessed by apps), which is inspired by similar recent access nudge contents in current mobile operating systems, is one of the least effective nudge contents. Additionally, we show that we can design nudge contents that are more effective than the "recent access" nudge content in motivating users to review their app privacy settings. We show that more effective nudge contents inform users of: the unexpected contexts in which their sensitive information has been accessed, the purpose for apps' accessing their sensitive information, or the potential implications of how their sensitive information can be reused.

- We show that user engagement with nudges decreases as users receive them repeatedly. Specifically, fewer participants indicated their willingness to review their app settings in response to a repeated nudge than to a first nudge. In addition, participants made fewer app setting adjustments in response to repeated nudges. These results align with previous work on stimuli which suggest that user engagement with stimuli decreases with repeated exposure (e.g., [83, 179, 183]).

[3] Google introduced the "App permissions" manager in Android Marshmallow (Android OS version 6) and above. See Chapter 2 for more details.

- We show that receiving a repeated nudge with a different nudge content than the previous nudge (varying the nudge's content), does little to influence users' decisions in comparison to receiving a repeated nudge that shows the same nudge content as the previous nudge albeit updated (updating the nudge's content). Nonetheless, anecdotal evidence suggests that varying a nudge content does have some influence on users' decisions. Simultaneously, some participants did not notice the difference between the nudge contents, most likely due to the nearly identical look and feel of the two nudges. It could be the case that varying a nudge content does little to influence users' decisions when receiving repeated nudges in comparison to updating a nudge content. Alternatively, the difference between nudge contents was not conspicuous enough and, thus, the influence of varying nudge contents (if it exists) was diminished. We suggest that future work may explore the effect of varying nudge contents using two dissimilar nudge designs in contrast to the similar designs we used in our experiment.

- We show that a number of factors may motivate users to engage with repeated nudges. In particular, users are more likely to engage with repeated nudges (1) if users have engaged with previous nudges, (2) if repeated nudges contain new information (e.g., additional apps, not shown in earlier nudges, that accessed sensitive resources), or (3) if the nudge contents of repeated nudges resonate with users. App privacy nudge designers should take these factors into account when sending users repeated nudges. In addition, researchers should consider these factors when designing experiments which evaluate the long-term effect of app privacy nudges.

- We show that two benefits of app privacy nudges emerged from our empirical results. First, app privacy nudges increase users' awareness of privacy risks associated with their apps by highlighting such risks in nudge contents. Second, app privacy nudges make privacy the primary task for users at least briefly, switch users' attention to privacy management, and motivate users to review and possibly adjust their their apps privacy settings. The latter benefit extends the effect of app privacy nudges beyond the limited content of the nudge itself. We refer to this phenomenon as the "spill-over effect" of app privacy nudges. When users are presented with a nudge about a particular type of sensitive information (e.g., location information) or particular apps, users take advantage of the opportunity to also review app settings related to other types of sensitive information (e.g., phone contacts list) and review apps settings of other apps (i.e. other than the apps presented in the nudge).

The remainder of this dissertation is organized as follows. Chapter 2 looks at prior work on mobile app privacy and privacy nudges, and shows how this dissertation extends previous work.

Chapter 3 reports the results of the first study, which explores the effectiveness of the permission manager before and after introducing the privacy nudges. Chapter 4 reports the results of the second study, which designs and evaluates effective nudge contents. Chapter 5 reports the results of the third study, which evaluates user engagement with repeated nudges. Finally, Chapter 6 summarizes the lessons learned throughout this dissertation and describes future research directions beyond the work in this dissertation.

Chapter 2

Background & Related Work

2.1 Privacy Mechanisms in Mobile Platforms

Both Android and iOS provide tools and controls to help users protect their privacy; However, each platform adopts different approaches to protect users privacy. Below we describe how each platform handles users' privacy.

2.1.1 Android

In Android, app developers are required to declare the list of permissions that the app needs in order to access protected resources (e.g., the user's location) [6]. When a user chooses to install a new app from the Google Play store, the user is presented with a list of permissions that the

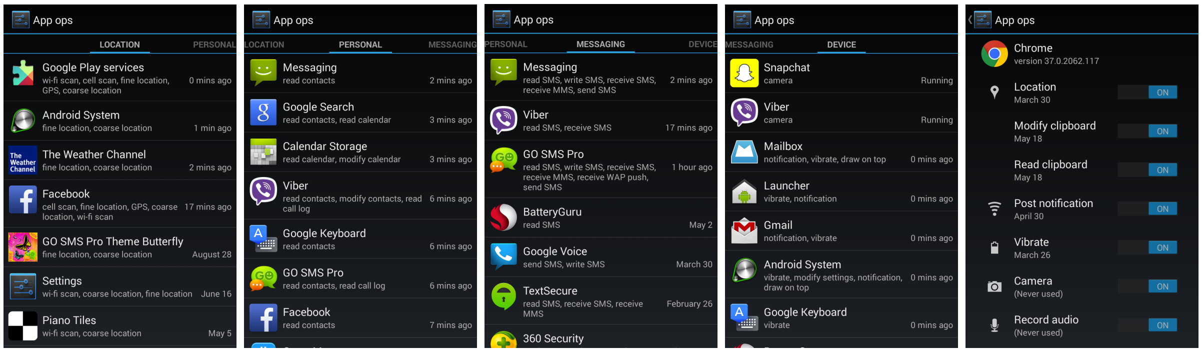

Google Play services	wi-fi scan, cell scan, fine location, GPS, coarse location	0 mins ago
Android System	fine location, coarse location	1 min ago
The Weather Channel	fine location, coarse location	2 mins ago
Facebook	cell scan, fine location, GPS, coarse location, wi-fi scan	17 mins ago
GO SMS Pro Theme Butterfly	fine location, coarse location	August 28
Settings	wi-fi scan, coarse location, fine location	June 16
Piano Tiles	wi-fi scan, coarse location	May 5

Messaging	read contacts	2 mins ago
Google Search	read contacts, read calendar	3 mins ago
Calendar Storage	read calendar, modify calendar	3 mins ago
Viber	read contacts, modify contacts, read call log	6 mins ago
Google Keyboard	read contacts	6 mins ago
GO SMS Pro	read contacts, read call log	6 mins ago
Facebook	read contacts	7 mins ago

Messaging	read SMS, write SMS, receive SMS, receive MMS, send SMS	2 mins ago
Viber	read SMS, receive SMS	17 mins ago
GO SMS Pro	read SMS, write SMS, receive SMS, receive MMS, receive WAP push, send SMS	1 hour ago
BatteryGuru	read SMS	May 2
Google Voice	send SMS, write SMS	March 30
TextSecure	read SMS, receive SMS, receive MMS	February 26
360 Security		

Snapchat	camera	Running
Viber	camera	Running
Mailbox	notification, vibrate, draw on top	0 mins ago
Launcher	vibrate, notification	0 mins ago
Gmail	notification, vibrate	0 mins ago
Android System	vibrate, modify settings, notification, draw on top	0 mins ago
Google Keyboard	vibrate	0 mins ago

Chrome
version 37.0.2062.117

Location	March 30	ON
Modify clipboard	May 18	ON
Read clipboard	May 18	ON
Post notification	April 30	ON
Vibrate	March 26	ON
Camera	(Never used)	ON
Record audio	(Never used)	ON

Figure 2.1: AppOps has four different tabs: location, personal (e.g., phone contacts and calendar), messaging(e.g., SMS), and device (e.g., vibration, camera). Each tab lists all apps that accessed the corresponding category of data, e.g., location (*left*). Selecting a specific app opens a screen showing all permissions accessed by the app (*right*).

app requires. Android shows the list of permissions required by the app as a primary feature to inform users about potential privacy and security risks associated with installing the app. The user has to accept all the required permissions to install the app.

At runtime, Android provides users with privacy control over location information. Users can turn off location information and prevent all apps from accessing it; However, users cannot exercise control over individual apps. Additionally, Android provides a list of apps that recently requested the user's location as a transparency feature. Traditionally, Android has not provided users with privacy controls over other types of personal information (e.g., phone contacts list or calendar)

In Android 4.3, Google introduced a permission manager, called *AppOps*, that was hidden by default and required an external app to access [16]. AppOps allows users to selectively grant or deny permissions for installed apps. AppOps, shown in Figure 2.1, organizes 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. In Android 4.3, AppOps allows users to control 31 different operations (e.g., would an app be able to read from the clipboard?) [113]. Google incrementally expanded the operations that AppOps controls to 43 in Android 4.4 and 48 in Android 5 [114, 115]. However, AppOps has been made inaccessible since Android 4.4.2 [94], unless the device is rooted.

In Android Marshmallow (Android OS version 6.0), Google replaced install-time permission screens with just-in-time permission requests and introduced the "App permissions" manager (as shown in Figure 2.2) which allows users to selectively grant or deny permissions for installed apps [5, 43]. When an app requests access to a sensitive resource for the first time (e.g., location information), the user is prompted to either grant or deny such a request. If the request is granted, the app can access the resource anytime. The user can revert the decision anytime either through the "App permissions" manager (as shown in Figure 2.2, left) or the per-app permission setting (as shown in Figure 2.2, right).

2.1.2 iOS

iOS handles privacy and security primarily through the app review process [68]. Apple provides detailed guidelines to app developers to ensure that their apps incorporate the appropriate security and privacy measures before the apps are approved and listed on the app store. For instance, an app would be rejected if the app does not function unless the user shares her personal information

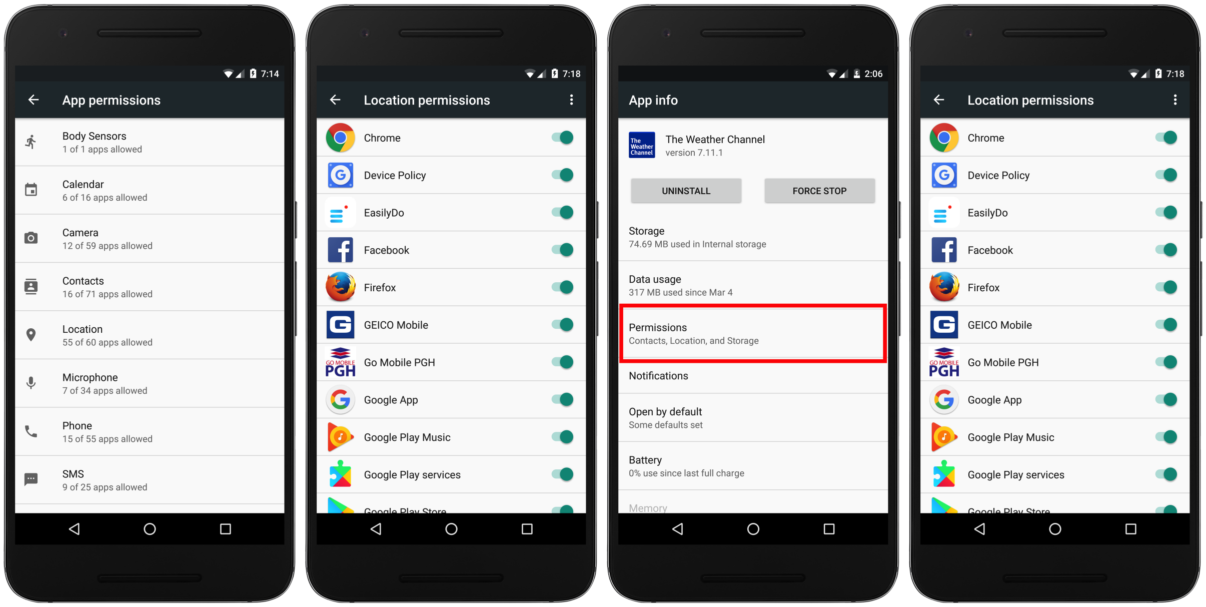

App permissions

Body Sensors
1 of 1 apps allowed

Calendar
6 of 16 apps allowed

Camera
12 of 59 apps allowed

Contacts
16 of 71 apps allowed

Location
55 of 60 apps allowed

Microphone
7 of 34 apps allowed

Phone
15 of 55 apps allowed

SMS
9 of 25 apps allowed

(a)

Location permissions

- Chrome
- Device Policy
- EasilyDo
- Facebook
- Firefox
- GEICO Mobile
- Go Mobile PGH
- Google App
- Google Play Music
- Google Play services
- Google Play Store

(b)

App info

The Weather Channel
version 7.11.1

Storage
74.69 MB used in Internal storage

Data usage
317 MB used since Mar 4

Permissions
Contacts, Location, and Storage

Notifications

Open by default
Some defaults set

Battery
0% use since last full charge

Memory

(c)

Location permissions

- Chrome
- Device Policy
- EasilyDo
- Facebook
- Firefox
- GEICO Mobile
- Go Mobile PGH
- Google App
- Google Play Music
- Google Play services
- Google Play Store

(d)

Figure 2.2: Android Marshmallow (Android OS version 6) and above provides users with tools to manage permissions of individual apps. The "App permissions" manager (a) lists categories of permissions that apps requested access to. If the user clicks the location category, the users is presented with location permissions (b). The user can also adjust permissions for individual apps through the Settings screen (c). If the user selects the Permissions setting (highlighted in a red rectangle), the user is presented with all permissions for this particular app (d).

with the app [69]. In addition, the guidelines include specific requirements if apps use personal information that is considered sensitive. For instance, an app is required to provide a privacy policy if the app uses health information through the HealthKit or financial information through the Apple Pay frameworks [69].

iOS follows an opt-in model and therefore an access request by an app is denied by default unless the user explicitly allows it. When an app requests access to the user's personal information for the first time, the user will be prompted to allow or deny that request [46]. If the user accepts the access request, she can later revoke it through the the privacy controls (i.e. a permission manager) that iOS provides [46]. iOS gradually has been providing users with privacy controls to enable them to selectively grant/deny apps' access to personal information. iOS 4 allowed users to selectively control how individual apps access their location only [46]. The controls expanded to include phone contacts lists in iOS 5 [190] and other types of personal information (e.g., calendar, microphone, camera, HealthKit, HomeKit) in iOS 7 and iOS 8. In iOS 8, users are provided with an even finer-grained control over location information. The user can allow the app either to use her location all the time or only while the user is using the app (i.e. preventing apps from accessing the user's location in the background). 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 [48, 52].

2.2 Evaluating and Enhancing Privacy Mechanisms in Mobile Platforms

2.2.1 Android

When a user chooses to install a new app, the user is presented with a list of permissions that the app requires. Android shows the list of permissions required by the app to inform users about potential privacy and security risks associated with the app. As a result, an area of research has been focusing on evaluating the utility of this feature, including how to improve its effectiveness.

Felt et al. evaluated the advantage of the permission model adopted by Android over the full-privileges model in desktop operating systems (i.e. once the app is installed, it can access virtually everything on the the device.) [103]. By surveying 100 paid apps and 856 free apps, the authors showed that Android apps are less privileged than apps in desktop operating system. Therefore, the permission model in Android is relatively effective in protecting users' security and privacy.

Listing the permissions that the app requests before installation is only effective if the user reads, understands, and makes a decision based on the requested permissions. Kelley et al. conducted semi-structured interviews with 20 Android users to explore whether users read the permissions presented at the installation time, comprehend them, and understand what the permissions allow the apps to do [133]. Although half of the participants reported that they notice or read the list of permissions before installing new apps, most of the participants were unable to understand what these permissions meant or what they allowed apps to do [133]. In a more comprehensive study, Felt et al. conducted an online survey and a lab study to evaluate the utility of Android permissions in warning users about potential security and privacy risks associated with installing new apps [105]. The majority of the participants failed to notice the list of the permissions before installing apps and had limited comprehension of what the permissions meant and entailed. However, about 20% of participants noticed the permissions before installing new apps, and had a fairly good level of permission comprehension. An even smaller number of participants (8%) self-reported that they refrained from installing some apps in the past based on the comprehension of app permissions [105].

Prior work also proposed and evaluated several designs (e.g., “privacy facts”, “privacy granules”, “privacy summary”) to effectively communicate the privacy risks associated with apps that users want to install [85, 118, 134, 142, 162]. The goal of these designs is to nudge users away from privacy-invasive apps or to consider apps that are more privacy-friendly than the ones chosen by users. We discuss each one of these designs in detail in Section 2.4.

When apps access the GPS sensor to obtain the user’s location, Android shows a blinking icon in the notification bar presumably to inform the user. Fu et al. [110] designed and proposed a full-screen location access notification and evaluated it against the existing Android location access disclosure method (i.e. the blinking icon in the notification bar). The authors showed that their proposed notification is more effective than Android existing notification. In addition, they showed anecdotal evidence of how the proposed notification affected participants’ behavior (e.g., influence users to stop using intrusive apps).

Traditionally, Android has not provided users with fine-grained privacy controls to enable users to selectively allow/deny apps’ access to their personal information. As a result, a body of prior work has been focusing on developing *fine-grained privacy controls* per individual app. Apex [152] modified Android to enable users to selectively grant or deny access to specific permissions on a per-app basis. In addition, Apex allows users to specify constraints on how apps use resources (e.g., frequency or time of the day) [152]. Enck et al. developed “TaintDroid” a modified version of Android that tracks in real-time how third-party applications access and transfer sensitive information [99]. The authors evaluated TaintDroid by selecting 30 apps that require access to both the Internet (i.e. to transfer data) and sensitive data (e.g., location, device ID). The authors found that 20 of those apps potentially violate users' privacy by transferring sensitive data off the phone to destinations including advertisers [99]. An extension of TaintDroid is AppFence, which empowers users with privacy controls to either provide fake data to third-party apps or block those apps from using the Internet to transfer sensitive data off the phone especially to advertisers [121]. Because blocking apps' access to such data might effect users' experiences, the authors evaluated how 50 apps would work both with and without the privacy controls. They found that users can exercise control over their sensitive information without affecting their experience with most apps. However, a third of the tested apps would not function with the privacy controls which necessitated a trade-off between usability and privacy [121]. In a similar approach, Beresford et al. developed "MockDroid" a modified version of Android that allows users to provide fake information to third-party applications [76]. However, when fake data is provided to an app, the user is notified and given the chance to provide real data, if needed. The goal is to give the user the ability to choose between usability and privacy [76]. Zhou et al. developed "TISSA" a modified version of Android that provides "privacy mode" which allows users to control how apps can access sensitive data [200]. The privacy settings manager allows users to provide apps with anonymized or fake data instead of the actual real data. The authors evaluated the proposed system with 24 free apps and found it to be effective in capturing privacy leaks [200]. Jeon et al. [125] developed two tools, "Dr. Android" and "Mr. Hide", that provide users with privacy controls over apps but with no modifications to the Android system. Mr. Hide works as a service to facilitate how installed apps access sensitive data, whereas Dr. Android modifies (i.e. rewrites) those apps to access sensitive data through Mr. Hide. The authors evaluated the approach using both automated and manual testing on real apps, and they found the approach to work with minimal issues [125].

2.2.2 iOS

Although the app vetting process is the primary technique to handle privacy and security issues in iOS, the vetting process was not documented at the beginning and therefore was not transparent. In addition, iOS had not yet provided privacy controls to enable users to manage their app privacy settings[1]. Due to the lack of transparency and privacy controls, researchers have examined privacy risks associated with apps on iOS. Egele et al. developed "PiOS" which uses static analysis to automatically identify potential privacy leaks of users' personal information to

[1] This was true back in 2011. However, since then Apple made the vetting process more transparent and also gradually introduced privacy controls.

third-party destinations [95]. Using PiOS, the authors analyzed 1407 apps (825 free apps from the app store and 582 apps from Cydia[2]) and found that 55% of the analyzed apps leaked the device's unique identifier (e.g., IMEI) to third-party destinations [95]. However, the authors found that most apps, including apps in Cydia, did not leak other types of users' personal information (e.g., contacts list, photos, call logs, etc.) [95].

iOS gradually has been providing users with privacy controls to manage their app privacy settings. To examine the effectiveness of the privacy controls, Fisher et al. [107] asked 273 iOS users to take screenshots of the apps' location privacy settings on their devices, and examined whether users allowed or denied apps access to their location. The authors also explored how to use the privacy decisions made by users to predict their future decisions. The results showed that indeed users utilize privacy controls in iOS. For instance, most users denied at least one app from accessing their location. The results also showed that users' decisions to grant/deny apps the access to location information differ by apps. For instance, while almost all users granted the Google Maps app access to their location, only half of the users, who had installed the Shazam app [3], granted it access to their location [107].

ProtectMyPrivacy (PMP) was introduced to enable iOS jailbroken users to selectively grant or deny apps access to an extended list of users' personal information beyond what iOS provided (e.g., control over device unique identifier (e.g., IMEI)) [63]. PMP also allows users to provide fake information to apps to avoid effecting users' experiences (e.g., some apps might crash if denied access to some information) [63]. In addition, PMP provides users with privacy setting recommendations for new apps via crowdsourcing as we discuss in detail in Section 2.4 [63].

When an app in iOS requests access to users' personal information (e.g., the user's location) for the first time, the user is prompted with a dialog box and asked to grant or deny that request. The developer of the app is encouraged to provide an explanation for the purpose of the request. Tan et al. explored the prevalence of developer-specified explanations for access requests and evaluated their effectiveness on iOS users' behavior [180]. By analyzing 4400 apps, the authors found that only 19% of the access requests included explanations by developers. The authors, then, compared the effectiveness of access requests with and without explanations in a 772-participant within-subject AMT online experiment. Participants were presented with a screenshot of an access request dialog with an explanation, then another screenshot of a different app with no explanation. To examine the effect of the content of the explanation, participants were asked to perform a third task in which they were presented with a screenshot of a request dialog with

[2] The iOS app store for jailbroken phones

[3] A popular music recognition app. URL:`http://www.shazam.com`

a explanation from a list of potential explanations that covered a wide range of scenarios (e.g., accessing and sending the data to the provider's servers or accessing the data to provide near-by offers). The authors found that access requests were more likely to be granted when an explanation is provided regardless of the actual content of the explanation [180]. They also found users experience is improved when the purpose of the access request is presented to the user [180]. However, the provided explanation by apps from the app store did not help users understand why their personal information was being requested and whether it would be subjected to secondary uses [180].

2.3 Privacy Nudges

Nudges are "soft-paternalistic" behavioral interventions that do not restrict choice, but attempt to enhance and influence the user's choice by accounting for decision making hurdles such as asymmetric information, heuristics (e.g., anchoring) and bounded rationality, and cognitive and behavioral biases (e.g., overconfidence) [58, 59, 61, 181]. For instance, to nudge students to eat healthy food, we can design the school's cafeteria menu to make the healthy food attractive (e.g., putting the healthy food first in the menu) [181]. This is considered a nudge because it does not restrict choice (e.g., students can still choose unhealthy food) but attempts to influence students' decisions by making healthy food salient [181]. Nudges are shown to be useful in different domains such as retirement plans, health care, and organ donations [181]. 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 or users' privacy and security preferences that change over time [58, 176].

Of all decision making hurdles, we are particularly interested in asymmetric and incomplete information and how nudges can help overcome this hurdle in the context of mobile app privacy.

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 [61]. This phenomenon has been attributed to ineffective communication of privacy risks and protections to users in privacy policies and notices [124]. In the context of mobile privacy, researchers recently have identified a similar information asymmetry, in which users are unaware of the data collection performed by mobile apps [55, 99, 104, 105, 142, 174, 185]. This has led to alternate proposals for presenting privacy risks to users in a manner that is more readable and salient [131]. However, research focused on privacy contexts finds that information disclosures lead 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. [62] 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 such as privacy nudges for communicating relevant privacy information to users [58].

In the mobile context, the potential for nudges that account for asymmetric information to support privacy decision making is appealing. It may include notifications that, in contrast to traditional notices, highlight the recipients, contexts, or types of personal information being shared via a mobile device. [65, 72, 73, 110].

2.4 Nudging Users to Change their Privacy & Security Behaviors and Decisions

Thaler and Sunstein posit that any design choice (deliberate or not) influences people's decisions and nudges them toward one decision or another [181]. By adopting this broad definition we categorize as research on nudges, any prior work that developed approaches to influence users' privacy and security behaviors and decisions. To navigate previous work easily, we arrange it into two categories: (1) the mobile context and (2) other contexts (e.g., online social networks, or web browsers). [140] provides an extensive literature review of security and privacy nudges.

2.4.1 Mobile Privacy & Security Nudges

Prior work has looked at the effectiveness of feedback in helping users review and possibly adjust their privacy settings in mobile location-sharing systems. Tsai et al. explored the effect of feedback on users' privacy preferences for a mobile location sharing application that allows users to share their location with friends on Facebook [186]. In a 56-participant 4-week field study, participants had the ability to create and edit rules to control how their locations are shared. The treatment group (feedback condition) was given access to a list of people who requested their location, at what time, and where (feedback condition); The control group was not given any feedback. The authors found that participants in the feedback condition were more comfortable sharing their location even with strangers [186].

Jedrzejczyk et al. explored how real-time feedback notifications affect users' privacy settings and behaviors in a location sharing application [123]. The authors conducted a field study with 12 participants divided in two groups: 5 and 7 people respectively, related to one another within each group (e.g., family members and close friends). The results showed that real-time feedback notifications reduced the number of unreasonable/unnecessary requests because feedback notifications made requesters accountable for their location requests. However, the participants never used the privacy settings provided by the application nor adjusted them in response to the feedback notifications they received during the study, presumably because participants have close relationships to one another [123]. Patil et al. explored how the timing of the feedback (i.e. before vs. after the disclosure occurs) and the availability of controls to act upon the feedback (i.e. available vs. no controls) affect users' privacy decisions in location sharing social networks [161]. Using experience sampling, 35 participants received hypothetical location requests from pre-defined groups (e.g., participants' family, friends, colleagues). Participants in the feedback condition were informed that their location was shared according to their privacy settings (that they identified at the beginning), asked to specify their level of comfort with the decision, and asked whether the request should have been denied. Participants in the decision condition were prompted to grant or deny a location request and to specify their level of comfort with the decision. When comparing participants' decisions in situ to their initial decisions, participants in the feedback condition felt that they were oversharing, whereas the participants in the decision condition tended to be more comfortable with sharing (i.e. their in situ settings were less protective than their prior settings) [161].

Prior work has primarily focused on nudging users towards less privacy-invasive apps during the installation process. Kelley et al. designed and evaluated a "Privacy Facts" interface to help users select apps that request fewer permissions [134]. The "Privacy Facts" interface consists of two sections: what information the app collects, and what third-party libraries the app uses (e.g., advertisement or analytic). The authors evaluated the privacy fact interface against the Android default permission screen. Using a 20-participant lab study and a 366-participant online experiment, participants were asked to select a number of apps and were shown either the privacy facts or the Android default permission screen. The authors found that privacy facts can influence users to select apps that request fewer permissions [134]. Harbach et al. enriched permission dialogs with personalized examples from the user's device (e.g., showing a personal photo stored in a user's device if an app requests access to files stored in the device.) to make risks more salient and overcome habituation [118]. Using experiment settings similar to Kelley et al. [134], the authors conducted a 36-participant lab study and a 332-participant AMT online study to compare the proposed approach to the Android default permission screen. The authors found that personalized examples are effective in helping users to either choose apps with fewer requested permissions or to not install intrusive apps altogether [118]. Lin at al. designed and evaluated a "Privacy Summary" interface to bridge the gap between users' expectations and how apps are actually collecting and using their information [142]. Through crowdsourcing, the privacy summary shows the percentage of users who feel surprised that the app is collecting some types of personal information. Using an online experiment, the authors evaluated the privacy summary interface against the Android default permission screen by asking participants whether they would recommend a friend to install a given app. They found that users in the privacy summary interface condition were more aware of privacy implications of installing the apps. In addition, the privacy summary interface was easier to comprehend. Choe et al. explored how privacy rating affects users' perceptions of apps and whether the positive or negative framing of the rating has an influence on such perception [85]. In a 332-participant between-subject online experiment, the participants were asked to search for a weather app and presented with the results (e.g., app description page) along with the privacy rating. The authors found that apps that received high privacy ratings where trusted more than apps with lower privacy rating regardless of the framing. However, when the privacy rating of an app is low, the negative framing has a counter effect (i.e. participants trusted the app more) [85]. Paturi et al. designed and evaluated "Privacy Granules" which are icons aimed to communicate the potential risk of the ads library included in Android mobile apps [162]. The authors focused on the risk of capturing three data types: location, device unique identifier (e.g., IMEI), or web search queries. Through both static and dynamic analysis, the authors analyzed the top 50 free Android apps to identify what personal information is collected and what libraries are bundled with the apps. In a within-subject online experiment, 272 participants were presented with the privacy granules followed by the Android default permission screen and asked about four different apps in term of easiness to find privacy threats from the apps and from third-party libraries. Users reported finding threats to identity and search queries easier in the privacy granules than the Android default interface; however, both interfaces were effective in helping users find privacy threats related to location information. In addition, the privacy granules interface was effective in informing users about privacy threats from third-party ads libraries [162].

Researchers have also explored using recommendations to assist users in making better app privacy decisions. Protect My Privacy (PMP)[4] used crowdsourcing to recommend privacy settings (both granting and denying) for each app [63]. The recommendations followed a majority rule and focused on representative apps (e.g., installed by 5 users) and users (e.g., made decisions for 10 other apps). The total number of recommendations shown by PMP to users is slightly less than 2M and the acceptance rate is 67% [63]. To move from one recommendation-fits-all

[4] Protect My Privacy was described earlier in 2.2

to more fine-grained recommendations, [142, 143, 144] clustered users into a small number of profiles. Liu et al. analyzed how 4.8M Android users made decisions to grant 12 different permissions to apps installed on their mobile devices [144]. By focusing on representative users (e.g., installed more than 20 apps) and apps (installed by at least 10 users), the authors used machine learning clustering algorithms to put users into a small number of profiles based on their decisions to grant or deny apps the access to different permissions. Using a small number of labeled apps (i.e. apps in which a user granted/denied access to permissions), the authors were able, with relatively high accuracy, to assign the user to one of the profiles and, in turn, predict privacy decisions for apps in which the user has not yet labeled [144]. Lin et al. developed a similar approach but, in addition, utilized static analysis to identify why an app requests access to a particular permission (e.g., to provide a function or to deliver ads) [142, 143]. Using AMT, the authors asked 725 participants to express their comfort level with apps accessing different permissions for different purposes. Using participants' responses, the authors clustered users into four profiles using machines learning techniques. The four profiles included two extreme groups (unconcerned and privacy conservatives) and two groups of users in between [142, 143]. Wilson et al. evaluated the effectiveness and influence of privacy profiles on users' privacy preferences for location sharing applications [195]. In a field study, 33 participants installed and used "Locaccino" - a mobile location sharing application - on their smartphones for three weeks. The participants were randomly assigned to two conditions: (1) a profile condition in which participants' privacy preferences were created using a small number of profiles (e.g., people who I want to always share my location with), and (2) a control group where participants created their privacy preferences using standard rule-based approach. Participants were required to complete a daily questionnaire to indicate their satisfaction with location sharing events (both real and artificial) that occurred during the day. The authors found that participants in the profile condition were more likely to share their location than those in the rule-based condition, and participants in both conditions were similarly satisfied [195].

While the majority of research has examined install-time mobile privacy nudges, some work has focused on helping users better manage their privacy for apps that users previously installed (i.e. at runtime). Fu et al. [110] designed and evaluated a runtime location access notification. When an app that the user is currently using accesses the user's location, the user is notified through a full-screen notification. The authors showed that their proposed notification is more effective than Android's existing notification mechanism (i.e. the blinking icon in the notification bar). However, they reported users being annoyed by the full-screen notification especially when apps accessed users' location frequently. Balebako et al. [73] conducted a lab study with 19 participants to evaluate their perception and concerns about data collection practices of two popular game apps. Data collection practices were communicated to participants through two mechanism: a notification at the moment the apps collected and shared the information, and a report that showed how frequently each app shared different types of participants' personal information (e.g., location, device ID). The authors showed that users are unaware of, and surprised by, apps' data access practices, especially frequency [73]. Bal et al. designed and evaluated "Styx": a system that monitors apps' data collection practices and informs users of what these apps are able to infer based on the data they accessed and collected [71]. For instance, Styx informs the user that an app can potentially predict the user's gender when the app obtains the list of installed apps on the user's device [71]. In a between-subject lab study with 55 participants, the authors evaluated the effectiveness of a proof-of-concept of Styx against a basic interface that provides the user with a chronological list of data collection events by apps. Styx was shown to be more effective in informing users about potential privacy risks [71]. Protect My Privacy (PMP) utilized crowdsourcing to recommend privacy settings (both granting and denying) for each app [63]. PMP prompts the user with privacy settings recommendation while the user is using the app (e.g., a recommendation to allow the app to use the user's location but not her phone contacts list). If the user accepts the recommendation, PMP applies the recommended settings instantly. Thompson et al., evaluated two types of nudges to identify apps that access system resources (e.g., vibration and changing wallpaper): a passive nudge in which the user needs to navigate to the Settings app to identify apps that has accessed system resources, and a semi-active nudge in a form of a notification which the user receives in the notification bar [182]. Via a lab experiment, the authors showed that the semi-active nudge was more effective than the passive nudge in grabbing users' attention and in helping users to identify apps which has accessed system resources. In addition, the results of the experiment also showed that these two types of nudges were more effective in helping users to identify apps which has accessed system resources than the control condition (i.e. no accessible mechanisms to identify apps that has accessed system resources) [182]. In Chapter 3, we describe the design of runtime mobile app privacy nudges to inform users about how frequently apps access their personal information, and how we evaluated the effectiveness of the nudges in a field study.

2.4.2 Privacy & Security Nudges in other Domains

Privacy and security nudges have also been proposed and evaluated in domains other than the mobile. In this section, we briefly describe some of these nudges. Although we cite as many related references per concept as we can, we only discuss few references to highlight the concept.

Prior work in online social networks has looked at the effectiveness of *feedback* in helping users review and possibly adjust their privacy settings [100, 126, 192]. [126] designed and evaluated two types of feedback: (1) showing users who viewed their profiles, and (2) showing users what content was viewed recently (e.g., a post or a picture). However, both types of feedback were artificial because thi information is not available on Facebook. By asking participants to install a Facebook application, the authors conducted a 107-participant between-subject experiment for 24 days with three conditions: two types of feedback and a control condition. The authors found that the privacy setting changes made by participants were not statistically significant across conditions, which the authors attributed to the artificial nature of the feedback [126]. Wang et al. designed and evaluated a Facebook privacy nudge to help users consider the audience and content of their posts to avoid regrettable posts [192]. When the user posted some content on Facebook, the nudge showed both the profile pictures of some of the recipients, the total number of recipients, and delayed posting the content to Facebook for 10 seconds. The authors tested the nudge with 28 Facebook users in a 6-week field study, a 3-week control phase followed by a 3-week treatment phase. Although participants did not adjust their privacy settings in response to the nudges, they found the nudges useful especially the audience nudge [192].

Researchers have also designed and evaluated *social cues* to help users make better privacy and security decisions [77, 90, 111, 160]. Besmer et al. designed and evaluated a prototype that uses social cues (both positive and negative) to assist users to configure their privacy access control regarding applications on Facebook [77]. The authors asked 309 participants on AMT to authorize a number of Facebook applications to access their personal data on Facebook. Throughout the authorization process, participants were shown the percentage of other users who also authorized (or did not authorize) the app to access different pieces of their personal information. The authors found that negative social cues influenced users' decisions [77]. In a large-scale study, Das et al. explored how social cues increased users' awareness and use of security tools [90]. Using seven different variations, the authors informed 50K Facebook users of the number of their friends who use either one of three security features: login notifications, approvals, or trusted contacts. The authors found that this simple social cue was effective in nudging users to consider these security features [90].

Prior work in security dialogs and warnings explored approaches to assist users in making *better security decisions* [64, 66, 82, 83, 84, 106, 170, 179]. Brustoloni et al. explored how polymorphic dialogs (dialogs that alway change changing what input is required from the user and how it is elicited) and polymorphic combined with audited security dialogs assist users in choosing less risky decisions [84]. In a 26-participant between-subject lab study, participants (as employees in a company) received emails with attachments related to two scenarios (recruiting applicants and processing customer requests) using Mozilla Thunderbird with an extension to show the modified dialogs. Attachments were associated with justified and unjustified risks. Both approaches were shown to be effective to assist users in making better security decisions when the risk is unjustified [84]. Bravo-Lillo et al. designed and evaluated "attractors" which are interfaces modifications that bring users' attention to critical information to help them make informed security decisions [82]. In two AMT between-subject experiments, the authors evaluated 11 different attractors using two different scenarios: high security risk and no security risk. The authors found attractors to be effective in bringing users' attention to critical information presented in the security dialogs and, thus, helped users to make informed security decisions [82]. Felt et al. evaluated how the manipulation of the appearance of the SSL warnings would improve their effectiveness [106]. In a large-scale field study, the authors manipulated the warning by changing the default image of the warning (using a policeman, a criminal, a traffic light), adding an extra step before the user makes the final decision, changing the style of the warning, and reusing the Firefox's SSL warning on Chrome. The results show that only reusing the Firefox's warning on Chrome led to significant improvement in the effectiveness of warning [106].

Scholars have also developed and evaluated different approaches to assist users in choosing *stronger passwords* [98, 108, 188]. Forget et al. proposed and evaluated a mechanism - persuasive text passwords (PTP) - to help users create stronger passwords [108]. PTP takes the password that the user initially created and suggests, as a replacement, a modified and stronger version of the original password. In a 83-participant lab study, the authors evaluated four different variations of PTP. The authors found that indeed PTP nudged participants to adopt stronger passwords with minimal usability effects [108]. Similarly, Ur et al. conducted online experiments to measure how strength meters affect password created by users [188]. In a two-part 2931-participant online study via AMT, the authors evaluated 14 different password strength meters. In the first part, they asked participants to play a role in which they need to create new passwords for their emails as their providers changed the passwords requirements. In the second part, participants were invited via email to login using their new passwords. The authors found that meters nudged participants to create longer passwords overall. However, only aggressive meters (e.g., requiring an additional class character or longer passwords) were effective in nudging participants to create stronger and harder to be guessed passwords [188].

2.5 Distinction From Prior Work

Prior work has focused primarily on designing app privacy nudges that help users make better privacy decisions when *choosing new apps to install* (install-time app privacy nudges) [85, 118, 134, 142, 162]. The install-time nudge works by informing users of privacy risks associated with the app that the user chooses to install and recommending a more privacy-friendly app than the one chosen. In experimental settings, install-time nudges are shown to be effective in informing users of privacy risks associated with apps that users want to install, and effective in influencing users to consider installing more privacy-friendly apps [85, 118, 134, 142, 162]. However, an install-time nudge can only be effective (a) if there is an alternative app that is more privacy-friendly than the app that the user chooses to install, or (b) if the nudge can persuade the user to refrain from installing the privacy-invasive app altogether, even without alternatives. To the best of our knowledge, outside the experimental setting there is no accessible way to find alternatives to privacy-invasive apps assuming alternatives exist at all. In addition, prior work has suggested that privacy notices are less likely to garner users' attention at install-time in comparison to notices at runtime (e.g., during or after using an app) [74], and that users are less likely to choose another app once they have made up their minds to install a given app [174]. Thus, if the user chooses to proceed and install the app, the install-time app privacy nudges lose their effect.

To extend previous work on mobile app privacy nudges, this thesis focuses on designing and evaluating runtime app privacy nudges to help users better manage their app privacy settings. Specifically, runtime app privacy nudges aim to (1) increase users' awareness about privacy risks associated with apps that users have previously installed, and (2) motivate users to review their app privacy settings and adjust them as needed. As such, runtime app privacy nudges do not burden the user to find alternative privacy-friendly apps. Instead, runtime nudges are intended to help the user to better manage the privacy settings for apps that the user has already installed and used. Furthermore, runtime nudges may also motivate users to uninstall privacy-invasive apps [65, 110].

Few studies have looked at runtime mobile app privacy nudges. Balebako et al. [73] proposed mobile privacy nudges based on apps' access frequency to specific data and evaluated them in a lab study. Although our work in Chapter 3 builds on the work of Balebako et al., our work differs in two dimensions. First, our work measures the effect of privacy nudges in triggering users to review and adjust their app permissions, whereas the work of Balebako et al. measured perception and feeling. Second, our study in Chapter 3 evaluates privacy nudges in situ, with participants using their own devices "in the wild", whereas Balebako et al. conducted a lab study, and participants did not use their own devices. Fu et al. [110] proposed a per app runtime location access notification and contrasted it with the existing location access disclosure method in Android (i.e. a blinking icon on the notification bar). Our work in Chapter 3 extends the work of Fu et al. 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 et al. focused only on transparency without providing users with tool to adjust their app privacy settings. Second, Fu et al. focused on location access, whereas we also examine phone contacts, calendar, and call logs.

Protect My Privacy (PMP) utilized crowdsourcing to recommend privacy settings (both granting and denying) for each app [63]. In contrast, our focus is on informing users of apps behavior instead of informing them of other users' privacy settings. Bal et al. designed and evaluated "Styx": a system that monitors apps' data collection practices and informs users of what these apps are able to infer based on the data they accessed and collected [71]. Our work in Chapter 4 builds on the work of Bal et al. by using inferences/consequences of data collected by apps to design effective runtime mobile app privacy nudges.

In contrast to [110] which nudges the user while using the app (app-usage dependent nudges), our work primarily focuses on privacy nudges that are independent of app-usage. An advantage of our approach is to also inform users of privacy risks associated with apps that access users' sensitive information in the background (i.e. while users are not actively using the app). Only 22% of mobile users are aware of apps' capabilities (and thus potential privacy risks) while apps in the background [182]. Nonetheless, future work should explore which approach is more effective and when.

In sum, this thesis attempts to extend previous work by introducing runtime mobile app privacy nudges, evaluating their effectiveness, designing effective content of privacy nudges, and evaluating approaches to keep users engaged with repeated nudges.

Chapter 3

Permission Managers & Mobile App Privacy Nudges

In this chapter, we focus on two research questions: (1) Is access to an app permission manager an effective way of helping users review and modify their app permissions? (2) Can app privacy nudges, that regularly alert users about sensitive data collected by their apps, enhance the effectiveness of an app permission manager? To address these 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 confirm that users are generally unaware of mobile app data collection practices. Our findings also 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.

This chapter makes the following contributions. First, our study highlights the benefits of a permission manager like AppOps and quantifies the additional benefits that result from supplementing such functionality with privacy nudges. Second, we designed mobile privacy nudges that increase user awareness of data collection practices and effectively motivate users to review and often revise their app permissions. Third, we evaluated the effectiveness of the privacy nudges in a real-world setting. Finally, we derive recommendations for the design of effective

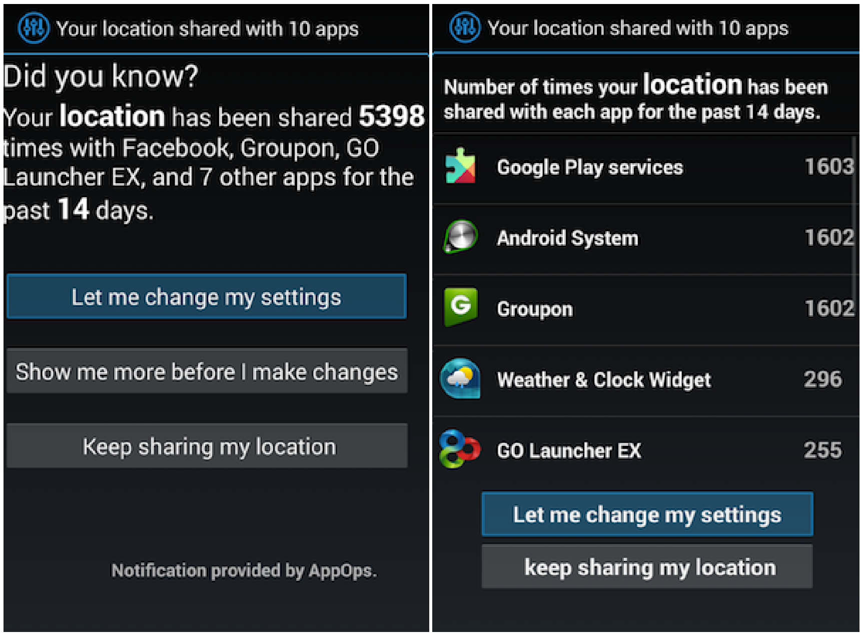

Your location shared with 10 apps

Did you know?

Your **location** has been shared **5398** times with Facebook, Groupon, GO Launcher EX, and 7 other apps for the past **14** days.

Notification provided by AppOps.

Your location shared with 10 apps

Number of times your **location** has been shared with each app for the past 14 days.

Google Play services	1603
Android System	1602
Groupon	1602
Weather & Clock Widget	296
GO Launcher EX	255

Figure 3.1: 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.

privacy nudges on mobile devices based on our results.

3.1 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.

3.1.1 Nudge Content

Prior work has shown that users are unaware of, and surprised by, apps' data access practices and frequency [73, 110, 127], which suggests nudging users to review their app permissions has utility. Therefore, we designed the mobile privacy nudge shown in Figure 3.1 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 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 3.1).

3.1.2 Nudge Response Options

The privacy nudge provides targeted response options to facilitate privacy management (see Figure 3.1). The “Let me change my settings” option opens AppOps directly. We hypothesized 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 3.1, which lists each app’s access frequency for the nudge’s particular information type, in descending order. The detailed report enables 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 [73].

The third option (“Keep sharing my [data]”) allows users to indicate the status quo is acceptable. Keller et al. [130] 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.

3.2 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 nudges. We implemented our privacy

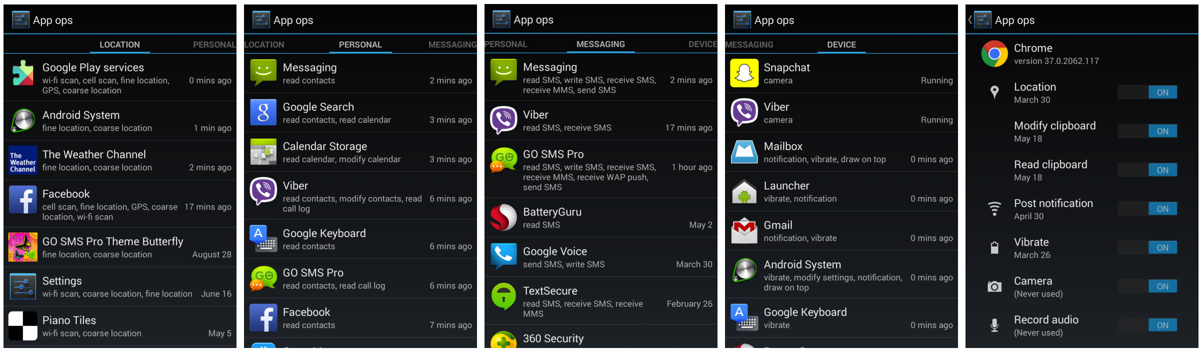

Google Play services — wi-fi scan, cell scan, fine location, GPS, coarse location — 0 mins ago
Android System — fine location, coarse location — 1 min ago
The Weather Channel — fine location, coarse location — 2 mins ago
Facebook — cell scan, fine location, GPS, coarse location, wi-fi scan — 17 mins ago
GO SMS Pro Theme Butterfly — fine location, coarse location — August 28
Settings — wi-fi scan, coarse location, fine location — June 16
Piano Tiles — wi-fi scan, coarse location — May 5

Messaging — read contacts — 2 mins ago
Google Search — read contacts, read calendar — 3 mins ago
Calendar Storage — read calendar, modify calendar — 3 mins ago
Viber — read contacts, modify contacts, read call log — 6 mins ago
Google Keyboard — read contacts — 6 mins ago
GO SMS Pro — read contacts, read call log — 6 mins ago
Facebook — read contacts — 7 mins ago

Messaging — read SMS, write SMS, receive SMS, receive MMS, send SMS — 2 mins ago
Viber — read SMS, receive SMS — 17 mins ago
GO SMS Pro — read SMS, write SMS, receive SMS, receive MMS, receive WAP push, send SMS — 1 hour ago
BatteryGuru — read SMS — May 2
Google Voice — send SMS, write SMS — March 30
TextSecure — read SMS, receive SMS, receive MMS — February 26
360 Security

Snapchat — camera — Running
Viber — camera — Running
Mailbox — notification, vibrate, draw on top — 0 mins ago
Launcher — vibrate, notification — 0 mins ago
Gmail — notification, vibrate — 0 mins ago
Android System — vibrate, modify settings, notification, draw on top — 0 mins ago
Google Keyboard — vibrate — 0 mins ago

Chrome
version 37.0.2062.117

Location — March 30 — ON
Modify clipboard — May 18 — ON
Read clipboard — May 18 — ON
Post notification — April 30 — ON
Vibrate — March 26 — ON
Camera — (Never used) — ON
Record audio — (Never used) — ON

Figure 3.2: AppOps has four different tabs: location, personal (e.g., phone contacts and calendar), messaging(e.g., SMS), and device (e.g., vibration, camera). Each tab lists all apps that accessed the corresponding category of data, e.g., location (*left*). Selecting a specific app opens a screen showing all permissions accessed by the app (*right*).

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.

3.2.1 The Permission Manager - AppOps

In Android 4.3, Google introduced a permission manager, called *AppOps*, that was hidden by default and required an external app to access. AppOps allowed users to selectively grant or deny permissions for installed apps. AppOps, shown in Figure 3.2, organizes 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. In Android 4.3, AppOps allows users to control 31 different operations [113]. Some of these operations map directly to one or more permissions (e.g., read contact control maps to "android.Manifest.permission.READ_CONTACTS" permission). Other operations do not map to any permission (e.g., the existence of audio control although accessing audio does not require a permission).

3.2.2 Implementation of Study Client

We designed a study client app to act as a launcher for AppOps and installed it on participants' phones. 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 logs show 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 easy access to specific interaction events and modifying AppOps would have required rooting and flashing participants' devices, which we deemed undesirable. 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.

3.2.3 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 affirm 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 look 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 between 11am and 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 [104], 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 ordering 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. All interviews were audio recorded and then coded with categorical codes (i.e. AppOps, the privacy nudges, reasons for app permission changes). We, then, conducted thematic analysis on each of these three categories.

3.2.4 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.

3.2.5 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. In addition, we acknowledge a potential self-selection bias in our participant pool. In spite of these limitations, we still observed an effect even with a smaller number of nudges; and therefore our findings still hold.

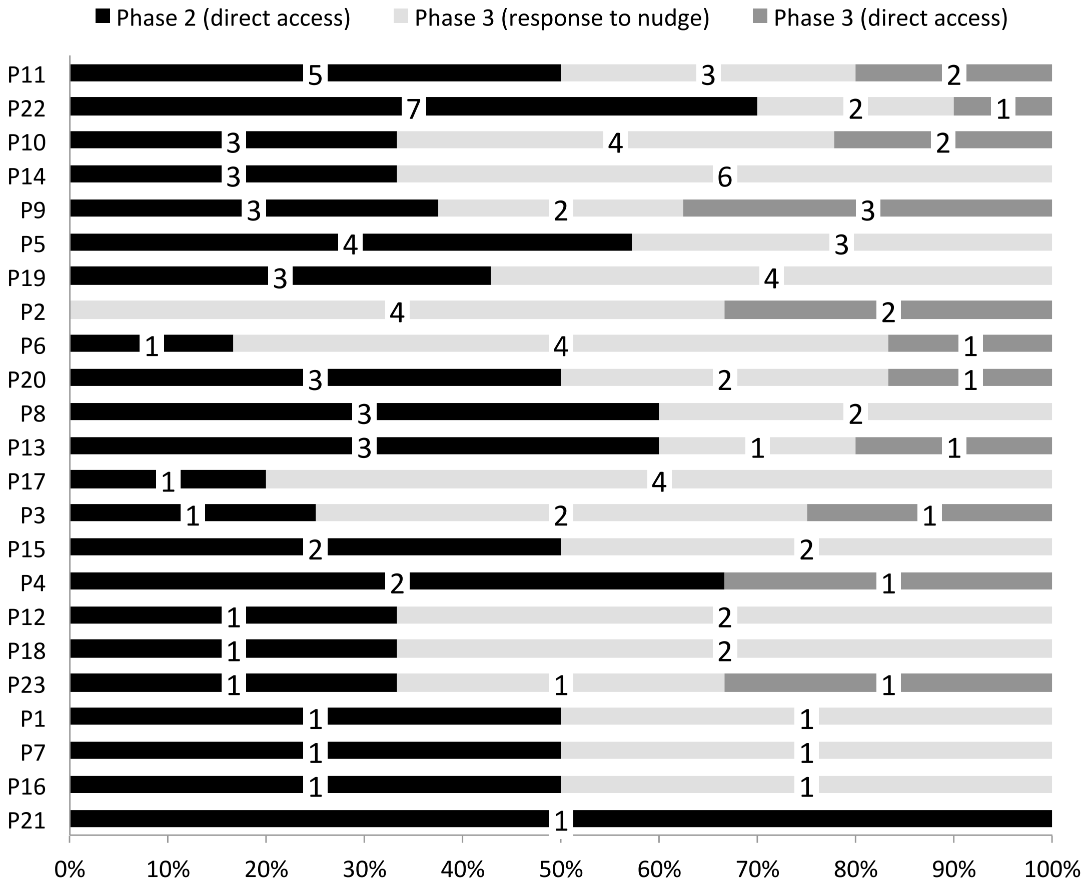

Figure 3.3: The number of times each participant reviewed their app permissions by opening AppOps in phase 2 & 3. Participants are ordered based on the frequency of reviewing app permissions. In phase 3, participants reviewed their app permissions 53 (77%) times in response to a nudge. The privacy nudges were the primary trigger for participants to review their app permissions in phase 3.

3.3 Results

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

3.3.1 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. Six (26%) of our participants are in information technology related fields and only one of them specializes in computer security. However, our results do not indicate that their background had a significant effect on their behavior. P2 never adjusted his app permissions throughout the study, P13 only adjusted the permissions in phase 2, P6&P7 only adjusted the permissions in phase 3, and P10&P12 adjusted their app permissions in both phases.

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 an app access to a permission, and (b) *permissive adjustments*, i.e., how often participants permitted an app access to a restricted permission.

3.3.2 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 3.3 shows, 22 participants (95.6%) reviewed their permissions at least once. Of those, 12 participants reviewed their apps' permissions multiple times. Only P2 did not review his permissions in phase 2.

Adjusting apps' permissions: As shown in Figure 3.4, 15 (65%) participants restricted 272 app-permission pairs for 76 distinct apps, including both participant-installed and pre-installed apps, see Figure 3.5. 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.

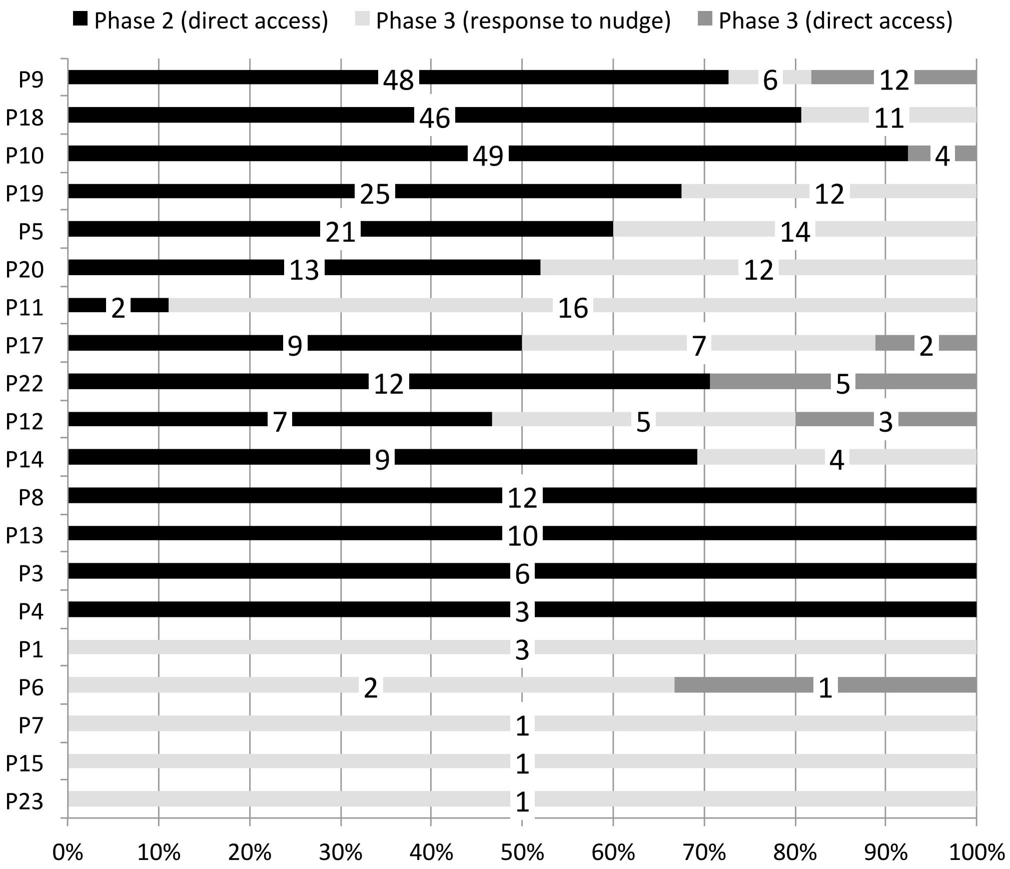

Figure 3.4: Number of permissions restricted by each participant. Participants are ordered based on (a) the restrictive adjustments they made either in phases 2 & 3, phase 2, or phase 3, and (b) the frequency of adjustments within each group. P2, P16, & P21 never adjusted their app permissions. Ninety-Five (78%) app permission adjustments in phase 3 were made in response to a nudge. The privacy nudges were the primary trigger for participants to adjust their app permissions in phase 3.

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

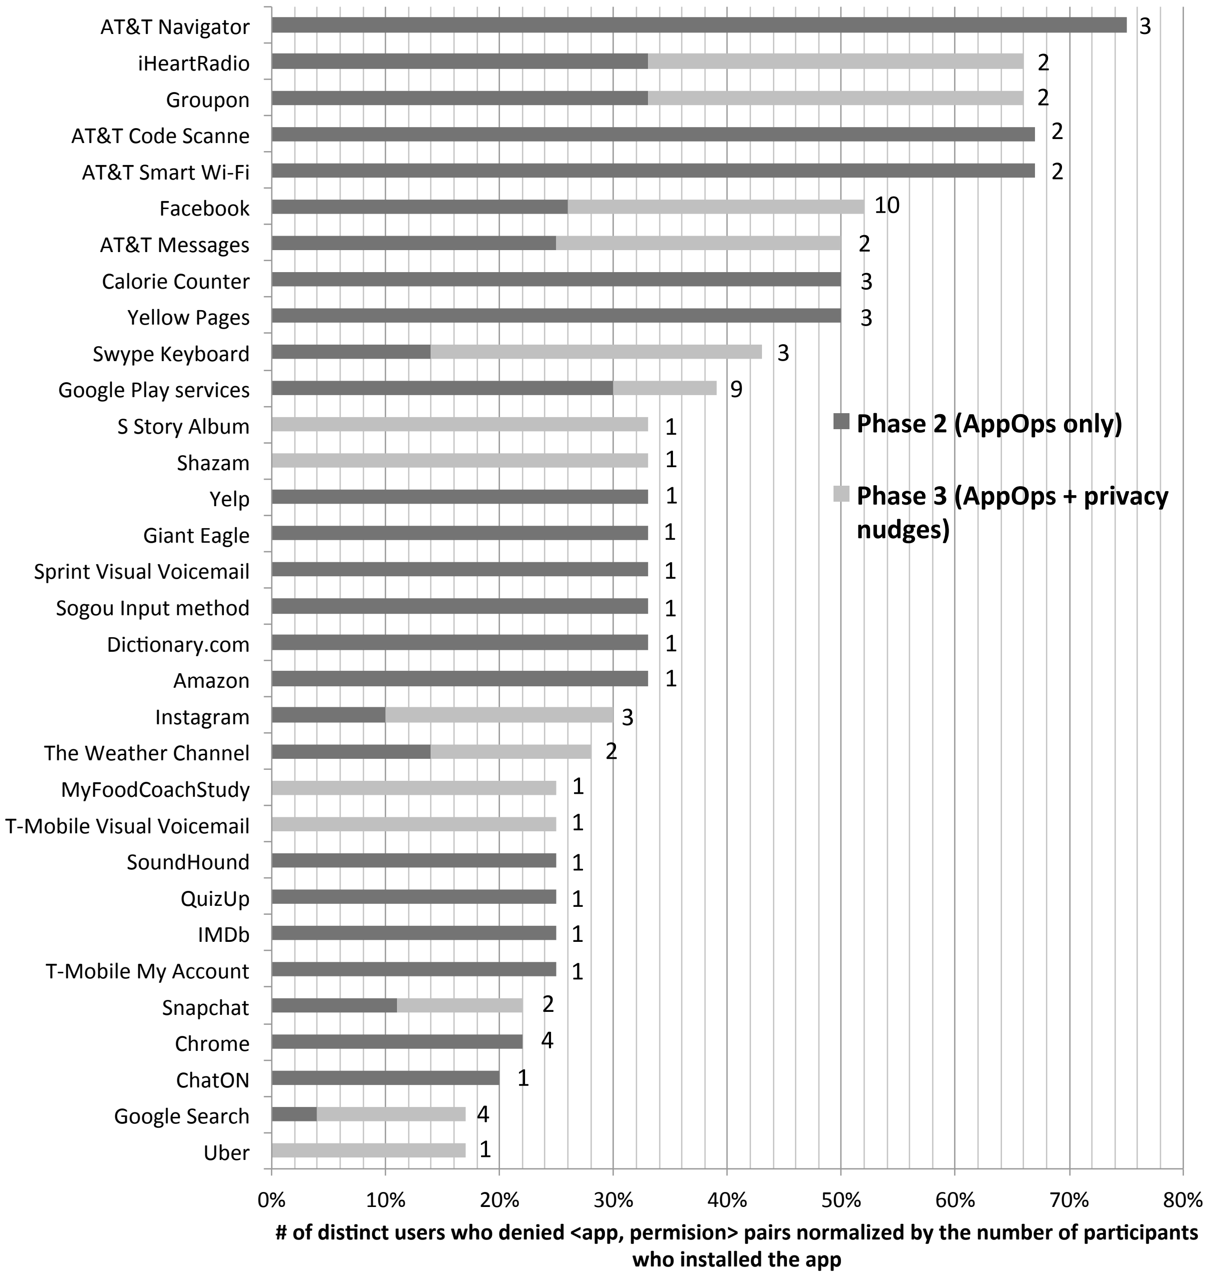

Figure 3.5: The apps that were revoked permissions in phase 2 & 3. The numbers to the right of the bars are the absolute number of distinct users who revoked at least one permission per app.

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.

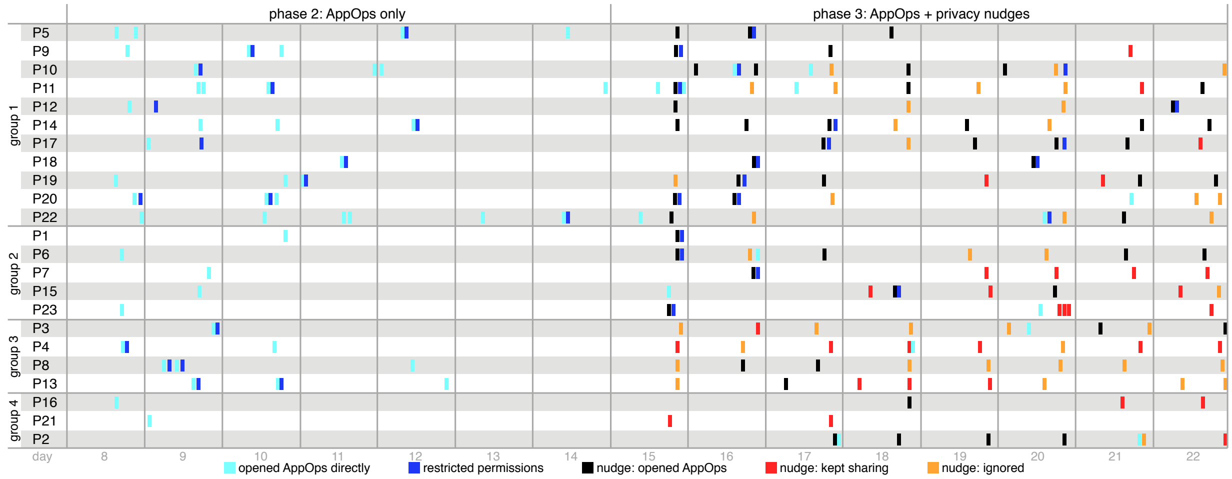

Figure 3.6: Timeline of participants' interactions with AppOps and the privacy nudges during phase 2 and phase 3. In the X-axis, each column represents a day during phase 2 (7 days) & phase 3 (8 days).

3.3.3 Effectiveness of AppOps with Privacy Nudges

The goal of our nudges was to get users to review their app 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 3.3 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 trigger for participants to review their apps' permissions.

Figure 3.6 shows that participants' interaction with AppOps declined sharply after day ten: 39% (day 10), 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 3.3). 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 3.4 shows that 16 (70%) participants restricted 122 app-permission pairs, 14 in direct response to nudges. Ninety-Five (78%) of the restrictive app permission 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 3.5. 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 the 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 five 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 four 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 three 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.

3.3.4 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 3.7 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." To mitigate the timing issue, future work can utilize frameworks such as Attelia [156] to find suitable times to deliver nudges to users.

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 helped 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 AppsOps 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 whether 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 the AppOps UI is structured. 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 3.2. 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 par-

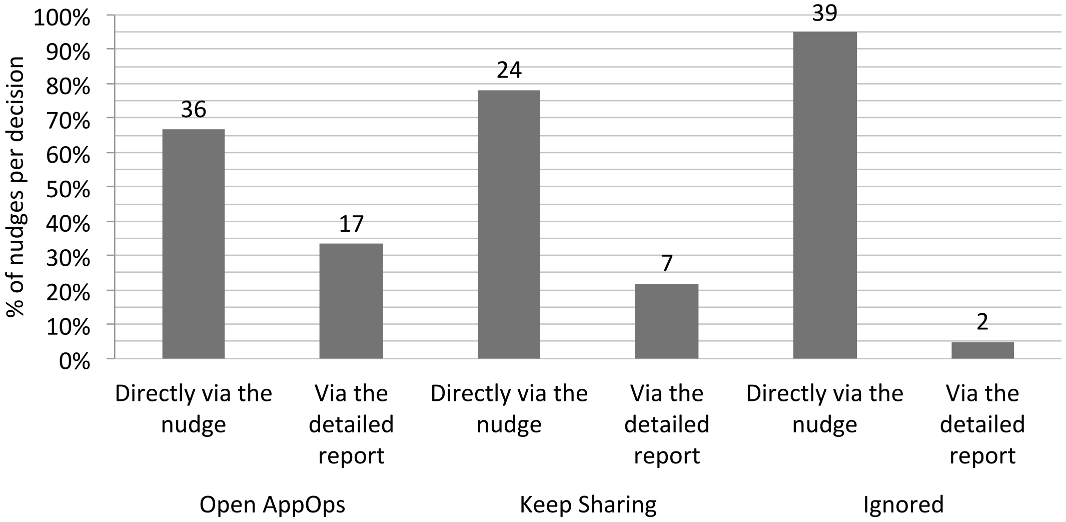

Figure 3.7: How participants responded to the nudges. The numbers on top of the bars are the absolute numbers.

ticipants 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 3.8 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 more pertinent nudges including only unidentified apps.

The detailed report: The main goal of providing a detailed report, as shown in Figure 3.1, 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 3.7.

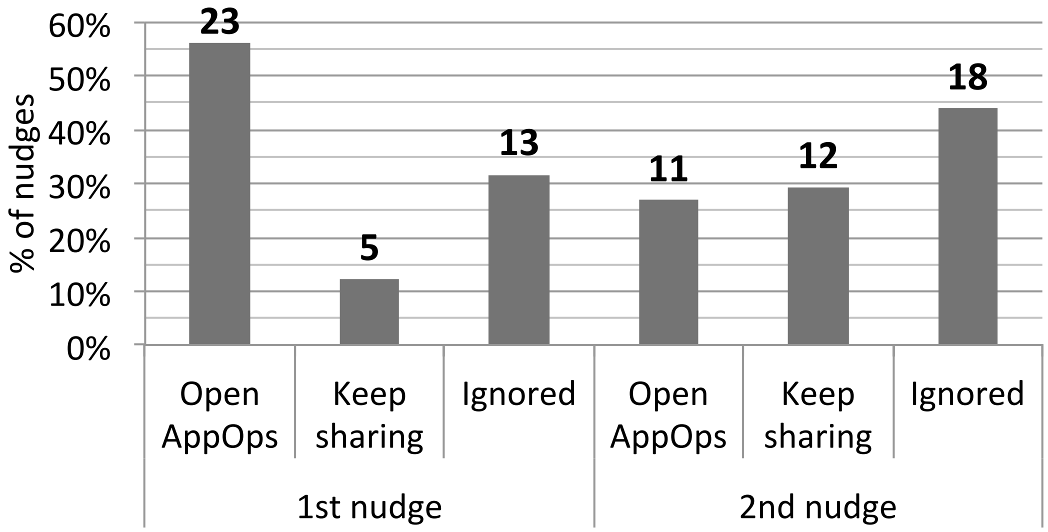

Figure 3.8: How participants responded to the 1st and the 2nd nudge of the same data type. The numbers on top of the bars are the absolute numbers.

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 adjustments as the detailed report provided a closer look at the data collection practices of individual apps.

Frequency of access: We explore whether an increase in frequency of personal information accesses by apps correlates with an increase in participants' likelihood to choose "let me change my settings." Using a random effects linear regression, 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.

During 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."

3.4 Discussion

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

3.4.1 Permission Managers Help

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 more privacy controls in iOS [190]. On the other hand, Google notoriously removed AppOps from Android phones in 2013 [94]. While no official explanation was provided, this move might be indicative of tensions between Google's advertising-based revenue model and its stated objective of giving users more control over their information.[1] Another possible interpretation might be that this move was motivated by usability concerns, whether concerns over the number of settings made available to users or concerns that users may deny permissions that are necessary for apps to function. Our study however showed that, despite the large number of settings, users generally benefited from having them and we did not see many instances of users complaining about their apps stopping to work.

3.4.2 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

[1] Analysis 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.

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 are starting to find their way into mobile platforms. Recently, iOS 8 introduced what can be viewed as a form of privacy nudge. Specifically, if a user allows an app to access her location even when not using the app (i.e. in the background), iOS will occasionally ask her whether she wants to continue doing so [48]. This approach is consistent with our findings, which show the benefits of combining permission manager functionality and nudges.

Chapter 4

Toward more Effective App Privacy Nudge Contents

Results from Chapter 3 indicate that, even with access to a permission manager, Android users have limited awareness of the privacy risks associated with their apps. Therefore, users do not take full advantage of the permission manager's functionality. Our results show that users benefit from receiving nudges that effectively: (1) increase users' awareness about privacy risks associated with apps they have previously installed on their devices (e.g., inform users of how frequently apps collect their information), and (2) temporarily make privacy the primary task for users and motivate users to review and adjust their app privacy settings as needed. Chapter 3 shows that the privacy nudges led all but two of our participants to review their app permissions at least once within eight days. More importantly, the privacy nudges led more than half of our participants to exercise control over the personal information that apps were allowed to access. In sum, the results from Chapter 3 suggest that mobile app privacy nudges are a promising approach to help users better manage their app privacy settings.

Nonetheless, it is not clear how effective different types of nudge contents are in motivating users to review and possibly adjust their app privacy settings. Should we nudge users by simply informing them that their personal information has been *recently* accessed and used by apps, which has been used in current mobile operating systems as shown in Figure 4.1? Is this the best nudge content to help users manage their app privacy settings given that mobile users have expressed their desire to take control of how apps access their personal information [135]? If not, as we hypothesize, which nudge content is more likely to motivate users to review and possibly adjust their app privacy settings? Should a nudge content inform users of how frequently their information has been used as we did in Chapter 3? Can we do better by utilizing attributes that users consider when making privacy decisions about mobile apps (e.g., purpose of accessing users' information as prior work has suggested [142, 173, 180])? Our goal in this chapter is to evaluate and compare the effectiveness of different app privacy nudge contents. This investigation is important for two reasons. First, privacy is typically a secondary task for users [79, 148, 198]. Therefore, we posit that we can only occasionally nudge users to switch their attention to privacy management. Thus, for such infrequent opportunities, we should use a nudge content that is more likely to grab users attention and to influence them to review their app privacy settings and ensure that these settings match their privacy preferences. Second, our results should assist mobile operating system providers and mobile privacy tool developers in rethinking the current nudge contents that they utilize.

The central research question in this chapter is: Which nudge contents are more effective in motivating mobile users to review and possibly adjust their app privacy settings?

4.1 Methodology

4.1.1 Research Questions

Both Android and iOS provide a run-time app privacy nudge which simply informs users that their location has been recently accessed by apps installed on their phones as shown in Figure 4.1. We refer to this nudge as the "recent access" nudge content. However, such a nudge content has not been evaluated and its effectiveness is not known to the best of our knowledge. Furthermore, we hypothesize that such a nudge content might not be effective because it provides very little information, which may not be sufficient for users to make informed privacy decisions. Therefore, we explore whether we can identify and design more effective nudge contents that help users to better manage their privacy on mobile devices. We posit that a nudge content is more effective than other nudge contents if it is more likely to motivate users to review and possibly adjust their app privacy settings. To this end, we focus on two research questions related to nudge contents.

1. How effective is the "recent access" app privacy nudge content in motivating users to review their app privacy settings in comparison to other nudge contents? The "recent access" nudge content simply informs users that their personal information has recently been accessed by apps on their phones.

2. Can we design nudge contents that are more effective than the "recent access" nudge content?

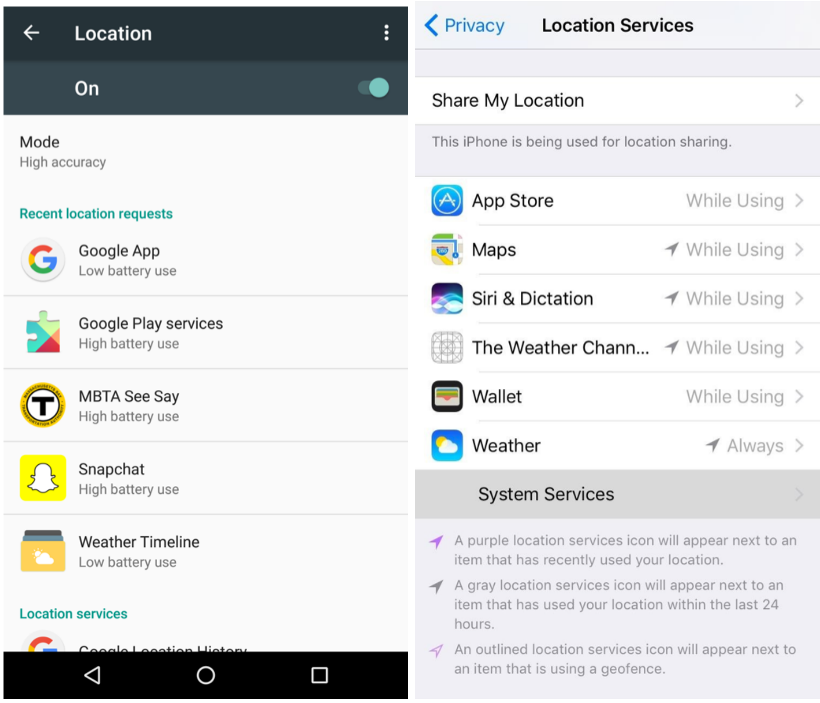

Location

On

Mode
High accuracy

Recent location requests

Google App
Low battery use

Google Play services
High battery use

MBTA See Say
High battery use

Snapchat
High battery use

Weather Timeline
Low battery use

Location services

Google Location History

(a) Android

Privacy | Location Services

Share My Location

This iPhone is being used for location sharing.

App Store	While Using
Maps	While Using
Siri & Dictation	While Using
The Weather Chann...	While Using
Wallet	While Using
Weather	Always

System Services

A purple location services icon will appear next to an item that has recently used your location.

A gray location services icon will appear next to an item that has used your location within the last 24 hours.

An outlined location services icon will appear next to an item that is using a geofence.

(b) iOS

Figure 4.1: Both Android and iOS provide features to enable users to identify apps that have recently accessed users' location information. Android (left) shows the list of apps that have made recent location requests. iOS (right) uses a purple arrow to identify apps that recently used users' location information, and uses a gray arrow to identify apps that used users' location information within the past day. We refer to these features as the "recent access" app privacy nudge content. The baseline nudge content in our experiment is inspired by these "recent access" nudge contents in Android and iOS.

4.1.2 Nudge Contents

We hypothesized a nudge content is effective if it can garner the user's attention, make privacy the primary task for the user by switching the user's attention to privacy management, and motivate the user to review and possibly adjust their privacy settings to better match their preferences. To this end, we designed nine nudge contents that inform mobile users of : (1) apps' data collection practices, (2) purposes of data collection by apps, and (3) potential implications of reusing collected data.

(1) Data Collection Practices

Both Android and iOS inform users that their location information has been recently accessed by apps installed on their phones. Android shows its users the list of apps that recently (i.e.

within the past 15 minutes [7]) requested access to users' location information (Figure 4.1 (a)). Similarly, iOS shows the list of apps that are allowed to access the user's location in the privacy settings. Next to each app, a purple arrow refers to an app that has recently accessed the user's location, and a gray arrow refers to an app that had accessed the user's location within the past day (Figure 4.1 (a)). The nudge content of our first (and baseline) nudge is inspired by the these "recent access" nudge contents in current mobile operating systems. The baseline nudge content simply informs users that their personal information has recently been accessed by some apps on their phones. The exact wording of the baseline nudge is shown in Table 4.1.

Two other data collection practices by apps that may be less known to users are: how frequently apps access their personal information and whether apps can access users' information even when users are not actively using the apps.

Previous work including Chapter 3 has shown that apps access users' location in high frequency but users are unaware of and surprised by such practices [73, 110, 127]. Additionally, as Chapter 3 shows, a nudge content based on how frequently apps access users' information effectively motivates the majority of our participants to review their app privacy settings. Therefore, we chose this nudge content for the second condition in our experiment. The exact wording of the nudge content is shown in Table 4.1.

Furthermore, apps on mobile devices are capable of running and accessing users' personal information even when users are not actively using the apps (e.g., through services [9]). Wijesekera et al., has shown that about three quarters of requests to sensitive resources (including users' information) occur while the user is oblivious and not actively using the apps [194]. However, the majority of users indicated that they would have denied at least one of the access requests that happened in the background [194]. Similarly, Micinski et al., has shown that the majority of access requests for location information in particular occurs either while the user is not actively using the app nor actively using any functionality that needs location information [149]. Although users may have previously authorized access to information while actively using the app, they may not expect the app to access the information again in the background [149]. Additionally, a similar nudge content found its way to the industry as Apple has recently started notifying iOS users when their location is accessed in the background [48, 52]. This suggests that informing users about apps accessing their information in the background is potentially a promising nudge content. Therefore, we designed a nudge content for our third condition to inform users of this practice. The wording of the third nudge nudge is shown in Table 4.1.

(2) Purposes of Data Collection

Prior work has suggested that the purpose of data collection is an important factor that users consider when making privacy decisions about mobile apps [134, 142, 162, 173, 180]. However, users' ability to identify why their information has been accessed by apps is limited, especially when their information is used for purposes other than providing the apps' main functionality [142]. By highlighting why apps access users' information, we may help users make better privacy decisions [142, 173]. Thus, we see potential value in designing nudge contents that highlight purposes not related to apps' main functionality. Amongst these unrelated purposes, users are particularly concerned about using their personal information for advertising [142, 143, 173]. Therefore, we decided to design two variations of nudge contents related to the purpose of data collection: a generic and a specific. The generic variation informs users that apps have used their information for purposes other than providing a main function. The specific nudge content additionally provides a concrete example of such an unrelated purpose: location-based advertising. The exact wording of these two nudge contents are shown in Table 4.1: (4) & (5), respectively.

(3) Potential Implications of Reusing Collected Data

Personal information which is collected for a particular purpose (e.g., facilitating a transaction between the user and the service provider) can be later reused to infer additional and unknown characteristics of users [57, 71, 118]. Such inferred characteristics can then be used to build profiles and may be used in discriminatory judgmental transactions against users [20, 34, 35, 91, 117]. For example, Flurry, the popular provider of mobile analytics and advertising solutions, has a product called: "Flurry Personas" which uses behavioral information (e.g., apps that the user has installed) to profile and categorize users [20]. Examples of those profiles include: "High Net Individuals", "Slots Players", "New Mothers", and "Mobile Payment Makers" [20]. In light of these potential implications, privacy guidelines (e.g., FIPPs and OECD) [18, 29] and standards (e.g., HIPAA and COPPA) [15, 155] have been explicit about the importance of limiting secondary usage of personal information, of informing users of such practices, and of seeking their consent. Furthermore, the practices and implications of reusing collected data are generally opaque to users due to the complexity of the information broker ecosystem [109]. We hypothesized that bringing such practices to users' attention will increase users' awareness and will motivate them to review their app privacy settings. Therefore, we decided to design nudge contents that inform users of potential privacy sensitive implications of how information collected by apps might be reused.

Potential implications of secondary usage of users' information are unlimited. However, we restricted ourselves to practices that have been documented by prior work and are reasonably plausible. Such a limit will assist us in reaching balanced nudge contents that inform users of potential (sensitive) implications of reusing their personal information without getting into the game of scaring users, which may have an unpredictable (short and long term) impact on users' trust in nudge contents (e.g., accused of exaggeration in the short term and crying wolf effect in the long term).

We decided to design four nudge contents about potential implications of reusing users' information: a generic nudge content and three specific nudge contents. In the generic nudge content, we inform users that apps are able to infer additional details about them based on their location information, but we do not provide examples of such inferred details. The goal is to examine whether such a generic nudge content would be sufficient to motivate users to review their app privacy settings. Additionally, the generic content will allow us to see if providing particular examples in the specific nudge contents (as we will see shortly) affects how users respond to nudges. For the three specific nudge contents, we similarly inform users that apps able to infer additional details about them based on location information but we provide examples of such details. In particular, we inform users that apps may be able to (1) infer where users live, (2) infer where users live and then use such information to predict users' income, and (3) infer where users live and then use such information to predict users' income which may make users subject to price discrimination[1]. As previous work has suggested, users are concerned about their home location [187], they only share it with a limited group of people (e.g., family and friends) [86], and they are less willing to share such information with advertisers [132]. However, prior work has shown that inferring where users' live based on location traces is attainable (e.g., [120, 138, 141]) and prevalent mobile services such as Google Now predict one's homes based on location history [26]. Given the sensitivity of this information and the ability to infer it from location traces, we hypothesized that a nudge content that highlights such an ability will likely motivate users to review their apps' permission settings. The exact wording of the nudge is shown in Table 4.1.

To build on this nudge content, the second and third nudge contents informs users of how such an inferred characteristic (e.g., where they live) can be used to predict an unknown demographic information (e.g., their level of income) and of how such a prediction may be exploited to affect users' experience (e.g., make users subject to price discriminatory transactions), respec-

[1] We use price discrimination here to refer to practices that are more generic than the strict definition of price discriminatory (i.e. showing the same product with different prices for different populations). These practices include tailoring the list of products in a website or a mobile app so that pricey products are presented to certain populations (e.g., affluent users) and at the same time other less expensive products are made harder to find (e.g., by increasing the number of steps needed to find such products) [147].

tively. Predicting demographic information based on previously collected personal information has been explored by prior work (e.g., [78, 89, 93, 112, 172, 199]). In particular, prior work has shown that users' income can be predicted based on browsing history [112, 151], social networks (posts [32, 164] or users whom you follow on Twitter [89]), shopping data [191], and installed apps on users' phones [36, 146]. Additionally (as mentioned earlier), analytics and advertising companies, such as Flurry, identify affluent users and assign them to specific profiles (e.g., "High Net Individuals") [20]. As such, the second nudge content informs users that apps, which have accessed their location, can infer where they live and then predict their income. We choose income particularly for its sensitivity [88, 163] and because we want to utilize it for introducing price discrimination in the third nudge content. In the third nudge content, we took a step further and designed it to inform users that a prediction of their income can be exploited to make users subject to online price discrimination. Research has shown that online services actually engage in price discriminatory practices based on previously collected personal information [147, 189]. For example, it has been reported that Staples[2] changes prices of items based on users' location and whether there is a Staples' competitor near by the user's location [189]. Similarly, Orbitz[3] profiles visitors, who are Mac users (i.e. using pricey products from Apple), as affluent and in turn shows them deals with higher prices [147]. The third nudge content is based on these documented practices and the exact wording of the nudge content is shown in Table 4.1. Although the nudge content was designed to motivate users to consider privacy and review their app settings, the content was carefully and iteratively worded to be as neutral as possible. In the pilot testing, some users predictably pointed out that they responded to the nudge because they object to price discrimination, whereas others perceived tailored prices as a positive practice. Furthermore, the three nudge contents were intentionally designed so that each nudge content provides incremental implications to the previous one which allowed us to design comparable nudge contents. It is important to note that we do not claim that apps or advertising companies engage in the exact practices that the nudge contents describe. Rather, the nudge contents are inspired by similar documented practices and has been designed to show potential (negative) implications of reusing users' location by apps.

4.1.3 Designing Nudges

As Figure 4.2 shows, each nudge consists of the location icon, the title, the main content of the nudge, the list of apps, and the buttons to act upon the nudge. In the experiment, only the main content of the nudge was replaced by one of the nine nudge contents shown in Table 4.1. Other

[2] https://www.staples.com

[3] https://www.orbitz.com/

Condition	**Nudge Wording**
(1) Location access (Baseline)	These apps have **accessed** your location in the past week.
(2) Frequency of access (Frequency)	These apps have accessed your location **1,865 times** in the past week.
(3) Access while not using the app (Background)	These apps have accessed your location in the past week although **you did not use them**.
(4) Access not to provide main function (Purposes)	These apps have accessed your location in the past week for purposes **not related to the apps' main function**.
(5) Condition (4) + an example, location-based ads (Purposes+Example)	These apps have accessed your location in the past week for purposes **not related to the apps' main function, such as location-based advertising**.
(6) Potential inferences (Inferences)	These apps have accessed your location in the past week. With this information, apps can **infer additional details about you**.
(7) Condition (6) + an example, inferring the address where you live (Inferences+Example)	These apps have accessed your location in the past week. With this information, apps can **infer additional details about you, such as the address where you live**.
(8) Condition (7) + a prediction example, predicting your income (Predictions)	These apps have accessed your location in the past week. With this information, apps can **infer additional details about you, such as the address where you live, and use it to predict your income**.
(9) Condition (8) + potential implications of prediction (Predictions+Implications)	These apps have accessed your location in the past week. With this information, apps can **infer additional details about you, such as the address where you live, and use it to predict your income. Knowing your income can affect prices and discounts you see in ads**.

Table 4.1: This table shows the wording of each nudge.

components remained consistent.

The location icon informs the user immediately that the screen is about location information. We used an icon with which users are familiar [162] and that has been used by maps and navigation applications in both Android and iOS.

The title of the nudge reads: "Your location accessed." The title is grammatically condensed by removing auxiliary verbs (i.e. "has been"). This was a deliberate decision to ensure that the title is short and fits in one line. In the pilot testing, we evaluated the title closely and found that users were able to easily comprehend its meaning.

The main content of the nudge matches one of the nine conditions as shown in Table 4.1 that the respondent is randomly assigned to. The words and phrases in bold in Table 4.1 are highlighted in blue on the nudge screen to guide the respondent's attention to the essential part of the nudge content. We used blue color for highlighting because it has been shown to be a neutral color [175] and thus we avoided introducing an additional confounding factor.

The three apps are listed in alphabetical order. Each row shows the app's icon and the app's name. This resembles how apps are typically listed in both Android and iOS settings.

The respondents are given two buttons that correspond to two options: (1) "Restrict which apps get my location" to enable respondents to review and possibly adjust their app privacy settings in response to the nudge, and (2) "Keep sharing my location with the apps" to keep things as is and close the screen. The captions of the two buttons are intentionally long because we want to ensure that respondents can easily comprehend what the buttons do. Furthermore, we used "enhanced active choice" when we chose the wording of the second button ("Keep sharing my location with the apps"). Keller et al. [130] recommend employing enhanced active choice to emphasize the desired option (button 1 in our design) by "highlighting the losses incumbent in the non-preferred alternative [button 2 in our design]."

If the respondent is an Android user, the nudge is presented using Nexus 5 as shown in Figure 4.2. If the respondent is an iOS user, the nudge is presented using an iPhone 6 frame as shown in Figure 4.2. The dimensions of the two frames are close to each other. However, iPhone six is slightly taller and Nexus 5 is slightly wider. We make the frames consistent across all respondents of an operating system to avoid introducing any biases related to the length or width of a frame.

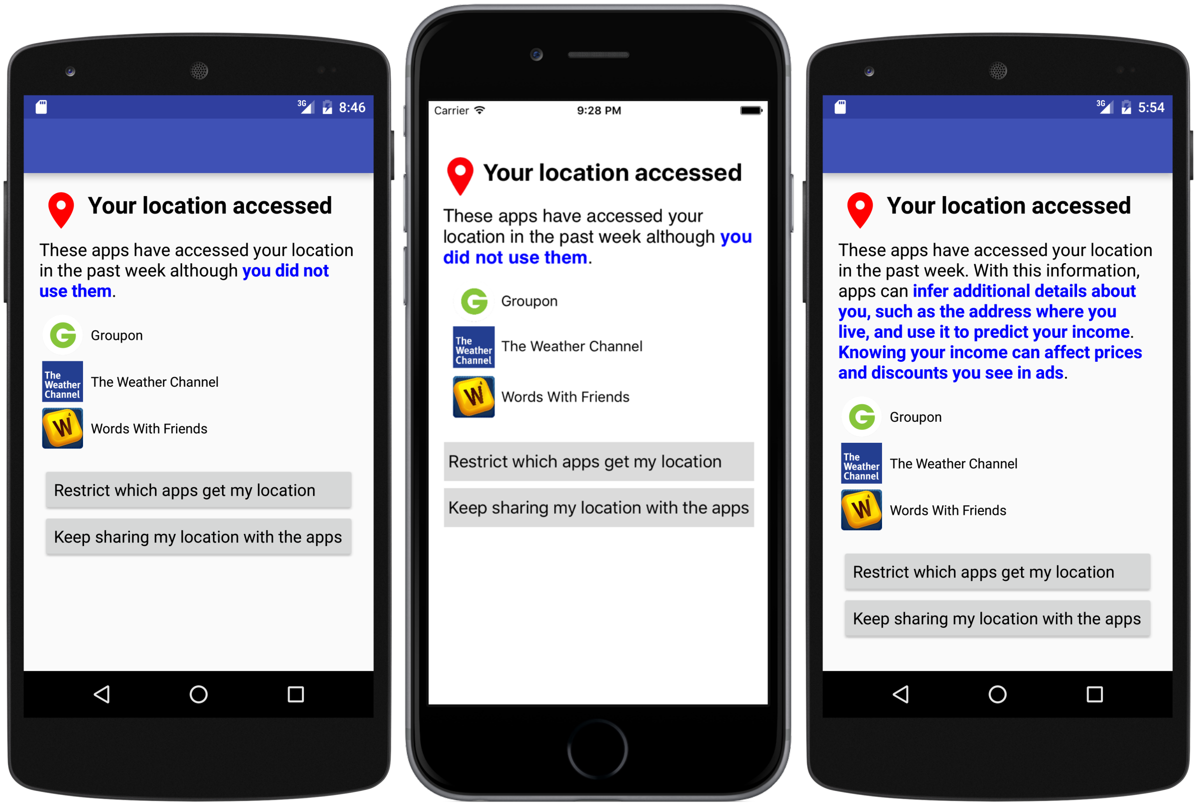

8:46
Carrier
9:28 PM
5:54

Your location accessed

These apps have accessed your location in the past week although **you did not use them**.

Groupon
The Weather Channel
Words With Friends
Restrict which apps get my location
Keep sharing my location with the apps

Your location accessed

These apps have accessed your location in the past week although **you did not use them**.

Groupon
The Weather Channel
Words With Friends
Restrict which apps get my location
Keep sharing my location with the apps

Your location accessed

These apps have accessed your location in the past week. With this information, apps can **infer additional details about you, such as the address where you live, and use it to predict your income**. **Knowing your income can affect prices and discounts you see in ads**.

Groupon
The Weather Channel
Words With Friends
Restrict which apps get my location
Keep sharing my location with the apps

Figure 4.2: Examples of nudge contents used in our experiment. A nudge content about apps accessing the users' personal information even if the users did not use the apps (*left: Android, middle: iOS*). Another nudge content about potential implications (e.g., price discrimination) of apps accessing and reusing users' location information (*right: Android*). We used a Nexus 5 frame for Android nudges and an iPhone 6 frame for iOS nudges.

Location Information

A large body of prior work has shown that users are typically concerned about sharing their location information (e.g., location-sharing services [86, 187] and mobile apps [81, 104]). However, location information is one of the most frequently requested piece of information by apps [102, 135] and is frequently used for advertising purposes [177]. As such, we decided to design our nudge contents around location information. Furthermore, we focused only on one type of information (i.e. only location information) rather than multiple types for a practical reason: to ensure that we have a manageable number of conditions in our experiment.

Apps Selection

We selected three popular apps because we wanted the majority of our respondents to be familiar with these apps and have used or are still using them. The goal is to minimize the effect of how unknown apps influence respondents' decisions to interact with nudge contents. In addition, we selected apps that request access to users' location information but are flexible enough to also function without location information or with alternative and less privacy sensitive information (e.g., a zip-code or a city-level instead of an exact location). This flexibility enables users which are especially privacy conscious or advanced users to continue utilizing apps' functionality without privacy loss. The goal is to minimize how trade-offs between usability/utility and privacy affect respondents' decisions to interact with the nudges. Furthermore, the apps that we selected collect location information for advertising purposes at minimum but may also use location information to provide apps' functionality. We selected three apps: The Weather Channel, Groupon, and Words With Friends. All three apps have shown between 50M-100M installs on the Google Play app store [27, 39, 40]. The three apps request access to users' location information but manual inspection of these apps shows that they can also function without location information or with alternative and less privacy sensitive location information (The Weather Channel and Groupon can function with zip-code or city-level; Words With Friends does not use location to provide any major functionality and works without location information). Words With Friends accesses users' location for advertising purposes [41]. The Weather Channel accesses users' location for both advertising and weather forecasting [99]. Similarly, Groupon accesses the users' location for advertising and to show near-by coupons. The Weather channel, Groupon, and Words With Friends belong to three different categories: weather, shopping, and games, respectively. For the frequency condition, we calculated the frequency of access per app as follows. We used averages of how frequently The Weather Channel and Groupon accessed location information from the first study (in Chapter 3). For Words With Friends, manual inspection shows that the app accesses location information per session. To match the level of frequency for the other two apps, we calculated the frequency for Words With Friends by assuming that the user is an active game player (approximately 5 sessions each two days [21].

We selected three apps rather than just one app because we wanted to minimize the extent to which one app may influence (due to trust or familiarity) users' decisions to respond to nudge contents. We also did not select more than three apps to avoid overwhelming the users with too many choices. We believe that the three apps strikes a good balance.

4.1.4 The Main Task

The main goal of our experiment is to test whether mobile users will be motivated to review and possibly adjust their app privacy settings when presented with different privacy nudge contents. Thus, the main task shows users one of the nine nudge contents and records how users responded to their nudges. To this end, we instructed respondents to imagine the following scenario: "*Imagine that you have* ***installed an app*** *on your phone. The app occasionally* ***informs you about the performance and behavior of the other apps installed*** *on your phone.*" The goal of describing a scenario in the introductory prompt is to provide a context for respondents and to establish a trust between respondents and the app that shows the nudges. Next, the respondent is presented with this prompt: "Now imagine that this app shows you the screen below" and a screenshot (as shown in Figure 4.2) that matches both the condition the respondent has been randomly assigned to, and the self-reported mobile operating system that the respondent uses. At the same time, the respondent is presented with two questions. First, "What would you do?" with four choices: (a) press the "Restrict which apps get my location" button, (b) press the "Keep sharing my location with the apps" button, (c) close the screen (by pressing the "Home" or the "Back" buttons[4]), and (d) I do not know. Second, "Why?" with a short-essay.

If the respondent chose options (b), (c), or (d), the main task is completed. If the respondent chose option (a), the survey asked follow-up questions to explore what app privacy settings the respondent would change.

First, we want to examine whether the respondent indeed wanted to adjust their app privacy settings and what those settings were. Thus, the respondent is prompted: "After pressing the "Restrict which apps get my location" button, you are presented with the screen below." (as seen in Figure 4.7). The survey asked "What would you do?", and presented three options: (a) Adjust

[4] In the iOS version, it reads: by pressing the "Home" button. iOS does not have a back button. This difference holds for other questions that contain a similar option.

apps' access settings, (b) Close the screen (by pressing the "Home" or the "Back" buttons), (c) I do not know.

If the respondent chose option (a), the respondent was asked "What apps' access settings would you adjust?" and presented with the list of three apps with the option to allow or deny each one of them. The screenshot with the privacy controls (as seen in Figure 4.7) is also presented to the respondent for reference.

After completing the main task, the respondent is asked follow-up questions to further determine what factors influenced the respondent's decision. The details of the survey is described in Section 4.1.6.

4.1.5 Experiment Design

We chose a between-subject design for our experiment. Each respondent was randomly assigned to one of nine conditions and was asked to complete the main task (as described in Section 4.1.4) only once. We chose between-subject design instead of within-subject design to specifically avoid the learning and order effects [139]. Such effects will strongly biased the results in ways we deemed unacceptable. Additionally, a within-subject design will require respondents to complete the main task more the once and perhaps as many times as the number of conditions in our experiment. Thus, respondents may exhibit fatigue which will additionally bias their answers [139].

One limitation of the between-subject design is potential variance within each group. Such a variance is anticipated in privacy experiments because it has been shown that users exhibit diverse privacy preferences (e.g., [143, 195]). We mitigated this limitation by increasing the number of respondents per condition (we had approximately 100 respondents per condition) and by randomly assigning respondents to one of the conditions in our experiment.

4.1.6 Survey Walk-through

We created 18 versions of our survey tailored for the nine conditions and the two mobile operating systems: Android and iOS. All versions of the survey were identical in structure and wording of both questions and prompts. However, when respondents were presented with the main task (see Section 4.1.4), they were shown a screenshot that matches the condition that they were randomly assigned to, and the self-reported mobile operating system that they use. Next, we describe in detail the steps to complete the survey.

1. Mobile operating system. The survey asked “What is the operating system of your phone?” We listed five mobile operating systems: iOS, Android, Windows phone, Blackberry. We also gave the respondent the option to provide her own answer or to say “I do not know.” If the respondent indicated she uses Android or iOS, the survey asked about the phone device and the version of the operating system. Respondents were given a link which shows the needed steps to identify the version of the operating system of their phone.

2. The main task as described in Section 4.1.4.

3. Likelihood, Awareness, and Concern. Next, the survey asked three consecutive 5-point likert-scale questions: (1) “How likely or unlikely is it that the following occurs:”, (2) “Before this survey, to what extent were you aware or unaware about the following (assume the statement is true):”, and (3) “To what extent are you concerned or unconcerned about the following (assume the statement is true):”. With each question, the respondent is presented with the the textual content of the nudge they were assigned to. We asked these questions because we hypothesized that these three factors (i.e. Likelihood, Awareness, and Concern) may affect how respondents choose to respond to different nudge contents.

4. Familiarity with apps on the nudge. We asked “How familiar or unfamiliar are you with the following apps?” We listed the three apps shown in the nudge and give the respondent four options: “I have never heard of the app”, “I have heard of the app but I have never installed it on my phone”, “I had installed the app on my phone at some point in time but it is not on my phone anymore”, or “I have installed the app on my phone and it is still on my phone”. If the respondent chose the last option, we asked: “How often do you use the following apps?” and provided the following options: “Never”, “Rarely”, “Sometimes”, “Often”, or “Always.” These two questions will enable us to examine whether familiarity and app usage affect respondents decisions to respond to different nudge contents.

5. Familiarity with privacy controls on the phone. The survey asked “Can you selectively allow or deny individual apps installed on your phone from accessing your location?” and showed four options: “No, I cannot”, “Yes, I can but I do not know how”, “Yes, I can and I know how”, or “I do not know.”. If the respondent chose the third option (“Yes, I can and I know how”), we further asked “Please briefly describe the steps that you can take to selectively allow or deny individual apps installed on your phone from accessing your location?” Next, we asked “What factors do you consider when allowing or denying individual apps installed on your phone from accessing to your location?” We gave four 6 options: “How frequently I use the app”, “How much I trust the company/developer of the app”,

"Why the app accesses my location", "Whether I use the app's functionality that needs my location", "others", and "No particular reason." These question help us understand whether respondents are aware of privacy controls available on their smartphones.

6. Privacy scale and demographics. Finally, we asked respondents to complete the Internet Users' Information Privacy Concerns (IUIPC) and demographic questions. We will use these questions to analyze how respondents' privacy concern level and different demographic attributes correlate with what options respondents chose when interacting with different nudges.

4.1.7 Recruitment

We recruited participants from Amazon Mechanical Turk (AMT) by posting a Human Intelligence Task (HIT) entitled: "Performance and Behavior of Mobile Apps (Android/iPhone Survey)." Participants were required to be 18 or older, live in the United States, have 95% approval rate for all previously completed tasks, and have completed at least a 100 tasks. Participants were instructed to complete the HIT only once and were provided with a link to check whether they have already participated. Any additional submissions by respondents were rejected. Furthermore, we asked two test questions to increase the quality of the results: (1) right after completing the main task (see Section 4.1.4), respondents were asked: "As shown in the previous screen, what types of personal information were accessed by the three listed apps?" and were given five options with only being correct: "Location information", (2)"Which of the following apps were listed on the screen shown in the previous question?" and were given six options and three of them were correct. We rejected respondents who answered the first test question incorrectly. We accepted respondents who correctly chose at least two of the three options in second question. We compensated participants $2.5 for a 20-minute survey which is slightly above the minimum federal wage.

4.1.8 Limitations

Our experiment has a number of limitations that we want to acknowledge. As typical with online experiments, respondents interacted with privacy nudges in unrealistic settings. There is no real privacy risk on respondents because they are not using their real devices and thus respondents may not have felt the need to act in a privacy conscious manner. However, the goal of the experiment is not to evaluate the absolute effectiveness of nudge contents. Rather, we focused on how effective nudge contents are in relevant to the baseline and to one another. As such, the bias of the unrealistic setting applies to all conditions equally.

Additionally, privacy is typically a secondary task to users in day-to-day life [79, 148, 198]. However, privacy was a primary task in our experiment. Respondents interacted with active privacy nudges that enforced them (by design) to stop whatever they were doing and interact with the nudges in order to proceed. However, passive privacy nudges are the dominant approach in current mobile operating systems. Therefore, privacy nudges in our experiment may have been more effective than than if respondents were to interact with passive nudges as prior work has suggested [97, 197].

Furthermore, only three apps were listed on the nudge in our experiment. Although the apps were carefully selected, these apps (collectively or individually) may have had an effect (positive or negative) on how respondents decided to interact with nudge contents. However, our analysis shows no significant correlation between how familiar respondents were with those apps and how they interacted with the nudges. Nonetheless, the effect of not listing other apps (other than these three apps) cannot be measured and it is a limitation of the design of our experiment.

Another limitation in our experiment only pertains to iOS respondents. The recent iOS privacy controls for location information is different than the privacy settings screen as shown in Figure 4.7. iOS privacy settings for location information give three options: Always, While using, never [51]. However, the privacy settings screen in our experiment gives the user only two options: allow or deny. We decided to offer only two options to ensure that both iOS and Android respondents have the same options. Additionally, the app settings is not the primary focus of our experiment.

Finally, we want to acknowledge that self-selection bias [119], which is typical in AMT experiments, is also a limitation in our study.

4.2 Results

We first describe respondents' demographics including their privacy concern level through the IUIPC scale. Then, we report how respondents interacted with different nudge contents with emphasis on the percentage of respondents per condition who became motivated to review and adjust their app privacy settings in response to different nudge contents. Furthermore, we look into factors that may have influenced respondents decisions when interacting with nudges. Finally, we report qualitative analysis of why respondents made their decisions about different nudge contents.

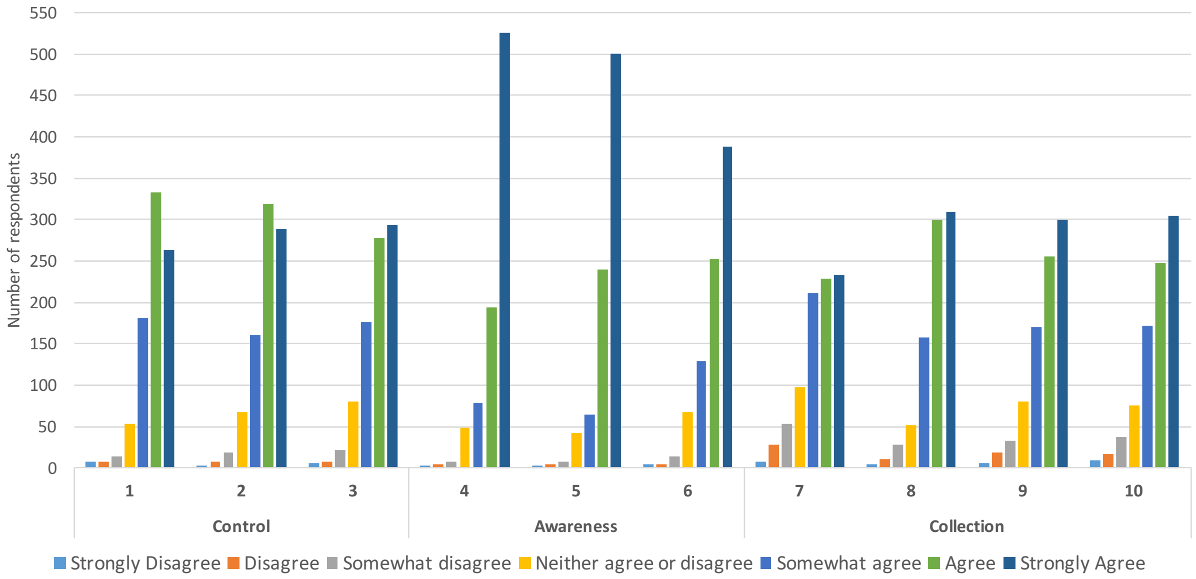

Figure 4.3: Distribution of respondents across factors of the Internet Users' Information Privacy Concerns (IUIPC).

4.2.1 Demographics

A total of 866 submissions passed the quality check questions. Of those, we removed 5 submissions. Two respondents indicted that they are Windows phone users. Although the other three submissions passed the quality check questions, their responses to the qualitative questions are incomprehensible and thus we decided to exclude them from the analysis. Accordingly, we focus on 861 submissions throughout this analysis.

Of the 861 respondents, 413 (48%) are females. Respondents' ages are between 18–70 (mean=34.1, median=32, SD=9.73). Of all respondents, 405 (47%) have bachelors degrees or higher. Statistical analysis shows no significant differences across conditions in terms of age (F = 1, p-value = 0.44), gender ($\tilde{\chi}^2$ = 5.26, df = 8, p-value = 0.73), or education level ($\tilde{\chi}^2$ = 7.53, df = 8, p-value = 0.48). Furthermore, IUIPC score analysis shows that there is a tendency toward privacy conscious attitude and behavior as Figure 4.3 depicts. The average scores for IUIPC factors are as follows: control is 17.54 (median=18, SD=2.79, max=21, min=6), awareness is 18.74 (median=20, SD=2.71, max=21, min=3), and collection is 22.77 (median=24, SD=4.66, max=28, min=4). Prior work has shown that workers on AMT are more concerned about privacy than the US general population [128]. Nonetheless, IUIPC scores are not statistically different across conditions (control: F = 0.65, p-value = 0.74; awareness: F = 0.81, p-value = 0.59; collection: F = 0.63, p-value = 0.75).

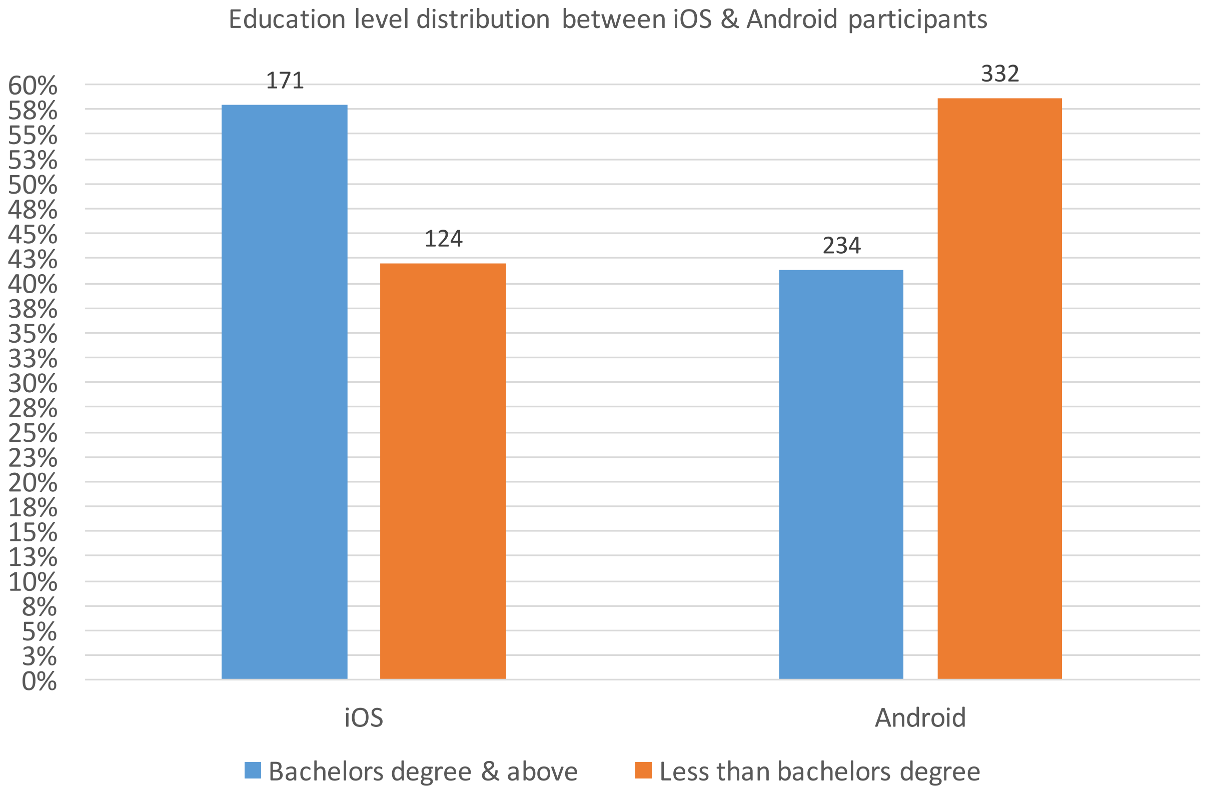

Figure 4.4: Distribution of eduction levels between respondents who use Android and iOS. The number of iOS respondents who have bachelors degree or above is significantly more than Android respondents.

For our experiment, we recruited respondents who use either one of the two most popular mobile operating systems, Android and iOS. Of all respondents, 566 (65.74%) use Android and 295 use iOS (34.26%). Of Android respondents, 262 (46.29%) are female respondents, whereas 151 (51.19%) of iOS respondents are female respondents. Although Android's sample is skewed toward male and iOS's sample is skewed toward female, the difference is not statistically significant ($\tilde{\chi}^2 = 1.67$, df = 1, p-value = 0.2). Of Android respondents, age range is 18–70 (mean=34.52, median=33, SD=9.55), whereas age range is 19–61 for iOS respondents (mean=33.31, median=31, SD=10). Respondents who use iOS are younger than Android's respondents and the age difference is statistically significant (t = 2.03, p-value = 0.04). For education levels, Android and iOS respondents differ significantly. As shown in Figure 4.4, only 234 (41.34%) of Android respondents have bachelors degree or higher, whereas 171 (58%) of iOS respondents have bachelors degree or higher ($\tilde{\chi}^2 = 20.85$, df = 1, p-value < 0.001). Furthermore, we analyzed the IUIPC scores for Android and iOS respondents as shown in Table 4.2. However, our analysis shows no statistical difference between between Android and iOS respondents with regard to IUIPC (Control: t = 0.73, p-value = 0.47; awareness: t =0.67, p-value = 0.51; collection: t =0.68, p-value =

Operating System		Control	Awareness	Collection
Android	Mean	17.59	18.78	22.85
	Median	18	20	24
	SD	2.76	2.77	4.78
	Max	21	21	28
	Min	6	3	4
iOS	Mean	17.44	18.65	22.63
	Median	18	19	23
	SD	2.87	2.59	4.43
	Max	21	21	28
	Min	6	0	4

Table 4.2: IUIPC scores for Android and iOS respondents. There is no statistical difference between Android and iOS respondents with regard to IUIPC scores.

0.5).

4.2.2 Responding to Different Nudge Contents

We posit that a nudge content is more effective than other nudge contents if it is more likely to motivate users to review and possibly adjust their app privacy settings. In this section, we show how respondents were motivated to review their app privacy settings by different nudge contents. Additionally, we show whether the "recent access" nudge content (that simply tells users that their location information has been recently used), which is used in current mobile operating systems, is more or less effective in motivating respondents to review their app privacy settings (and possibly adjust them) than other nudge contents.

When presented with a nudge, the respondent had four options to choose from: press the "Restrict which apps get my location" button ("Adjust settings"), press the "Keep sharing my location with the apps" button ("Keep settings"), close the screen ("Close screen"), or choose I do not know ("Don't know"). Overall, all nudge contents motivated more respondents to review their apps privacy settings than the baseline. The overall difference between nudge contents is statistically significant ("Adjust settings" and "Keep settings", two options[5]: $\tilde{\chi}^2 = 26.24$, df = 8, p-value < 0.001;

[5] We focus on the "Adjust settings" and "Keep settings" options to ensure that the statistical analysis is reliable by excluding cells with small counts. Chi square test becomes less reliable when cell counts are small [171]. This holds throughout the results section. Nonetheless, for completion we also report statistical results when all four options are included.

Condition	N	Adjust Settings	Keep Settings	Close Screen	Don't Know	$\tilde{\chi}^2$	p
(1) Baseline	96	56 (58.33%)	33 (34.38%)	7 (7.29%)	0 (0%)	-	-
(2) Frequency	97	67 (69.07%)	22 (22.68%)	8 (8.25%)	0 (0%)	2.63	0.1
(3) Background	99	75 (75.76%)	20 (20.2%)	4 (4.04%)	0 (0%)	5	0.025
(4) Purposes	97	75 (77.3%)	15 (15.5%)	3 (3.1%)	4 (4.12%)	8.5	0.004
(5) Purposes+Example	94	78 (83%)	15 (16%)	0 (0%)	1 (1%)	9.23	0.002
(6) Inferences	94	58 (61.7%)	34 (36.17%)	2 (2.13%)	0 (0%)	0	1
(7) Inferences+Example	92	65 (70.65%)	21 (22.83%)	6 (6.52%)	0 (0%)	2.72	0.1
(8) Predictions	96	69 (71.88%)	18 (18.75%)	7 (7.29%)	2 (2.08%)	4.97	0.026
(9) Predictions+Implications	96	74 (77.08%)	15 (15.63%)	5 (5.21%)	2 (2.08%)	8.24	0.004

Table 4.3: This table shows how respondents interacted with each nudge content. For each nudge content, the table shows the number of respondents, the distribution of their responses to the nudge based on the available options (including the actual numbers and percentages), and how significantly different those responses are from the baseline. The overall difference is statistically significant across conditions. Particularly, conditions 3, 4, 5, 8, and 9 (highlighted in green) differ significantly from condition 1, the baseline.

four options: $\tilde{\chi}^2 = 54.87$, df = 24, p-value < 0.001). Five of these nudge contents motivated a statistically significant number of respondents more than the baseline: conditions 3 (Background), 4 (Purposes), 5 (Purposes+Example), 8 (Predictions), and 9 (Predictions+Implications). Seventy-eight respondents (83%) in condition 5 (Purposes+Example) chose to adjust their app privacy settings. This is the highest percentage of respondents who became motivated by a nudge content. Conditions 4 (Purposes), and 9 (Predictions+Implications) come next with the same percentage of respondents who chose to restrict their app privacy settings: 75 (77.3%) in condition 4 (Purposes), and 74 (77.08%) in condition 9 (Predictions+Implications). The fourth and fifth nudge contents that are statistically significant from the baseline are conditions 3 (Background) and 8 (Predictions) with 75 (75.76%) respondents and 69 (71.88%) respondents, respectively, who wanted to adjust their app privacy settings in response to the nudge. Pairwise comparison with false discovery rate (FDR) correction [75] shows no statistical differences between these five nudge contents. Table 4.3 shows in detail how respondents interacted with different nudge contents. In sum, our results suggest that the "recent access" nudge content (which is used in current mobile operating systems) is less effective (in comparison to other nudge contents) in motivating respondents to review and possibly adjust their app privacy settings. Additionally, nudge contents that provide more information than the baseline (e.g., conditions 3, 4, 5, 8, and 9) are more effective in helping users manage their app privacy settings.

Conditions 2 (Frequency), 6 (Inferences), and 7 (Inferences+Example) were not more effective than the baseline. The nudge content in condition 2 (Frequency) did not significantly motivate more respondents to review their app settings than the baseline nudge. Notably, the frequency nudge in Chapter 3 (which is similar to the nudge content in condition 2) resonated with users and motivated them to review their app privacy settings. Perhaps, the frequency nudge is more effective in a real-world setting (as seen in Chapter 3) than in an experimental setting (the present chapter). The nudge content in condition 6 (Inferences) also did not significantly motivate more respondents than the baseline likely because it did not provide specific implications. The results of the nudge content in condition 7 (Inferences+Example) were unexpected. Although the nudge highlights a sensitive type of personal information (i.e. one's home address) [187], it was not more effective than the baseline nudge. It is possible that respondents did not perceive home address as sensitive or they did not believe that apps are able to make such inferences.

To explore whether respondents who use different mobile operating systems responded differently to nudge contents, we analyzed how Android and iOS respondents interacted with each nudge content. For Android respondents, the overall difference between nudge contents is statistically significant (two options: $\tilde{\chi}^2 = 23.52$, df = 8, p-value = 0.003; four options: $\tilde{\chi}^2 = 44.43$, df = 24, p-value = 0.007). In addition, five conditions motivated a statistically significant number of respondents more than the baseline (see Table 4.5): conditions 3 (Background), 4 (Purposes), 5 (Purposes+Example), 8 (Predictions), and 9 (Predictions+Implications). Android results are very similar to what we reported earlier in Table 4.3 for all respondents combined and regardless of their mobile operating system.

For iOS respondents, the interaction with nudge contents is drastically different across all conditions. The overall difference between nudge contents is not statistically significant (two options: $\tilde{\chi}^2$ = 10.02, df = 8, p-value = 0.26; four options: $\tilde{\chi}^2$ = 26.63, df = 24, p-value = 0.32). That is, the nudge contents in the treatment conditions are not more effective than the baseline, which simply informs users that their location information has been recently accessed by apps. In addition, overall and across all conditions, iOS respondents were more motivated to adjust their app privacy settings in response to nudge contents than their counterpart in Android. It is notable that the baseline motivated 71.4% of iOS respondents to adjust their app privacy settings, which is 18 percentage points more than the corresponding condition in Android. Similarly, condition 2 (Frequency) in iOS motivated 78.9% of respondents to adjust their app privacy settings, which is 16 percentage points more than the corresponding condition in Android. Two exceptions are conditions 5 (Purposes+Example) and 8 (Predictions) in which the percentage of respondents who chose to adjust settings is more for Android respondents than iOS (the differences are 3 percentage points for condition 5 and 1 percentage point for condition 8). Nonetheless, the difference between iOS and Android respondents from corresponding conditions is not statistically significant (The test only included two options: adjust and keep settings. The results of the tests are in Appendix A in Table A.1). This drastic difference between iOS and Android respondents could potentially be attributed (among other reasons) to two reasons. First, it is possible that iOS respondents are more accustomed to exercising control over apps than Android respondents. Privacy controls (which enable users to selectively allow or deny apps from accessing users' sensitive information) has been part of iOS for a long time (i.e. when iOS 4 and iOS 5 were released back in 2010 and 2012, respectively [122, 190]). However, similar functionality was only introduced recently in Android Marshmallow (i.e. Android OS version 6). Therefore, such familiarity with privacy controls by iOS respondents might have made them more inclined to exercise such control and choose to adjust their settings in response to nudge contents. It is probably relevant to point out that IUIPC privacy scale, however, is not statistically significant between iOS and Android respondents as shown in Section 4.2.1. Second, iOS respondents in our experiment are more educated and younger than Android respondents as reported in Section 4.2.1. Therefore, it is possible that younger and highly educated respondents were more able to make informed (complex) privacy decisions. However, our analysis shows no significant correlation between options chosen by participants and age when blocking by the operation system (The results of the statistical test are in Appendix A in Table A.2). In other words, changing the operating system does

Category of nudge contents	N	Adjust Settings	Keep Settings	Close Screen	Don't Know
(1) Data collection practices	292	198 (67.8%)	75 (25.7%)	19 (6.5%)	0 (0%)
(2) Purposes of data collection	191	153 (80.1%)	30 (15.7%)	3 (1.6%)	5 (2.6%)
(3) Potential implications of reusing collected data	378	266 (70.4%)	88 (23.3%)	20 (5.3%)	4 (1%)

Table 4.4: We grouped the nine nudges into one of three categories: (a) Data collection practices, (b) Purposes of data collection, and (c) Potential implications of reusing collected data. This tables shows how respondents interacted with combined nudge contents in each category. Nudge contents based on purposes of data collection (category (b)) motivated significantly more respondents to review and adjust their app privacy settings.

not change how age correlates with options chosen by respondents. Similarly, statistical analysis shows no significant difference between options chosen by respondents and education level when blocking by mobile operating system (two options: Cochran-Mantel-Haenszel M^2=2.25, df =1, p-value = 0.13; four options: Cochran-Mantel-Haenszel M^2=5.13, df =3, p-value = 0.16). These results suggest that it is still an open and probably an interesting research question to explore why Android and iOS users may differ in privacy behavior and attitudes.

As described in Section 4.1.2, the nine nudge contents can be assembled into one of three categories: (a) Data Collection Practices, (b) Purposes of data collection, and (c) Potential implications of reusing collected data. Thus, we were interested in how these three categories of nudge contents motivated respondents to review and adjust their app privacy settings. We combined conditions 1 (Baseline), 2 (Frequency), and 3 (Background) into category (a), and conditions 4 (Purposes) and 5 (Purposes+Example) into category (b), and conditions 6 (Inferences), 7 (Inferences+Example), 8 (Predictions), and 9 (Predictions+Implications) into category (c) as Table 4.4 shows. The overall difference between the three categories is statistically significant (two options: $\tilde{\chi}^2$ = 7.78, df = 2, p-value = 0.02; four options: $\tilde{\chi}^2$ = 21.52, df = 6, p-value = 0.001). Pairwise comparison with FDR correction shows that the difference is significant between category (b) and category (a) (data collection practices vs. purposes of data collection: two options, adjusted p-value = 0.025; four options, adjusted p-value $<$ 0.001), whereas the difference is marginally significant between category (b) and (c) (Purposes of data collection vs. potential implications of reusing collected data: two options, adjusted p-value = 0.049; four options, adjusted p-value = 0.01). That is, the results suggest that nudge contents that highlight why apps are using and collecting personal information can more effectively motivate users to review and possibly adjust their app settings than other nudge contents.

Condition	N	Adjust Settings	Keep Settings	Close Screen	Don’t Know	$\tilde{\chi}^2$	p
(1) Baseline	68	36 (52.94%)	26 (38.24%)	6 (8.82%)	0 (0%)	-	-
(2) Frequency	59	37 (62.7%)	17 (28.8%)	5 (8.5%)	0 (0%)	0.94	0.33
(3) Background	68	50 (74%)	14 (21%)	4 (6%)	0 (0%)	4.96	0.026
(4) Purposes	70	52 (74.3%)	14 (20%)	2 (2.85%)	2 (2.85%)	5.46	0.02
(5) Purposes+Example	57	48 (84.2%)	9 (15.8%)	0 (0%)	0 (0%)	8.56	0.003
(6) Inferences	52	29 (55.8%)	22 (42.3%)	1 (1.9%)	0 (0%)	0	1
(7) Inferences+Example	67	47 (70.1%)	15 (22.4%)	5 (7.5%)	0 (0%)	3.64	0.056
(8) Predictions	65	47 (72.31%)	11 (16.92%)	6 (9.23%)	1 (1.54%)	6.38	0.01
(9) Predictions+Implications	60	43 (71.7%)	11 (18.3%)	4 (6.7%)	2 (3.3%)	5.23	0.02

Table 4.5: The table shows reported results from Android respondents. For each condition, the table shows the number of respondents, the distribution of their responses to the nudge based on the available options (including the actual numbers and percentages), and how significantly different those responses are from the baseline. The responses in conditions 3, 4, 5, 8, and 9 (highlighted in green) differ significantly from condition 1, the baseline.

Condition	N	Adjust Settings	Keep Settings	Close Screen	Don’t Know
(1) Baseline	28	20 (71.4%)	7 (25%)	1 (3.6%)	0 (0%)
(2) Frequency	38	30 (78.9%)	5 (13.2%)	3 (7.9%)	0 (0%)
(3) Background	31	25 (80.65%)	6 (19.35%)	0 (0%)	0 (0%)
(4) Purposes	27	23 (85.2%)	1 (3.7%)	1 (3.7%)	2 (7.4%)
(5) Purposes+Example	37	30 (81.1%)	6 (16.2%)	0 (0%)	1 (2.7%)
(6) Inferences	42	29 (69%)	12 (28.6%)	1 (2.4%)	0 (0%)
(7) Inferences+Example	25	18 (72%)	6 (24%)	1 (4%)	0 (0%)
(8) Predictions	31	22 (71%)	7 (22.6%)	1 (3.2%)	1 (3.2%)
(9) Predictions+Implications	36	31 (86.1%)	4 (11.1%)	1 (2.8%)	0 (0%)

Table 4.6: The table shows reported results from iOS respondents. For each condition, the table shows the number of respondents, the distribution of their responses to the nudge based on the available options (including the actual numbers and percentages), and how significantly different those responses are from the baseline. The overall difference across conditions is not statistically significant.

Adjust Settings	Keep Settings	Close Screen	Don’t Know
(1) Minimize personal information collection (2) Purpose of why apps access sensitive information (3) Generic privacy concerns (4) Phone’s resources consumption (5) Potential consequences of collecting sensitive information (6) Usage basis (7) Benefits of nudges (8) Issues with the nudges	(1) Purpose of why apps access sensitive information (2) Low privacy concerns (3) Issues with the nudges (4) Undecided	(1) Low privacy concerns (2) Issues with the nudges (3) Undecided (4) Purpose of why apps access sensitive information	(1) Issues with the nudges (2) Undecided

Table 4.7: This table shows a summary of themes for each option available in the nudge.

4.2.3 Why Respondents Made their Decisions

For each nudge, respondents were asked to choose one of four options: “Adjust settings”, “Keep settings”, “Close screen”, “Don’t know”. In addition, respondents were asked to explain why they chose such an option. To understand why respondents made certain decisions when interacting with the nudges, we qualitatively analyzed their free responses to this open-ended question.

We used open coding [178] to analyze respondents’ answers to the open-ended question. Two researchers went iteratively through the responses and labeled sentences and phrases with concepts. Concepts were iteratively improved (e.g., new concepts were added or similar concepts were combined) and a code book was created. The two researchers independently went through all the responses and coded them using the code book. Later, they reviewed the coding together and disagreements were reviewed and resolved (by adopting coding by both researchers, agreeing on one coding after discussion, or keeping the disagreement). Next, related concepts within the code book were combined into more generic themes (i.e. categories). After applying the generic themes to the coding, the two researchers sampled the coding and reviewed the generic themes to ensure consistency. It is important to point out that a respondent’s answer can be labeled by more than one theme.

In this section, we report the results of the qualitative analysis by organizing them into the four options chosen by respondents when interacting with the nudges. Table 4.7 shows the summary of themes per each option.

Adjusting App Privacy Settings

In response to the nudges, 617 (71.7%) respondents chose the “Restrict which apps get my location” option. Eight themes emerged from analyzing respondents’ answers to why they chose to adjust their app privacy settings in response to the nudges.

(1) Minimize Personal Information Collection

Respondents chose this option because they wanted to *minimize the collection of their personal information*. For example, a respondent explained in generic terms: “Too much information is shared today over the [I]nternet because people don’t take the time to read and follow instructions. I would like to limit the amount of my personal information that is being accessed.” Other respondents wanted to pick and choose and to exercise control over apps and information they access. For instance, a respondent explained: “I want to control which apps get information about my location. If I am at a mall I may want information about a store having a sale. However, I don't want to be bothered by it otherwise. This would allow me to control my information." Some respondents wanted to restrict the flow of personal information to times only when they were aware "I do not like sharing information unknowingly." Other respondents took it a step further and indicated that they did not want to share their personal information at all: "I can find the information I need without having the phone automatically collected [it]. I don't like sharing my personal information or my contacts information, including location." Others only wanted to restrict excessive collection of personal information: "I dont like that they can tell where I am a lot of the time (...)." Respondents also wanted to protect specific types of information (which often were highlighted in the nudges). For example, a respondent explained the reason for restricting privacy settings: "There are very few apps that need to know my address and need to predict my income (...)." Respondents minimized the collection of personal information by manually providing less granular location information. For example, a respondent described: "I don't like the idea of apps collecting my location information. I know it's more convenient for some apps (like weather) to know my location without my telling them, but I'd rather take a few extra seconds to enter the data than have them track my location all the time." Some respondents wanted to minimize the amount of information that they share with specific categories of apps. For example, a respondent described: "I would want to limit the apps which have my location. I would want to allow Groupon and the Weather Channel apps, but restrict the game app." Others decided to adjust their privacy settings to avoid unwanted revelation of sensitive information through social components within apps: "I would press the 'restrict which apps get my location' button because some apps, like words with friends, could maybe reveal location information to people you don't know."

(2) Purpose of Why Apps Access Sensitive Information

Respondents chose to adjust their app privacy settings and cited *purposes* of apps' accessing their location information. For example, a respondent explained in general: "I don't want applications that do not need my location to access it." Others pointed out example apps that do not need the information: "I don't think that Words with Friends has any reason to know my location." Or, another respondent: "I don't agree with [G]roupon having access to my location as it should not need my location so I would restrict that app (...)." Others were okay only when apps accessed their location information to provide related functionality: "If they are not using it for normal functions of the app, I would not like them to track where I am." Other respondents were more specific and were okay only when apps accessed their information to provide main functionality: "Because the apps have used my location for something other than what the apps' main function is (...)."

(3) Generic Privacy Concerns

Respondents also cited *generic privacy concerns* for why they decided to adjust their app privacy settings. For example, a respondent explained: “because I’m paranoid.” And, another respondent elaborated: “[I] feel like this is an invasion of my privacy and honestly it[’]s a bit creepy.” Other respondents wanted to avoid potential security and privacy harms as a result of collecting personal information. For example, a respondent explained: “(...) If their servers are hacked, they will get my information and make me a target which I do not want. It is better to be safe and restrict the location for the apps.” Others indicated that they do not trust some apps: “I don’t trust certain apps to have my location (...).”

(4) Phone’s Resources Consumption

Some respondents chose to adjust their settings because they were worried about *phone resources consumption*. For example, a respondent explained: “Because (at least on my phone) GPS drains the battery, so I don’t want my more apps than necessary accessing my location.” Another wanted to avoid consuming mobile data plans: “(...) Not only is it an invasion of privacy, but it is also a drain on the battery. It also uses data that I have to pay for.”

(5) Potential Consequences of Collecting Sensitive Information

Respondents wanted to avoid *potential consequences* of collecting their personal information such as reusing collected information for secondary purposes. For example, a respondent explained in a detailed answer: “I don’t like the idea of apps accessing information that they do not need. It makes me think, ‘Why does this app need this information? Why is this app collecting this information? How is this information used? Is my information safe? Is this not needed info protected under their privacy policy? Who are they sharing this extra information with? How much money is this company making by sharing this information that they don’t need?’ (...).” Some respondents did not like targeted advertising in particular: “I’m not a big fan of location-based advertising. It feels creepy and it makes me uncomfortable.” Others wanted to avoid sharing their information with 3rd party recipients: “I would want to make sure only certain apps had access to my location. Ones that do not make sense seems like they are using your info to sell to others.” Respondents also did not like apps’ ability to infer additional (unknown) characteristics about them: “I would not feel comfortable with apps figuring out that information [i.e. the address where I live and my income] about me.” And, others wanted to avoid potential (negative) implications: “I don’t want my privacy invaded and as a result I’ll miss out on stuff such as paying higher prices.”

(6) Usage Basis

Respondents also chose to adjust their app settings to limit apps' access to location information on *usage basis*. For instance, a respondent explained: "I only want to share my location when using the app." Another respondent elaborated: "If I don't use the app I don't see why they need to access my location." Others allowed apps they regularly use: "If they are [a]pps that I do not regularly use, then I would restrict them being able to access my location." Other respondents preferred a policy in which they turn on location information when needed: "I try not to have location on unless I am using google maps or mapquest. I always turn it back off after (...)."

(7) Benefits of Nudges

Respondents also referred to *the nudges themselves* for being a reason to choose this option. For example, a respondent explained: "I would believe the [nudge] app, since I don't see why it would lie to me." Notably, the nudges also switched respondents' attention to privacy management which is typically a secondary task for users. For example, a respondent became interested in reviewing privacy settings of other apps: "I would want to see if there are apps that are using my location that I might not be aware of, or that I have forgotten that I granted access to." Another respondent wanted to review additional permissions (beyond location information): "I don't feel that those apps need to know my location. I would use this opportunity to check other permissions."

(8) Issues with the Nudges

Some respondents chose to adjust their app privacy settings provisionally although they had some *issues with the nudges*. For example, a respondent explained: "I would restrict access to my location as a preventative measure until I can learn more about the behavior of the three apps listed on the screen. First of all, what purposes are related to the apps' main function? Who determines that? Next, for what purposes did these apps access my location? How does this app determine for what purposes those three apps accessed my location?" Another respondent described similar reasoning and pointed out that the information that the nudge provided might have not been sufficient: "I would restrict apps that get my location because many apps do require the location in order to function properly but this would allow me to pay better attention to which apps have access to my location depending on when I need to use them. I would also probably look further into what it means by predicting my income as I feel this info is not necessary to know, and how this can potentially be harmful to me."

Keeping App Privacy Settings Unchanged

In response to the nudges, 193 (22.4%) respondents chose the “Keep sharing my location with the apps” option. Four themes emerged from analyzing respondents’ answers to why they chose to keep their app privacy settings unaltered in response to the nudges.

(1) Purpose of Why Apps Access Sensitive Information

Respondents chose this option because they wanted the apps to keep accessing location information for a *purpose*. For instance, a respondent explained in generic terms: “I would assume that the apps would need to know my location to function properly.” Other respondents wanted to share their location information with the apps to enjoy tailored experience that may otherwise be unavailable. A respondent described: “I don’t mind these apps having access to my location because they need it to show me customized information based on my area.” Respondents sought tailored experience from The Weather Channel and Groupon in particular: “If I did want to use the Weather Channel [a]pp, I would want it to know my location so I could get an accurate forecast. If I was out somewhere and decided to look for a deal on Groupon, I would also want it to know my location (...).” Some respondents were torn between functionality and privacy trade-offs. They acknowledged the privacy concern but cited the importance of functionality and thus decided to keep sharing their location information with the apps. For instance, one respondent explained: “I am concerned about what they could possibly do with the info as stated in the warning, however, an app like the weather channel needs to know my location in order to function (...).” And, another respondent described similar reasoning more explicitly: “I feel like if I want to know what my weather forecast is going to be, I need to sacrifice a bit of privacy (...).”

(2) Low Privacy Concerns

Other respondents chose to keep sharing their location information with the apps because they have *low privacy concerns*. A respondent explained: “Because I don’t really care that apps are accessing my location.” Respondents also decided to keep sharing their location information with apps because they trusted them and were familiar with them. For instance, a respondent explained: “I don’t really mind those apps knowing my location. They are famous apps and well-known (...).” Another respondent linked trust to usage pattern: “because [I] would trust theses apps if I used them alot (...).” Other respondents expected such practices from apps: “This is normal behavior for phone apps.” Others did not perceive any privacy risks based on such practices: “Because I’m sure they wouldn’t do anything harmful with my information!” Other respondents were not concerned about targeted advertising in particular: “I don’t see any reason why I need to hide my location from most advertisers (...).” Some respondents wanted to avoid the hassle of configuring app privacy settings now: “Some of these apps need to see my location, and I can’t be bothered to go through individually and set them.” Others wanted to avoid doing so in the future: “I am not concerned with apps knowing my location, and do not want to have to deal with allowing them to know in the future.” Others believed privacy is already lost: “Since the privacy doesn’t exist anymore it’s okay for me to share my location with those apps.”

(3) Issues with the Nudges

Respondents also chose this option because they had some *issues with the nudges* themselves. For example, some respondents did not realize that they could adjust settings for individual apps: “(...) I don’t know why Words With Friends would need to know my location, but I can’t restrict it and allow the other two, so I would leave it alone.” Or, they indicated they wanted to adjust app settings for one app but they did not follow through: “My location is helping me get better info about me. I would restrict the words with friends as that is unnecessary.” Others were afraid that if they acted on the nudge contents, the apps or the phone would stop working: “I am afraid to change things and break the phone.” Other respondents preferred to manage their app privacy settings using *alternative approaches* (other than the nudges). Some respondents explained the approach they use to manage app privacy either before installing apps: “I pay close attention to the permissions that apps get and the reasons for them. I already know which apps are utilizing things like GPS location and what reason they have for doing so, and I would only install an app that has such permissions if it had a legitimate reason to (...).” Others explained a similar approach for after installation: “(...) When I first purchased my current phone (Samsung S5) I went through the applications to see which one was grabbing personal information and decided whether to keep the application [or] not (...).”

(4) Undecided

Some respondents chose this option because they were *undecided*. For example, a respondent explained: “It really depends on what apps I have running, at the time.” Respondents also were undecided because they needed to investigate the issue more: “(...) I don’t know what Words With Friends does, I’d have to look into it. ”

Choosing to Close the Nudge Screen

Forty-two (4.9%) respondents chose the “Close the screen” option in response to the nudges. Our analysis shows four themes that explain why respondents chose this option. It is notable that these four themes are similar to the themes that emerged when respondents chose to keep their settings unchanged.

(1) Low Privacy Concerns

Respondents, who chose this option, indicated that they generally have *low privacy concerns*. For example, a respondent's answer reads: "I don't really care enough about it." Other respondents thought that privacy is already lost: "I think a lot of them have that information, anyway." Some respondents were not concerned because they did not perceive any privacy risks in this context. A respondent explained: "I am not really too concerned about these apps accessing my location. I am not sure a lot of harm can really be done with an app knowing my location." Other respondents were familiar with the apps and trusted them: "I Would close the screen because the applications showing that had used my location was known to me and I was okay with that (...)."

(2) Issues with the Nudges

Respondents chose to close the nudge screen because they had *issues with the nudges* themselves. For example, some respondents did not like the idea of receiving app privacy nudges: "I'd actually uninstall the app, honestly. This is annoying and patronizing. I know how to control location access already if I f[***]ing care." Others seem to misunderstand available options in the nudge. A respondent explained: "(...) If it let me choose which ones I could block and keep then I might have selected accordingly." Other respondents found the nudge contents to be insufficient: "I do not know what to do with this information. I'm not sure if it's harmful for the apps to have my location (...)." Others did not find the content of the nudge assuring enough: "I wouldn't want to select any options that might cause problems with my phone (...)." In addition, respondents chose this option because they preferred *alternative approaches* to managing their app privacy than utilizing the nudges. For example, an iOS respondent explained: "I have the ability to allow each app permissions in IOS so it would be silly for me to get an app that does what a simple two or three pushes on my phone will allow me to do (...)." A respondent who uses Android Marshmallow (i.e. Android OS version 6) explained a similar reasoning: "Because I already have fine-grained control over what can access my location in my version of android. This application wouldn't be particularly useful to me, since it's now essentially built into android itself (...)." Others pointed out that they already made informed decisions about individual apps, which is their alternative approach to manage app privacy settings: "I don't like nanny apps and since I read everything before signing up for each app, I likely already know how much any individual app is accessing whatever."

(3) Undecided

Some respondents chose this option because they were *undecided*. For example, a respondent explained: “I wouldn’t be sure what to do so I would just make the screen go away!” Others wanted to investigate more before making any decision: “I would have to look into it more before deciding (...).”

(4) Purpose of Why Apps Access Sensitive Information

Respondents also chose to close the nudge screen because they believed that the apps requested access to location information for a *purpose*. A respondent described a generic reason: “I figure if an app needs my location there must be a reason.” Another respondent explained more specifically: “I already figure those apps need my location to work correctly for the most part (...).”

Choosing the “I do not know ” Option

Nine (1%) respondents chose the “I do not know” option when interacting with the nudges. Our analysis shows two themes for why respondents chose this option.

(1) Issues with the Nudges

Some respondents chose this option because they had some *issues with the nudge* itself such as questioning the content of the nudge or misunderstanding the options that the nudge provides. For instance, a respondents explained why she chose this option: “Well, I am not sure I trust the OS always knows whether these apps were using data for normal purposes or not (...).” Another respondent pointed out that the information that the nudge provided might have been insufficient: “I do not know why I should care that apps can discover info about my income. I would like to know why this is something I should care about before I act on it (...).” Some respondents appeared to think that the nudge does not allow respondents to adjust settings for individual apps (i.e. either allows or deny all apps). A respondent described what seems like a misunderstanding of available options: “It would really depend on which apps are accessing my location (...) However, Words with Friends does not need to know my location for the app to run.” Additionally, some respondents preferred *alternative approaches* to managing their app privacy than utilizing the nudges. A respondent explained: “I would manually go in and change the location services allowances for each of those apps using the iOS privacy functionality in settings.”

(2) Undecided

Respondents also chose this option because they were *undecided* and needed more information. A respondent explained: "I would need to know more details about how each company utilizes my information (...)."

4.2.4 Factors that Influenced Respondents' Decisions

After completing the main task, respondents were asked three questions, one after another: how likely or unlikely is it that the nudge content actually occurs, whether respondents were aware or not aware of the nudge content before the study, and how concerned or unconcerned they are about the nudge content. The goal of these questions is to explore how believability of the nudge contents, awareness of them, and concern about them affect respondents' decisions.

Likelihood of Nudge Contents

Table 4.8 shows how respondents think about the likelihood of each nudge content. The difference between nudge contents is statistically significant overall ($\tilde{\chi}^2$ = 47.5, df = 16, p-value < 0.001). Conditions 2 (Frequency), 6 (Inferences), 7 (Inferences+Example), 8 (Predictions), and 9 (Predictions+Implications) in particular (highlighted in green in Table 4.8) are significantly different from the baseline. Additionally, pairwise comparisons with FDR correction (between the likelihood levels) show significant differences between unlikely and likely in conditions 2 (Frequency: adjusted p-value < 0.001)[6], 6 (Inferences: adjusted p-value = 0.04), 8 (Predictions: adjusted p-value = 0.03), and 9 (Predictions+Implications: adjusted p-value = 0.01) in comparison to condition 1, Baseline. Our analysis shows no significant difference between unlikely and likely in condition 7 (Inferences+Example) in comparison to condition 1 (Baseline). That is, respondents in conditions 2 (Frequency), 6 (Inferences), 8 (Predictions), and 9 (Predictions+Implications) think that what these nudges' contents convey is less likely to occur than the baseline, which simply informs users that their location has recently been accessed by apps. Furthermore, we explored whether these five conditions differ from one another. However, pairwise comparison with FDR correction does not show statistical differences between the five conditions (See Appendix A Table A.4).

When designing different nudge contents, we hypothesized that a nudge content may become less effective if respondents believe that the content is less likely to occur in real life. As such, we examined how the options chosen by the respondents (e.g., "Adjust settings" vs. "Keep settings")

[6] There is also a significant difference between neutral and likely for condition 2 (Frequency: adjusted p-value < 0.001).

Condition	N	Unlikely	Neutral	Likely	$\tilde{\chi}^2$	p
(1) Baseline	96	3 (3.1%)	3 (3.1%)	90 (93.8%)	-	-
(2) Frequency	97	24 (25%)	11 (11%)	62 (64%)	26.1	< 0.001
(3) Background	99	5 (5%)	8 (8%)	86 (87%)	2.81	0.24
(4) Purposes	97	8 (8.2%)	9 (9.3%)	80 (82.5%)	5.86	0.054
(5) Purposes+Example	94	4 (4.3%)	7 (7.4%)	83 (88.3%)	2	0.37
(6) Inferences	94	13 (13.8%)	6 (6.4%)	75 (79.8%)	8.6	0.01
(7) Inferences+Example	92	11 (12%)	10 (11%)	71 (77%)	10.5	0.005
(8) Predictions	96	13 (13.5%)	11 (11.5%)	72 (75%)	12.8	0.002
(9) Predictions+Implications	96	15 (15.6%)	11 (11.5%)	70 (72.9%)	15.1	< 0.001

Table 4.8: This table shows how respondents think about the likelihood of each condition. Likely and extremely likely are combined into one column: likely. Similarly, unlikely and extremely unlikely are combined into one column: unlikely. The overall difference of likelihood between conditions is statistically significant. The likelihoods of conditions that are statistically significant from the baselines are highlighted in green.

Likelihood Level	Adjust Settings	Keep Settings	Close Screen	Don't Know
Unlikely	66 (68.8%)	22 (22.9%)	7 (7.3%)	1 (1%)
Neutral	43 (56.6%)	19 (25%)	12 (15.8%)	2 (2.6%)
Likely	508 (73.7%)	152 (22.1%)	23 (3.3%)	6 (0.9%)

Table 4.9: This table shows the relation between how respondents think about the likelihood of each nudge content and what options they chose when interacting with the nudges. Likely and extremely likely are combined into one column: likely. Similarly, unlikely and extremely unlikely are combined into one column: unlikely. How respondents think about the likelihood of each nudge content does not have a significant effect on how respondents interacted with the nudges. The same conclusion also holds when blocking by conditions.

were affected by how respondent perceived the likelihood of each nudge content. Table 4.9 shows a summary of the results. Our statistical analysis shows that the likelihood of nudge contents did not affect respondents' decisions when all conditions are combined (two options: $\tilde{\chi}^2$= 1.9, df = 2, p-value = 0.39; four options: $\tilde{\chi}^2$= 28.29, df = 6, p-value < 0.001[7]). The same results also hold even when blocking by condition (two options: Cochran-Mantel-Haenszel M^2=3, df =2, p-value = 0.22; four options: Cochran-Mantel-Haenszel M^2= 26.54, df = 6, p-value < 0.001). This analysis suggests that although respondents perceived the likelihood of nudge contents differently (as shown in Table 4.8), the chosen nudge contents were reasonable enough such that they did not affect how respondents interacted with the nudges.

Awareness of Nudge Contents

One of the objectives of designing app privacy nudges is to increase users' awareness of privacy risks associated with apps installed on their phones. Table 4.10 shows how aware respondents are about each nudge's content. The difference between the awareness of nudge contents is statistically significant overall ($\tilde{\chi}^2$ = 191.76, df = 32, p-value < 0.001). All treatment conditions are significantly different from the control condition. Additionally, pairwise comparisons (between awareness levels) with FDR correction show significant differences between "Not at all aware" and "Extremely aware" in conditions 2 (Frequency), 3 (Background), 4 (Purposes), 5 (Purposes+Example), 8 (Predictions), and 9 (Predictions+Implications) in comparison to condition 1 (Baseline)[8] as Table 4.12 shows. Furthermore, significant differences are also found between "Moderately aware" and "Extremely aware" in conditions 6 (Inferences) and 7 (In-

[7] Although p is significant, the $\tilde{\chi}^2$ approximation may not be accurate due to small counts in "Close Screen" and "I don't know" cells

[8] There are also significant differences between other awareness levels in conditions 2 (Frequency), 3 (Background), 4 (Purposes), 5 (Purposes+Example), 8 (Predictions), and 9 (Predictions+Implications) in comparison to condition 1 (Baseline).

Condition	N	Not at all aware	Slightly aware	Somewhat aware	Moderately aware	Extremely aware	$\tilde{\chi}^2$	p
(1) Baseline	96	0 (0%)	1 (1%)	5 (5.2%)	19 (19.8%)	71 (74%)	-	-
(2) Frequency	97	15 (15.5%)	27 (27.8%)	19 (19.6%)	19 (19.6%)	17 (17.5%)	80.44	< 0.001
(3) Background	99	8 (8.1%)	11 (11.1%)	17 (17.17%)	28 (28.28%)	35 (35.35%)	36.8	< 0.001
(4) Purposes	97	8 (8.2%)	15 (15.5%)	18 (18.6%)	32 (33%)	24 (24.7%)	54.16	< 0.001
(5) Purposes+Example	94	9 (9.6%)	11 (11.7%)	15 (16%)	26 (27.7%)	33 (35%)	37.29	< 0.001
(6) Inferences	94	3 (3%)	6 (6%)	10 (11%)	43 (46%)	32 (34%)	32.28	< 0.001
(7) Inferences+Example	92	3 (3.3%)	6 (6.5%)	9 (9.8%)	36 (39.1%)	38 (41.3%)	22.89	< 0.001
(8) Predictions	96	8 (8.3%)	8 (8.3%)	19 (20%)	34 (35.4%)	27 (28%)	45.6	< 0.001
(9) Predictions+Implications	96	20 (21%)	13 (13.5%)	29 (30%)	22 (23%)	12 (12.5%)	89.39	< 0.001

Table 4.10: This table shows how aware respondents are about each nudge's content. The overall difference of awareness level between conditions is statistically significant. The awareness level of conditions that are statistically significant from the baselines is highlighted in green.

ferences+Example) in comparison to condition 1 (see Appendix A Table A.3). Therefore, this analysis suggests that respondents in the treatment conditions (conditions 2 to 9) are less aware of the practices conveyed in the nudge contents than their corresponding respondents in the baseline, who were simply informed that their location has been recently accessed by some apps. In other words, the nudges in the treatment conditions increased respondents' awareness of privacy risks and unexpected data collection practices by apps. Furthermore, pairwise comparison with FDR correction shows significant difference between condition 2 (Frequency) and conditions 3 (Background), 5 (Purposes+Example), 6 (Inferences), 7 (Inferences+Example), and 8 (Predictions); and between condition 9 (Predictions+Implications) and conditions 3 (Background), 4 (Purposes), 5 (Purposes+Example), 6 (Inferences), 7 (Inferences+Example), and 8 (Predictions). The statistical results are shown in Table A.5 in Appendix A. This suggests that respondents in conditions 2 (Frequency) and 9 (Predictions+Implications) are generally less aware of practices conveyed in nudges than their corresponding respondents in other conditions. Specifically, respondents are less aware of access frequency by apps and being potentially subject to price discrimination based on location traces.

When designing different nudge contents, we also hypothesized that if respondents are unaware of practices conveyed in nudge contents, they are more likely to become motivated by such nudges. To this end, we explored the correlation between options chosen by respondents and their reported awareness of nudge contents. Table 4.11 shows a summary of the results. Statistical analysis shows significant correlation between awareness and respondents decisions when interacting with the nudges (two options: $\tilde{\chi}^2$= 17.29, df = 4, p-value = 0.002; four options: $\tilde{\chi}^2$= 28.27, df = 12, p-value = 0.005). Pairwise comparison with FDR correction shows significant differences between options chosen by respondents when they are "Not at all aware" of nudge contents and when they are "Extremely aware" of nudge contents (two options: adjusted p-value < 0.001; four options: adjusted p-value < 0.001). When respondents are "Not aware at all" of a nudge content, they are more likely to choose to adjust their app privacy settings than when respondents are "Extremely aware" of nudge contents. In addition, there is a significant differences between options chosen by respondents when they are "Extremely aware" of nudge contents than when respondents are "Moderately aware" of nudge contents (two options: adjusted p-value = 0.002; four options: adjusted p-value < 0.001)[9]. That is, when respondents are "Moderately aware" of nudge contents they are more likely to choose to adjust their app privacy settings than when respondents are "Extremely aware" of nudge contents. Additionally, we analyzed the correlation between respondents decisions and awareness when blocking by condition. Our analysis shows a significant difference between conditions (two options: Cochran-Mantel-Haenszel

[9] There are also significant differences between other awareness levels.

Awareness Level	Adjust Settings	Keep Settings	Close Screen	Don't Know
Not at all aware	61 (82.4%)	9 (12.2%)	2 (2.7%)	2 (2.7%)
Slightly aware	76 (77.5%)	18 (18.4%)	4 (4.1%)	0 (0%)
Somewhat aware	97 (68.8%)	31 (22%)	9 (6.4%)	4 (2.8%)
Moderately aware	198 (76.45%)	49 (18.9%)	11 (4.25%)	1 (0.4%)
Extremely aware	185 (64%)	86 (29.8%)	16 (5.5%)	2 (0.7%)

Table 4.11: This table shows the relation between how aware respondents are about each nudge's content and what options they chose when interacting with the nudges. How aware respondents are about each nudge's content has a significant effect on how respondents interacted with the nudges. The same conclusion also holds when blocking by conditions.

Not at all aware vs. Extremely aware	Adjusted p-value
Baseline : Frequency	< 0.001
Baseline : Background	0.002
Baseline : Purposes	< 0.001
Baseline : Purposes+Example	0.001
Baseline : Inferences	0.15
Baseline : Inferences+Example	0.22
Baseline : Predictions	< 0.001
Baseline : Predictions+Implications	< 0.001

Table 4.12: This table shows the pairwise comparison with FDR correction between two awareness levels (Not at all aware vs. Extremely aware) in the baseline in comparison to treatment conditions.

M^2=11.57, df = 4, p-value = 0.02; four options: Cochran-Mantel-Haenszel M^2= 19.34, df = 12, p-value = 0.08). However, pairwise comparisons did not show significant differences between conditions.

Concern about Nudge Contents

It is reasonable to assume that respondents are more likely to become motivated to adjust their app privacy settings if they are concerned about what nudge contents convey. Table 4.13 shows how concerned respondents are about each nudge's content. The difference between nudge contents is not statistically significant overall ($\tilde{\chi}^2$ = 40.57, df = 32, p-value = 0.14). This might be explained by the relatively high IUIPC scores reported in Section 4.2.1. In other words, respondents were concerned across the board.

Condition	N	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
(1) Baseline	96	13 (13.5%)	20 (20.8%)	27 (28.1%)	22 (23%)	14 (14.6%)
(2) Frequency	97	8 (8.2%)	18 (18.6%)	21 (21.6%)	31 (32%)	19 (19.6%)
(3) Background	99	11 (11.1%)	18 (18.2%)	18 (18.2%)	33 (33.3%)	19 (19.2%)
(4) Purposes	97	6 (6.2%)	11 (11.3%)	20 (20.6%)	38 (39.2%)	22 (22.7%)
(5) Purposes+Example	94	3 (3.2%)	9 (9.6%)	18 (19.1%)	39 (41.5%)	25 (26.6%)
(6) Inferences	94	9 (10%)	22 (23%)	15 (16%)	29 (31%)	19 (20%)
(7) Inferences+Example	92	9 (9.8%)	13 (14.1%)	19 (20.7%)	27 (29.3%)	24 (26.1%)
(8) Predictions	96	12 (12.5%)	14 (14.6%)	22 (22.9%)	34 (35.4%)	14 (14.6%)
(9) Predictions+Implications	96	6 (6.3%)	12 (12.5%)	26 (27.1%)	27 (28.1%)	25 (26%)

Table 4.13: This table shows how concerned respondents are about each nudge's content. The overall concern level between conditions does not differ significantly.

Concern Level	Adjust Settings	Keep Settings	Close Screen	Don't Know
Not at all concerned	10 (13%)	51 (66.2%)	16 (20.8%)	0 (0%)
Slightly concerned	56 (40.9%)	68 (49.6%)	12 (8.8%)	1 (0.7%)
Somewhat concerned	137 (73.7%)	38 (20.4%)	9 (4.8%)	2 (1.1%)
Moderately concerned	240 (85.71%)	30 (10.71%)	5 (1.79%)	5 (1.79%)
Extremely concerned	174 (96.1%)	6 (3.3%)	0 (0%)	1 (0.6%)

Table 4.14: This table shows the relation between how concerned respondents are about each nudge's content and what options they chose when interacting with the nudges. How concerned respondents are about each nudge's content have a significant effect on how respondents interacted with the nudges. The same conclusion also holds when blocking by conditions.

When designing different nudge contents, we hypothesized that nudge contents that concern respondents are more likely to motivate them to review and adjust their app privacy settings than nudge contents that are less concerning. To this end, we explored the correlation between respondents' concern about nudge contents and the options chosen by respondents when interacting with the nudges (e.g., "Adjust settings" vs. "Keep settings"). Table 4.14 shows a summary of the results. Our analysis shows significant differences between level of concern and options chosen by respondents (two options: $\tilde{\chi}^2$ = 261.55, df = 4, p-value < 0.001; four options: $\tilde{\chi}^2$ = 303.16, df = 12, p-value < 0.001). Pairwise comparison with FDR correction shows significant differences between all concern levels as Table 4.15 shows. That is, respondents are more likely to choose to adjust their app privacy settings as their level of concern about nudge contents increases. Additionally, our analysis shows significant differences between respondents' decisions and concern level when blocking by condition (two options: Cochran-Mantel-Haenszel M^2=235.1, df =4, p-value < 0.001; four options: Cochran-Mantel-Haenszel M^2= 275, df = 12, p-value < 0.001). However, pairwise comparisons with FDR correction shows no statistical differences between conditions.

Privacy Concern Level

We hypothesized that respondents who exhibit a higher privacy concern level are more likely to choose to adjust their app privacy settings in response to the nudges. To this end, we analyzed the correlation between IUIPC scores and the options chosen by respondents when interacting with the nudges. ANOVA analysis shows a significant difference among IUIPC scores and options (e.g., "Adjust settings" vs. "Keep settings") chosen by respondents (control: F = 8.59, p-value < 0.001; awareness: F = 12.69, p-value < 0.001; collection: F = 35.1, p-value < 0.001). Furthermore, pairwise comparison with Tukey correction shows statistically significant differences

Pairwise Comparison	Adjusted p-value (four options)	Adjusted p-value (two options)
Not at all concern : Slightly concern	< 0.001	< 0.001
Not at all concern : Somewhat concern	< 0.001	< 0.001
Not at all concern : Moderately concern	< 0.001	< 0.001
Not at all concern : Extremely concern	< 0.001	< 0.001
Slightly concern : Somewhat concern	< 0.001	< 0.001
Slightly concern : Moderately concern	< 0.001	< 0.001
Slightly concern : Extremely concern	< 0.001	< 0.001
Somewhat concern : Moderately concern	0.004	0.004
Somewhat concern : Extremely concern	< 0.001	< 0.001
Moderately concern : Extremely concern	0.004	0.005

Table 4.15: This table shows the pairwise comparison between concern levels and options chosen by participants. We used pairwise comparison with FDR correction with all four options that were available to respondents and with only two options: Adjust vs. Keep settings.

between IUIPC scores for respondents who chose "Adjust settings" vs. "Keep settings" (control: adjusted p-value < 0.001; awareness: adjusted p-value < 0.001; collection: p-value < 0.001) and for respondents who chose "Adjust settings" vs. "Close Screen" (awareness: adjusted p-value < 0.001; collection: p-value < 0.001; but the difference is not significant for control: adjusted p-value = 0.06). These results suggest that as respondents' privacy concern level increases (as indicated by IUIPC scores), the likelihood of choosing "Adjust settings" option increases as well. In other words, users, who are privacy conscious, are more likely to become motivated by nudges than users who are less concerned about privacy.

Familiarity with the Apps

Along with the nudge text, three apps were presented: The Weather Channel, Groupon, and Words With Friends. After completing the main task, respondents were asked to indicate how familiar or not familiar they are with each one of these three apps by choosing one of four options: "I have never heard of the app", "I have heard of the app but I have never installed it on my phone", "I had installed the app on my phone at some point in time but it is not on my phone anymore", or "I have installed the app on my phone and it is still on my phone." As Figure 4.5 shows, the majority of respondents recognized the three apps listed on the nudge. However, respondents who had used or were still using any of the apps are as follows: 380 (44.13%) for Groupon, 557 (64.7%) for The Weather Channel, and 455 (52.9%) for Words With Friends.

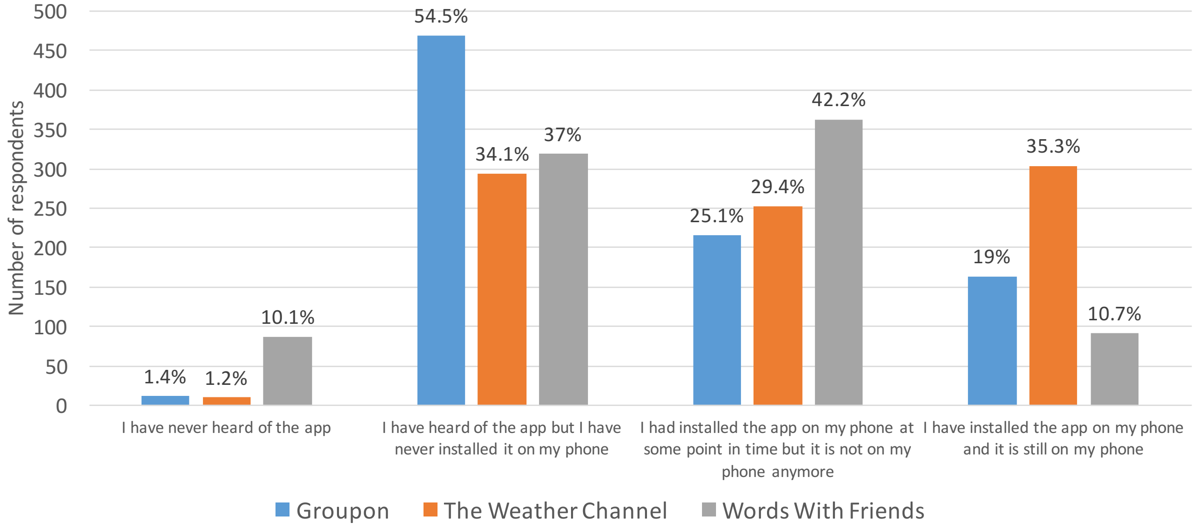

Figure 4.5: Respondents self-reported how familiar they are with the three apps listed in the nudges. The majority of respondents recognized these three apps.

We explored how familiarity with these apps affected respondents' decisions when interacting with the nudges. Our analysis shows no significant correlation between familiarity with apps and the options chosen by respondents when interacting with the nudges (Groupon: two options: $\tilde{\chi}^2$= 5.7, df = 3, p-value = 0.13; four options: $\tilde{\chi}^2$= 10.74, df = 9, p-value = 0.29; The Weather Channel: two options: $\tilde{\chi}^2$=0.68, df = 3, p-value =0.88; four options: $\tilde{\chi}^2$= 8.1, df = 9, p-value = 0.53; Words With Friends: two options:$\tilde{\chi}^2$= 0.95, df = 3, p-value = 0.81; four options: $\tilde{\chi}^2$= 3.75, df = 9, p-value = 0.93).

Of all respondents, 164 (19%) reported that they still have Groupon installed on their phones, 304 (35.3%) have The Weather Channel, and 92 (10.7%) have Words With Friends. Respondents who indicated that they still have any of the three apps installed on their phones were subsequently asked how frequently they use the app and were given the following options: "Never", "Rarely", "Sometimes", "Often", and "Always." As Figure 4.6 shows, the majority of these respondents use the apps sometimes or more: 118 (72%) for Groupon, 285 (93.75%) for The Weather Channel, and 67 (72.8%) for Words With Friends. Furthermore, we explored how app usage patterns affected respondents' decisions when interacting with the nudges. Our analysis shows significant correlation between app usage patterns and options chosen by respondents for Groupon and The Weather Channel apps but not for Words With Friends app (Groupon: two options: $\tilde{\chi}^2$ = 9.76, df = 4, p-value = 0.04; four options: $\tilde{\chi}^2$= 13.51, df = 9, p-value = 0.3; The Weather Channel: two options: $\tilde{\chi}^2$ = 10, df = 4, p-value = 0.04; four options: $\tilde{\chi}^2$= 17.22, df = 9, p-value = 0.14; Words

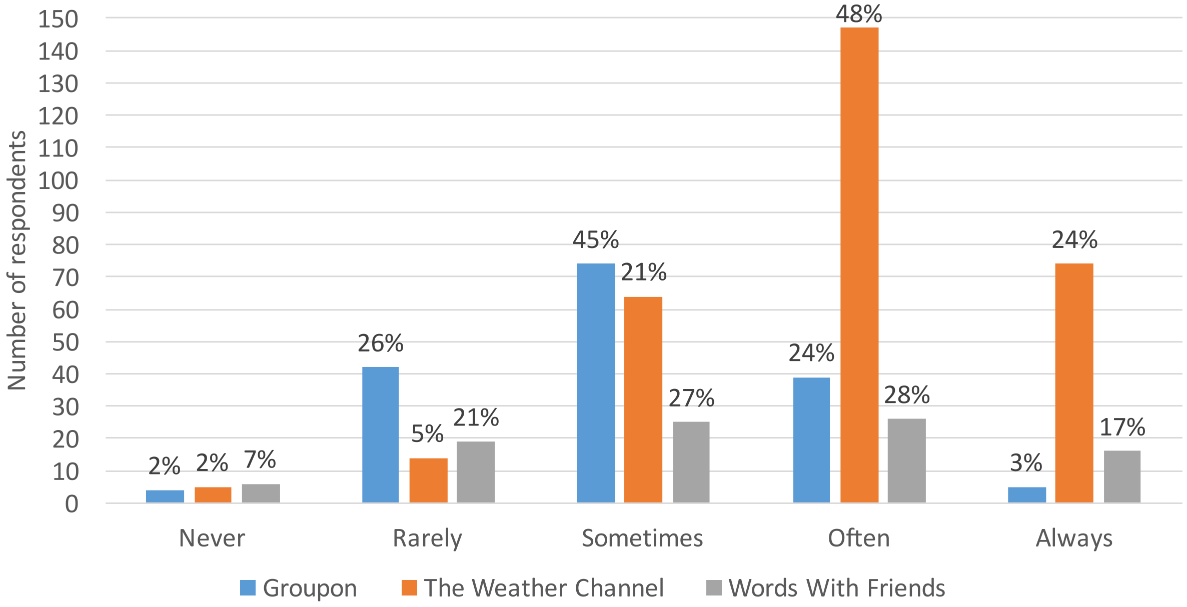

Figure 4.6: Respondents, who indicated that they still have any of these three apps installed on their phones, self-reported how frequently they use the apps. The majority of those respondents use these apps sometimes or more.

Gender	N	Adjust Settings	Keep Settings	Close Screen	Don't Know
Female	413	308 (74.6%)	78 (18.9%)	23 (5.5%)	4 (1%)
Male	448	309 (69%)	115 (25.7%)	19 (4.2%)	5 (1.1%)

Table 4.16: This table shows the distribution of options chosen by respondents and their gender. There was a significant correlation between gender of the respondent and the chosen option. Female respondents were more likely to choose "Adjust Settings."

With Friends: two options: $\tilde{\chi}^2 = 4.8$, df = 4, p-value = 0.3[10]). It seems that as usage of these two apps increases, the likelihood that the keep settings option was chosen by respondents increases as well. However, pairwise comparison with FDR correction between levels of app usage and options chosen by respondents did not show significant differences for both Groupon and The Weather Channel apps. The overall results of app familiarity and app usage suggest that these two factors did not have a strong effect on respondents' decisions. In other words, although our experiment only showed these three apps, their effect on respondents' decisions seem minimal.

[10] The statistical test is not applicable when using the four options for Words With Friends due to small counts in cells.

Demographics Factors

We also explored whether gender, age, or education level have an effect on options chosen by respondents when interacting with the nudges.

Our analysis shows that gender is correlated significantly with options chosen by respondents as shown in Table 4.16. Particularly, female respondents were more likely to choose "Adjust settings" than male respondents (two options: $\tilde{\chi}^2$ = 4.9, df = 1, p-value = 0.03; four options $\tilde{\chi}^2$ = 6.2, df = 3, p-value = 0.1).

Additionally, our analysis shows that average age for respondents who chose "Adjust settings" is higher than those who chose "Keep settings" and "Close screen" but not respondents who chose "Don't Know" option ("Adjust settings" = 34.7, "Keep settings" = 32.8, "Close screen" = 31.2, "Don't Know" = 34.7). ANOVA test shows a significant differences between these four options and age of respondents (F = 3.15, p-value = 0.02). However, pairwise comparison with Tukey correction does not show significant differences between any pairs of options. Nonetheless, the difference between age of respondents' who chose "Adjust settings" and respondents who chose "Keep settings" is statistically significant (F = 5.52, p-value = 0.02).

As for education level, our analysis shows no significant correlation between options chosen by respondents and education level (two options: $\tilde{\chi}^2$ = 3.3, df = 1, p-value = 0.07; four options: $\tilde{\chi}^2$ = 7.2, df = 3, p-value = 0.07).

4.2.5 Adjusting App Privacy Settings

Of all respondents, 617 (71.7%) chose to adjust their apps' access settings in response to the nudges. Subsequently, those respondents were presented with the screen shown in Figure 4.7 and were given three options: "Adjust apps' access settings", "Close the screen (by pressing the 'Home' or the 'Back' buttons)", or "I do not know." Of the 617 respondents, 584 (94.7%) chose "Adjust apps' access settings" option (the majority). Next, those respondents were presented with the three apps and were given the chance to indicate whether they want to allow or deny access to their location information by each app. All 584 respondents denied at least one of the apps from accessing their location information. Words With Friends was the most denied app with 568 (97.3%) respondents denying access to location information. Table 4.17 shows how respondents adjusted their app privacy settings.

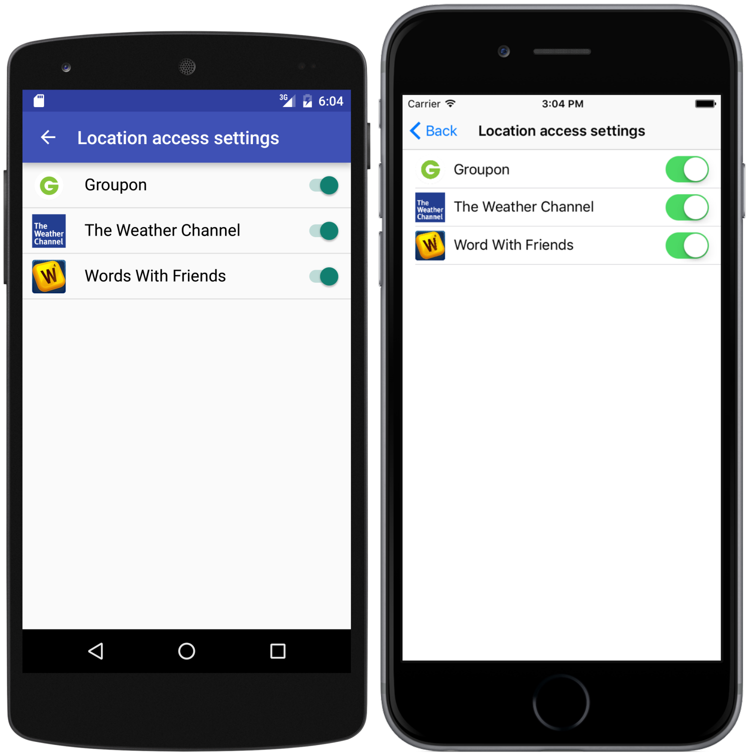

Location access settings

- Groupon ●
- The Weather Channel ●
- Words With Friends ●

Location access settings

- Groupon ●
- The Weather Channel ●
- Word With Friends ●

Figure 4.7: Apps location settings screens that were presented to respondents in our experiment (*left: Android, right: iOS*). The design of the screens is inspired by app permissions in Android and privacy controls in iOS.

Number of all respondents	861
Respondents who indicated their intention to adjust apps’ access settings	617 (71.7%)
Respondents who proceeded to adjust apps’ access settings	584 (94.7%)
The average number of apps which were denied access to location information per respondent	1.9 (out of 3 apps)
Number of respondents who denied one app only	192 (32.9%)
Number of respondents who denied two apps only	254 (43.5%)
Number of respondents who denied three apps	138 (23.6%)
Number of respondents who denied Groupon from accessing location information	393 (67.3%)
Number of respondents who denied The Weather Channel from accessing location information	153 (26.2%)
Number of respondents who denied Words With Friends from accessing location information	568 (97.3%)

Table 4.17: This table shows descriptive statistics about respondents’ decisions to adjust app privacy settings in response to the nudges.

4.3 Discussion

4.3.1 The Effectiveness of the “Recent Access” Nudge Content is Unsatisfactory

The “recent access” nudge content, which is used by Android and iOS, only informs users that their information has recently been accessed. Such a nudge content does not inform users of the contexts in which their information has been accessed, the purposes of accessing the information, or the potential implications of reusing users’ personal information. Prior work (including the results from Chapter 3) has shown that these missing pieces of information might resonate with users and motivate them to pay closer attention to the nudges [104, 118, 129, 137, 142, 173, 174, 194]. Our results suggest that indeed simply informing users that their information has been recently accessed is less effective in helping users make informed privacy decisions than providing users with more information about (unexpected) practices by apps.

In our experiment, we tested the baseline nudge content in an active setting. That is, the user was stopped from her primary task and was then shown the nudge. In reality, both Android and iOS provide the “recent access” app privacy nudge content (similar to the baseline in our experiment)

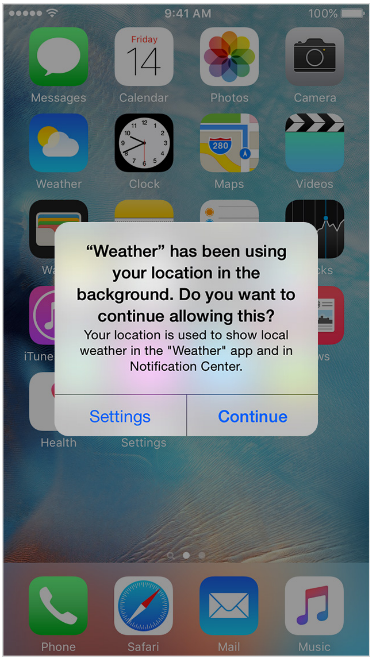

"Weather" has been using your location in the background. Do you want to continue allowing this?

Your location is used to show local weather in the "Weather" app and in Notification Center.

Figure 4.8: If an app on iOS uses the user's location in the background (i.e. while the user is not actively using the apps), the user will be shown a location privacy nudge (The source of the screenshot: [52]).

only in a passive setting. That is, the user needs to stop her primary task, switch her attention to privacy management, and then take the necessary steps to find the list of apps that has recently accessed her location information. In both Android and iOS, the user needs to take at least three to four steps to reach the location settings to find out those apps. Given that privacy and security are typically secondary tasks for users [198], prior work has shown that passive security and privacy warnings are less effective than active warnings [97, 197]. Thus, we expect the "recent access" nudge contents that are used in current mobile operating systems to be even less effective in motivating users to review their app privacy settings than what our experiment showed.

4.3.2 Additional Information Helps Users Make Informed Privacy Decisions

Our results show that the five different nudge contents that provide more information than the "recent access" nudge content are more effective in motivating users to review and possibly adjust their app privacy settings. These nudge contents inform users about context, purpose, and potential implications of accessing their personal information.

Of these five effective nudge contents, one nudge informs users of apps that have accessed the users' location even if they did not actively use the apps. This nudge is similar to and was inspired by a recent location privacy nudge introduced by iOS [48, 52]. If an app uses the user's location in the background (i.e. while the user is not actively using the apps), iOS will show an active location privacy nudge as shown in Figure 4.8. Although Apple has not published any results about the effectiveness of their privacy nudge, our results regarding the effectiveness of the background access nudge content suggest that the iOS nudge is effective as well. The effectiveness of the background nudge might be explained by two related concepts: privacy as contextual integrity [153, 194] and privacy as expectation [142, 167]. "Norms of information flow" when users are actively using apps might be different than when users are not using them (i.e. when the apps are running in the background) [153, 194]. Therefore, some users may not expect apps to access their sensitive information when they are not actively using the apps. When users are made aware of such practices through nudges, they act and adjust their app privacy settings.

The idea of giving users control over how apps access their information in the background is gaining traction in the mobile industry. For example, when iOS users make a decision to grant or deny apps access to their location information, they can set apps' location settings to "Always", "Only While Using [the app]", or "Never." To help users make informed privacy decisions, the upcoming iOS 11[11] nudges users when their location information is currently being used by an app in the background through a blue status bar at the top of the screen as Figure 4.9 shows [45]. This nudge may motivate more iOS users to change their apps' location settings from "Always" to "Only While Using" as analysts suggest [45][12]. Recent statistics by Apple showed that 21% of apps that use location information are allowed to access users' location all the time (i.e. apps' location settings are set to "Always") [45]. Additionally, inspection of Android 7.0 source code[13] shows that users' ability to control whether apps can run and, in turn, access users' information in the background may be coming to Android users in the foreseeable future [12].

Purpose-based nudge contents (which (1) inform users that their location information has been used for purposes other than providing app functionality, and (2) additionally highlight location-based advertising as an example of an unrelated purpose) were very effective. Furthermore, purpose-based nudge contents were more effective than the other two categories of nudge contents (i.e. behavior-based and implication-based nudge content) as our results show. In addition,

[11] iOS 11 is expected to be released Fall 2017.

[12] It is probably relevant to point out that iOS also shows a small arrow at the status bar when location information is requested by apps [47]. However, this arrow does not help users identify which apps have requested their location information.

[13] The inspection of the "AppOpsManager" class in particular shows that the ("OP_RUN_IN_BACKGROUND") operation has been added in Android 7.0 (Marshmallow). This operation enables the user to "[c]ontrol whether an application is allowed to run in the background." [12]

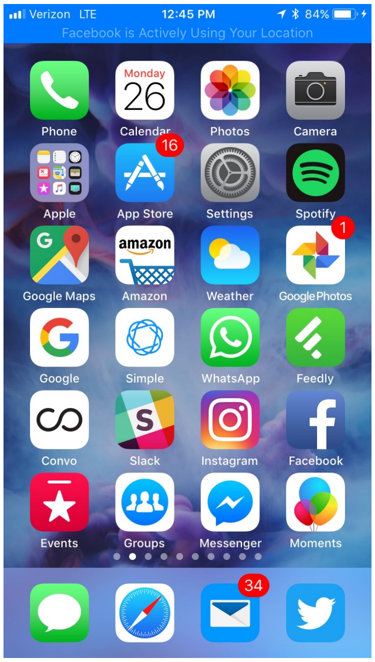

Figure 4.9: In the upcoming iOS 11, when an app is currently using the user's location in the background, iOS shows a blue status bar to inform the user (The source of the screenshot: [45]). In this example, the blue status bar reads: "Facebook is Actively Using Your Location."

the qualitative analysis of why respondents made certain decisions in response to the nudges showed that the purpose of apps' accessing users' personal information was an important factor that users considered when deciding whether to adjust app settings or keep them unaltered. Unfortunately, current mobile operating systems do not provide such a fine-grained control. Although users can selectively allow or deny an app from accessing their sensitive information, they cannot do so based on why the app needs to access users' information. As our results suggest, lacking such a fine-grained control makes users trade off functionality for privacy (or the other way around) which may lead to sub-optimal outcomes for users, app developers, and operating system providers.

The implication-based nudge contents showed mixed results. On the one hand, the nudge content in condition 6 (which reads "These apps have accessed your location in the past week. With this information, apps can infer additional details about you.") was not effective perhaps because it was a bit vague and did not provide specific implications. On the other hand, the nudge contents in conditions 8 (which highlights potentially predicting one's income) and 9 (additionally highlights potential price discrimination based on predicted income) were effective. Both highlight a sensitive type of personal information (i.e. income) and condition 9 additionally highlights (negative) implications of such a prediction. The nudge content in condition 7 (which highlights potentially inferring one's home address) is a bit puzzling. Although it highlights a sensitive type of personal information (i.e. home address) [187], it was not more effective than the baseline nudge content in our experiment. It is possible that respondents perceived home address as less sensitive (e.g., less sensitive than income). Future work may look closely at different wording and types of sensitive personal information and how these differences may effect users' decisions when responding to implication-based nudge contents.

4.3.3 Additional Benefits of Nudges

We designed app privacy nudges with two goals in mind: (1) increase users' awareness of privacy risks associated with their apps, and (2) switch users attention to privacy management. Our results show that the nudges indeed achieved these two goals.

Respondents in the treatment conditions were less aware of privacy risks and apps' data collection practices than their counterparts in the baseline condition, which simply informs users that their location information has been accessed by apps. In other words, nudge contents in the treatment conditions increased users' awareness of privacy risks associated with mobile apps. Such an increase in users' awareness correlated significantly with their decisions to review and adjust their app privacy settings. Users were more likely to choose to adjust their app privacy settings when they were less aware about what nudge contents conveyed. In sum, nudges increased users' awareness about privacy risks of mobile apps and in turn motivate them to review and adjust their privacy settings.

The majority of respondents in our experiment were motivated to review and adjust their app privacy settings in response to nudges including the least effective nudges. Furthermore, the qualitative analysis of why respondents made certain decisions in response to the nudges revealed a notable benefit of the nudges. Respondents wanted to review app privacy settings for other apps (i.e. other than what was listed in the nudge screen) and for other permissions (i.e. permissions other than location information). It seems that the nudges switched users attention to privacy management (at least briefly). Accordingly, respondents considered taking the opportunity to reviewing their privacy settings in general to ensure that they match respondents' privacy preferences. In other words, the influence of app privacy nudges may exceed their limited contents.

Chapter 5

User Engagement with Repeated App Privacy Nudges

Results from Chapter 3 and Chapter 4 indicate that runtime mobile app privacy nudges are a promising approach to help users manager their app privacy settings. App privacy nudges increase the utility of permission managers, highlight (unexpected and unknown) privacy risks associated with apps installed on users' phones, switch users' attention from their primary tasks to privacy management, and motivate users to review and possibly adjust their app privacy settings.

Nonetheless, results from Chapter 3 also suggest that receiving repeated nudges leads to less engagement with the nudges. About 56% of the first set of nudges led users to review their app permissions. In contrast, only 26% of the second set of nudges of the same data type led users to review their app permissions. It is possible that because users' preferences did not change within a short time or users reached a level of app privacy settings that users were comfortable with, they did not need to review their app permissions again and, thus ignored additional nudges. Alternatively, user engagement with repeated nudges might have decreased because users lost interest in them. Prior work has shown that users' attention drops sharply when they see a stimulus (e.g., a security warning) for the second time and onward [66, 67, 83, 179]. In our study, it is possible that the same effect occurred when users received nudges for the second time. However, the field study in Chapter 3 was not intended nor designed to understand *how* users respond to repeated nudges or *why* users respond in certain ways. Therefore, our focus here is to explore user engagement with repeated app privacy nudges.

This chapter has two objectives. The first objective is to examine whether and why user engagement with app privacy nudges decreases as users receive the nudges repeatedly. Because we hypothesized that user engagement with nudges decreases as users receive them repeatedly, the second objective of this chapter is to identify factors that may keep users engaged with repeated nudges. In particular, we aim to evaluate whether we can keep users engaged with repeated nudges by manipulating the contents of the nudges. In addition, we aim to identify other potential factors (that may keep users engaged with repeated nudges) based on users' behavior in situ and based on users' characteristics.

In this chapter, we describe and report results of a controlled mixed-design experiment in situ to gain insights on user engagement with repeated app privacy nudges and to evaluate approaches that may keep users engaged with repeated nudges. To this end, we conducted a six-day field controlled experiment in which participants installed a mobile app ("the study client") on their own Android-powered devices. During the period of the experiment, the study client sent two privacy nudges (three days apart) to inform participants of potential risks associated with apps installed on their phones. The nudges enabled participants to review apps' behavior and in turn decide to either adjust app privacy settings or keep them unaltered. Upon receiving the repeated nudge (the second nudge), participants were divided into two conditions: (1) participants who received a nudge with updated yet similar content to the first nudge, and (2) participants who received a nudge with a content that varies from the first nudge. We captured how participants responded to both nudges and we measured how participants responded to the repeated nudge (regardless of their condition) in comparison to the first nudge. In addition, we compared how participants in condition 1 (same nudge content) and participants in condition 2 (different nudge content) responded to the repeated nudge. To better understand why participants responded to the nudges, we also asked participants to complete an online survey about their experience with the nudges during the field experiment. In this chapter, we report the results of the experiment and the lessons we learned about user engagement with repeated nudges.

5.1 Methodology

5.1.1 Research Questions

In this experiment, we focused on two research questions.

- How do mobile users respond to repeated app privacy nudges?
- Can we keep users engaged with repeated nudges by updating or varying nudge contents? Which is more effective in keeping users engaged with repeated nudges, updating or vary-

ing nudge contents?

- What factors may motivate users to engage with repeated nudges?

5.1.2 Engagement Reduction & Repeated App Privacy Nudges

Prior work has shown that receiving a stimulus (e.g., a security warning or a privacy nudge) repeatedly reduces user engagement [64, 66, 67, 80, 82, 83, 84, 97, 179]. Attention reduction has been explored and shown when users received warnings for the second time [83, 179, 183], forth time [83], 12th time [184], 20th time [83], 23rd time [82], and 40th time [136]. Recent work in neuroscience has utilized more accurate approaches (e.g., fMRI imaging) to explore how the brain reacts when interacting with repeated stimuli [66, 67]. Such studies have shown that users' attention drops sharply when they see a stimulus (e.g., a security warning) for the second time and onward [66, 67]. In other words, these studies suggest that detecting a drop in users' attention to and engagement with stimuli (e.g., a security warning or a privacy nudge) is attainable by exposing users to a stimulus only twice. We build on these results to explore how users react to repeated app privacy nudges. In particular, we posit that we can measure reduction in user engagement and develop approaches to delay or minimize such a reduction by sending users two app privacy nudges. Earlier work has also suggested that changing aspects of the warning may delay or minimize the expected reduction in user engagement [67, 82, 83, 87, 136, 166]. Presuming that the brain may recall that the second app privacy nudge has been seen earlier (i.e. through the first nudge) [66, 67], we also evaluate how users react to a second but different nudge.

It is important here to point out two main issues. First, in the context of app privacy nudges, reduction in user engagement with repeated nudges might not be only attributed to earlier exposure to nudges. Rather, user engagement might also be reduced because the user has already reached a level of privacy settings that she is comfortable with after adjusting her settings in response to the first nudge. Reaching a comfortable level of privacy preferences may become attainable especially in the context of mobile app privacy given that apps' settings are finite. The second issue is that our study does not address the long-term habituation effect of app privacy nudges. Rather, our results may inform future work in understanding why and how user engagement with nudges changes upon receiving a repeated nudge.

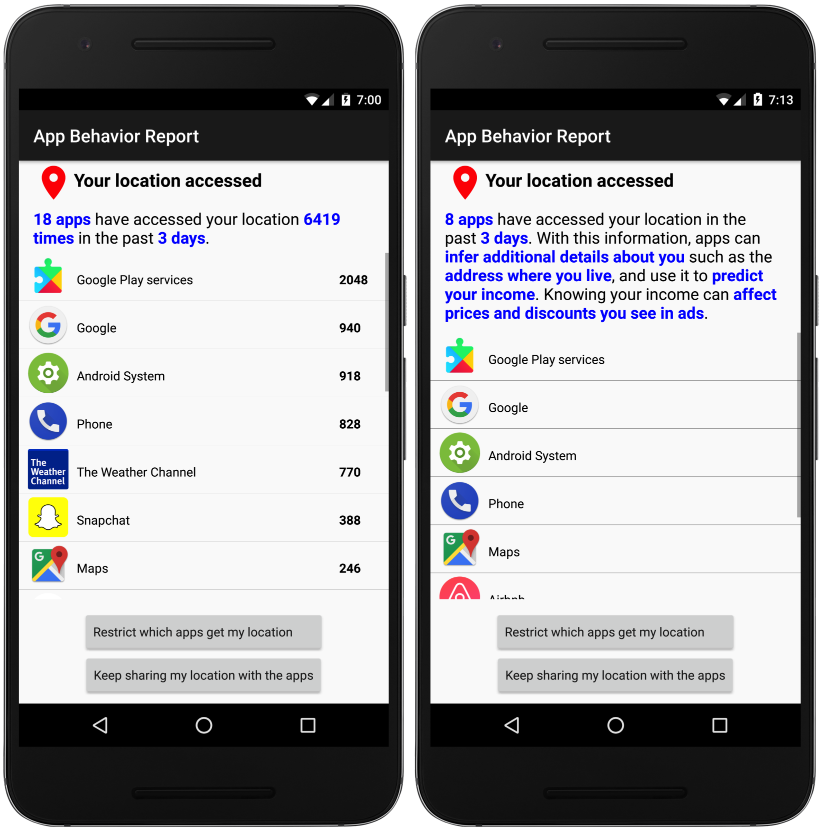
7:00
App Behavior Report
Restrict which apps get my location
Keep sharing my location with the apps

Your location accessed

18 apps have accessed your location 6419 times in the past 3 days.

Google Play services	2048
Google	940
Android System	918
Phone	828
The Weather Channel	770
Snapchat	388
Maps	246

7:13
App Behavior Report
Restrict which apps get my location
Keep sharing my location with the apps

Your location accessed

8 apps have accessed your location in the past 3 days. With this information, apps can infer additional details about you such as the address where you live, and use it to predict your income. Knowing your income can affect prices and discounts you see in ads.

- Google Play services
- Google
- Android System
- Phone
- Maps
- Airbnb

(a) The frequency nudge (b) The implications nudge

Figure 5.1: A screenshot of the frequency nudge (left) and the implications nudge (right).

5.1.3 Nudges Design

One goal of our experiment is to evaluate whether a subtle or a major change in nudge content will delay or minimize engagement reduction when users receive repeated app privacy nudges. Thus, we selected two nudges that we designed for the previous experiment in Chapter 4: frequency of access and potential implications of secondary usage of personal information. Figure 5.1 shows both nudges. We chose nudges that are drastically different because we want to evaluate whether a repeated but a clearly different content of a nudge would garner users' attention more than a repeated but a slightly different nudge content (e.g., the frequency of accessing user information increases or the list of apps which access users information expands). Thus, we can evaluate whether changing the nudge content in a subtle or major way may delay or minimize engagement reduction in response to repeated nudges.

We chose the frequency nudge for a practical reason. The frequency nudge is easily generated through the logs of the "App permissions" manager (as we describe in detail in Section 5.1.4). In addition, the look and feel of the nudge seems similar when received repeatedly, but the content gets updated as the frequency of accessing user information may increase over time or the list of apps which access the user's location information expands. Furthermore, we have shown in Chapter 3 that the frequency nudge resonated with users in the field and effectively motivated them to review and adjust their app privacy settings. Although the frequency nudge was not significantly more effective than the baseline nudge (which simply informs users that their location information has been recently accessed by apps) as shown in Chapter 4, the frequency nudge still provides more information. Therefore, we decided to use the frequency nudge as a baseline for our experiment.

The implications nudge has shown to be very effective in Chapter 4. We also chose the implications nudge for the practical reason that it does not require external information about app practices that may or may not be available at runtime (e.g., additional static or dynamic analysis such as TaintDroid [99] or PrivacyGrade [31] may be required to point out apps that collect users' location information for advertising purposes). Rather, the implications nudge brings users' attention to how apps may reuse location information that has been collected.

As Figure 5.1 shows, each nudge consists of the location icon, the title, the main content of the nudge, the list of apps, and the buttons to act upon the nudge.

The location icon informs the user immediately that the screen is about location information. We used an icon with which users are familiar [162] and that has been used by maps and navigation applications in Android.

The title of the nudge reads: “Your location accessed.” The title is grammatically condensed by removing auxiliary verbs (i.e. “has been”). This was a deliberate decision to ensure that the title is short and fits in one line. In the pilot testing, we evaluated the title closely and found that users were able to easily comprehend its meaning.

The main content of the nudge matches one of the two conditions that participants were assigned to: frequency or implications. As Figure 5.1 shows, we highlighted some phrases in blue on the nudge screen to guide participants’ attention to the essential parts of the nudge content. We used blue color for highlighting because it has been shown to be a neutral color [175] and thus we avoided introducing an additional confounding factor.

The nudge also lists apps that accessed the user’s location within a period of time (e.g., three days). Each row shows the app’s icon and name. When participants received the frequency nudge, we showed frequency of accessing location information by individual apps and listed the apps in a descending order based on how frequently each app accessed participants’ location information. When participants received the implications nudge, we similarly listed the apps in a descending order based on how frequently each app accessed participants’ location information although the frequencies were not shown to participants. We made the decision to list apps in the implications nudge based on frequencies (although frequencies were not shown) so that the list of apps would be consistent across conditions. Thus, we avoided introducing an additional confounding factor.

In each nudge, the respondents are given two buttons that correspond to two options: (1) “Restrict which apps get my location” (AdjustSettings) to redirect the participant to the “App permission” manager, and (2) “Keep sharing my location with the apps” (KeepSetting) to keep things as is and close the screen. The captions of the two buttons are intentionally long because we want to ensure that respondents can easily comprehend what the buttons do. Furthermore, we used “enhanced active choice” when we chose the wording of the second button (i.e. the “Keep sharing my location with the apps” button). Keller et al. [130] recommend employing enhanced active choice to emphasize the desired option (i.e. the “Restrict which apps get my location” button) by “highlighting the losses incumbent [e.g., apps will continue accessing your location] in the non-preferred alternative [i.e. the “Keep sharing my location with the apps” button].”

If the participant chooses AdjustSettings, the participant should ideally be redirected to the “App permissions” manager to review and possibly adjust their app privacy settings. However, Android prevents any 3rd party apps from directly opening the “App permissions” manager. To overcome

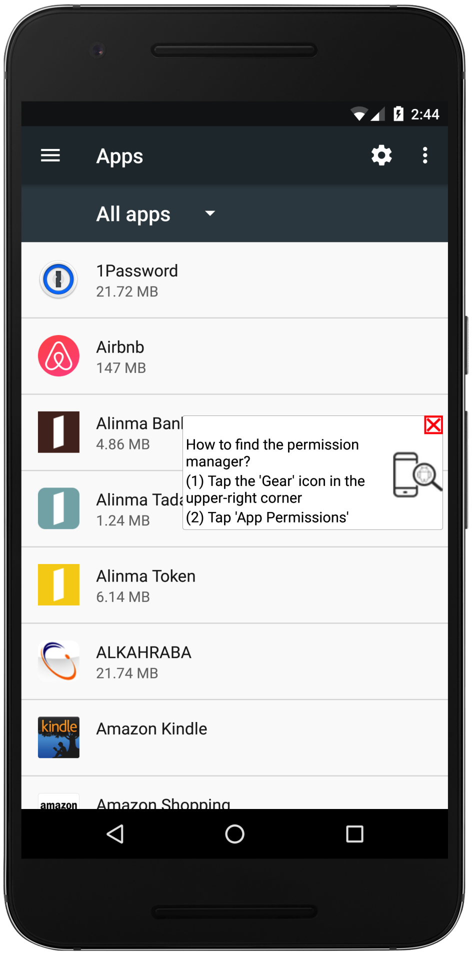

How to find the permission manager?
(1) Tap the 'Gear' icon in the upper-right corner
(2) Tap 'App Permissions'

Figure 5.2: Android does not allow a 3rd party app to directly open the “App permissions” manager. Thus, when a participants receives a nudge and clicks the “Restrict which apps get my location” button, the user is redirected to the “App” settings screen. To help the user find the “App permissions” manager, we show a small window that shows how the participant can navigate to the “App permissions” manager.

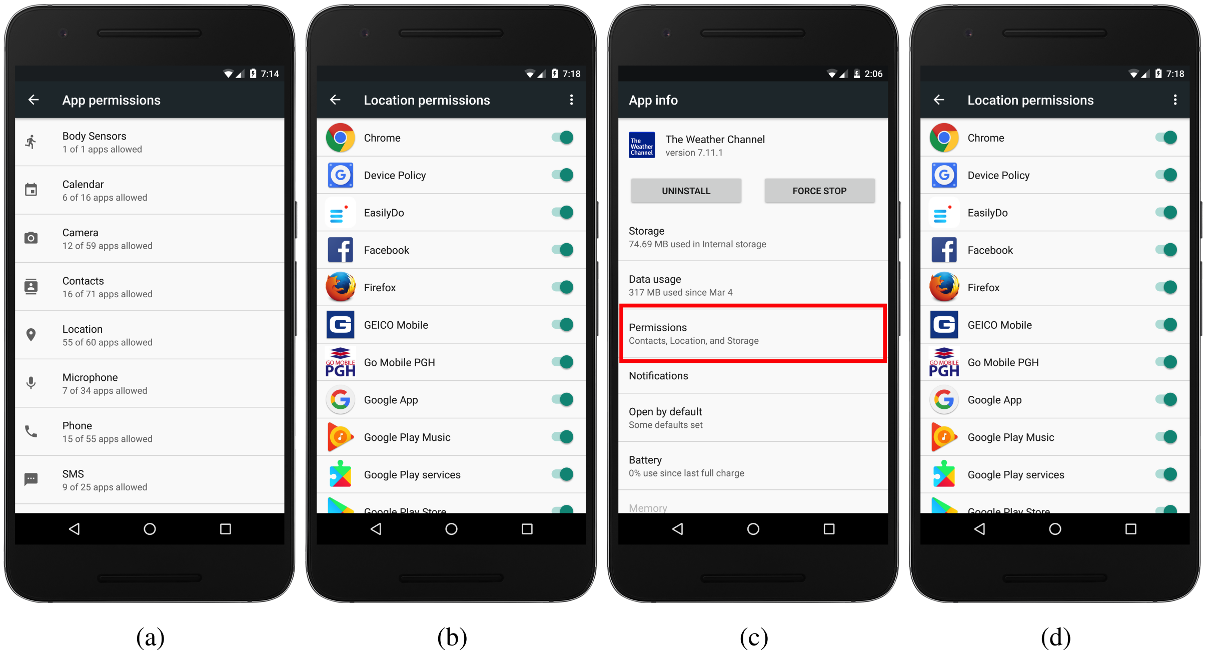

7:14

App permissions

Body Sensors
1 of 1 apps allowed

Calendar
6 of 16 apps allowed

Camera
12 of 59 apps allowed

Contacts
16 of 71 apps allowed

Location
55 of 60 apps allowed

Microphone
7 of 34 apps allowed

Phone
15 of 55 apps allowed

SMS
9 of 25 apps allowed

(a)

7:18

Location permissions

- Chrome
- Device Policy
- EasilyDo
- Facebook
- Firefox
- GEICO Mobile
- Go Mobile PGH
- Google App
- Google Play Music
- Google Play services
- Google Play Store

(b)

2:06

App info

The Weather Channel
version 7.11.1

Storage
74.69 MB used in Internal storage

Data usage
317 MB used since Mar 4

Permissions
Contacts, Location, and Storage

Notifications

Open by default
Some defaults set

Battery
0% use since last full charge

Memory

(c)

7:18

Location permissions

- Chrome
- Device Policy
- EasilyDo
- Facebook
- Firefox
- GEICO Mobile
- Go Mobile PGH
- Google App
- Google Play Music
- Google Play services
- Google Play Store

(d)

Figure 5.3: Android Marshmallow (Android OS version 6) and above provide users with tools to manage permissions of individual apps. The "App permissions" manager (a) lists categories of permissions that apps requested access to. If the user clicks the location category, the user is presented with location permissions (b). The user can also adjust permissions for individual apps through the Settings screen (c). If the user selects the Permissions setting (highlighted in the red rectangle), the user is presented with all permissions for this particular app (d).

this challenge and given that a participant may not know how to find the "App permissions" manager, we designed a small window that flows on top of other apps and guides the participant to find the permission manager. When a participant chooses AdjustSettings, the participant is redirected to the "App" settings screen and is shown the window, which guides the participant to find the "App permissions" manager as shown in Figure 5.2.

5.1.4 Implementation of Study Client

In Android Marshmallow (Android OS version 6.0), Google introduced the "App permissions" manager (as shown in Figure 5.3, left) which allows users to selectively grant or deny permissions for installed apps [5, 43]. When an app requests access to a sensitive resource for the first time (e.g., location information), the user is prompted to either grant or deny such a request. If the request is granted, the app can access the resource anytime. The user can revert the decision anytime either through the "App permissions" manager (as shown in Figure 5.3, left) or the per-app permission setting (as shown in Figure 5.3, right).

Our study client collected information about which apps accessed location information and how often, which was used to generate personalized privacy nudges for each participant's phone. The required information was obtained by periodically recording logs created by the "App permissions." The logs are organized by app-permission pairs. For each app-permission pair, the logs show whether the app is granted or denied access to the permission and the last time the app was granted or rejected access to the permission. 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 the "App permissions" manager. Collecting the logs and the permission changes had to be recorded periodically, since the "App permissions" manager does not permit 3rd party apps to access specific interaction events, and modifying the "App permissions" manager would have required rooting and flashing participants' devices, which we deemed undesirable. Accessing the logs of the "App permissions" manager requires a one-time runtime permission ("GET_APP_OPS_STATS"), which can only be granted through the Android Debugging Bridge (ADB) tool if the device is connected via USB after app installation.

In addition to recording access frequency and permission changes, the study client presented participants with two nudges and recorded how participants interacted with them. For each nudge, the study client recorded the option chosen by the participant and the length of the interaction session. 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.

5.1.5 Experiment Design

Our main goal in this study was to measure how users react to a repeated nudge. We hypothesized that users engagement will decrease when receiving the second, repeated nudge. Thus, we wanted to explore whether changing the nudge will delay or minimize such a reduction in user engagement. To this end, we conducted a six-day controlled experiment in situ. Participants installed the study client on their own Android-powered devices and received two app privacy nudges, three days apart. The three days enabled the study client to collect a reasonable amount of data about apps' location accesses. Simultaneously, participants received the repeated nudge within a relatively short time to ensure that they were still familiar with the nudges [183]. For each nudge, we captured how participants responded to the nudge, the length of the interaction session, and the app permissions adjustments that participants made.

In our experiment, we used a mixed-design: within-subject and between-subject. All participants received the same first nudge, the frequency nudge. Then, participants were randomly assigned to one of two conditions: (1) updating the nudge's content, or (2) varying the nudge's content.

In condition 1 (updating the nudge's content), participants again received the frequency nudge although the frequency of access had been updated to reflect apps' behavior in the past three days. In addition, the list of apps in the repeated nudge might have differed from the apps in the first nudge (e.g., an app accessed the user's location information in the last three days of the experiment but never accessed the user's location information in the first three days of the experiment). In condition 2 (varying the nudge's content), participants received a different nudge content, the implications nudge content, which highlights potential implications of secondary usage of personal information. The design of the experiment enabled us to measure how participants respond to the repeated nudge, regardless of its content. In addition, the design enabled us to compare how participants in condition 1 (updating content) versus participants in condition 2 (varying content) respond to the repeated nudge.

App privacy nudges in our experiment focused particularly on location information. Prior work has shown that users are typically concerned about sharing their location information with mobile apps (e.g., [81, 104]). Location information is also one of the most frequently requested piece of information by apps [102, 135]. We focused only on one type of information (i.e. only location information) rather than multiple types to ensure that we have a manageable number of conditions in our experiment.

5.1.6 Exit Survey

To further understand why users responded to the nudges in certain ways during the field experiment, we asked participants to complete an online survey. We designed the survey to be tailored to how participants responded to the nudges during the field study. We describe the survey briefly in this section.

The survey asked participants two consecutive open-ended questions about why they responded in certain ways to the first nudge and repeated nudge, respectively. For example, if the participant responded to the first nudge by choosing AdjustSettings, we presented this response and asked the participant why she responded this way. We asked a similar question about the repeated nudge. The goal of these questions is to understand the reasoning behind participants' decisions when interacting with the nudges. Then, we asked a third open-ended question to understand why participants responded similarly or differently to both nudges. For example, if a participant responded to the first nudge by choosing AdjustSettings and the repeated nudge by choosing KeepSettings, we presented both responses and asked why the participant responded differently. The goal of this question is to understand how participants perceived the interaction with the first and repeated nudge in comparison to one another.

To further understand participants behaviors during the field study, the survey also asked how participants perceived receiving a repeated nudge. For example, we asked participants to specify whether they agree or disagree (through a five-point likert scale) with four statements about receiving a repeated nudge (e.g., "Although report 1 was received a few days earlier, report 2 was necessary.", "Although report 1 was received a few days earlier, report 2 was unnecessary."). We also asked participants to compare the two nudges in terms of which is more likely to motivate participants to review and possibly adjust their app privacy settings. In addition, we hypothesized that awareness of and concern about what the nudges convey may affect how participants respond to the nudges. Therefore, the survey additionally asked participants to specify (through five-point likert-scale questions) their level of awareness of and concern about what each nudge conveyed.

Finally, the survey asked how familiar participants were withe the "App permissions" manager and the per-app permission setting and whether participants had used them before enrolling in our experiment. The survey concluded by asking participants about their privacy concern level through the Internet Users' Information Privacy Concerns (IUIPC), and by asking demographic questions.

5.1.7 Study Procedure

The study consisted of three phases: an entry phase, a six-day field experiment, and an exit survey. We collected behavioral data through the field experiment and supplemented this data with users' responses and explanations from the exit survey. Here, we describe in detail the procedure of our experiment.

Entry phase

Participants were invited to our lab to sign the consent form and to install the study client. When a participant arrived at our lab, a researcher handed the participant a consent form to read and sign. Then, the researcher randomly assigned the participant to either conditions 1 (updating the nudge's content) or 2 (varying the nudge's content) and installed the study client into the participant's phone. The study client was also deployed to the Google Play store to ensure that updates (if exist) are pushed automatically to participants during the study. After installation, the study client was granted the one-time runtime permission ("GET_APP_OPS_STATS") to collect the logs of the "App permissions" manager. In addition, the study client was granted permissions through the “Settings” app to collect information about apps’ usage[1], and to overlay the window[2] that guides participants to find the “App permissions” manager.

The study client was designed carefully and was tested in the lab and during the pilot study to ensure that it did not drain battery. The study client had to continue working in the background to collect data and send them to our servers. However, Android manages battery life by adding all apps (except a small number of system services) by default to the “Battery optimization” list. Battery optimized apps that are not active or have not been used by the user for sometime may receive fewer operating system resources (recall that participants were not expected to interact with the study client regularly except for brief sessions with two nudges, three days apart) [13, 14]. Therefore, we removed the study client from the battery optimization list (via the “Settings” app) to ensure its operations were not disrupted.

The entry phase took less than ten minutes.

Six-day Field Experiment

After installing the study client, participants went about their daily life and used their own phones as they usually do without interruption. Simultaneously, the study client worked silently in the background to collect data.

On the third day at 7pm, all participants received the first nudge, the frequency nudge (as shown in Figure 5.1), which showed how frequently apps had accessed participants’ location information in the past three days. The study client recorded how the participants interacted with the nudge and the length of the interaction session. If the participant ignored the nudge by pressing the home button to minimize the screen, the nudge was presented again in 15 minutes.

On the sixth day at 7pm, the participants received the repeated nudge. If the participant was in condition 1 (updating content), the participant received a frequency nudge similar to the first nudge (as shown in Figure 5.1), but with the frequencies updated to reflect how many times apps had accessed location information in the past three days. In addition, if a participant adjusted permissions for an app in response to the first nudge (e.g., denied an app access to location information in response to the first nudge), this app was not listed in the repeated nudge so that the participant was only presented with apps that might need further permissions’ revision. If the

[1] Although the study client declares this permission (“PACKAGE_USAGE_STATS”) in the manifest, it must be additionally enabled through the “Settings” app [11]

[2] Although the study client declares this permission (“SYSTEM_ALERT_WINDOW”) in the manifest, it must be additionally enabled through the “Settings” app [10]

participant was in condition 2 (varying content), the participant received the implications nudge (as shown in Figure 5.1), which showed the list of apps that had accessed the participant's location information in the past three days, and the potential implications of such behavior. Similar to participants in condition 1, if a participant adjusted permissions for an app in response to the first nudge (e.g., denied an app access to location information in response to the first nudge), this app was not listed in the repeated nudge. In both conditions, if the participant ignored the repeated nudge by pressing the home button to minimize the screen, the nudge was presented again in 15 minutes.

Twenty minutes after interacting with the repeated nudge, the participant was prompted through a regular Android notification to a screen in the study client. The screen guided the participant to send any collected data that had not yet been sent to our servers.

Exit Survey

After completing the field experiment, each participant was sent a link to complete an online survey about their experience during the field phase. As described in Section 5.1.6, the survey was tailored to how a participant interacted with the first and repeated nudge. After completing the survey, participants were compensated with a $20 Amazon gift card.

5.1.8 Recruitment

We conducted our study during July, August, September, and October in 2017. We recruited participants through a city-wide participants pool maintained by our university. We also recruited participants by using the data truck[3] which parks in different spots with the city of Pittsburgh, PA. We recruited participants who meet all of the following conditions: (1) are 18 years old or older, (2) have smartphones running Android Marshmallow (Android OS version 6) or above, (3) have Android versions that are not rooted (e.g,, CynogenMood), (4) have a mobile data plan of at least 2GB/month, (5) are able to visit our lab for the entry session, and (6) are speakers of the English language. We targeted participants who have phones running Android 6 and above because they are shipped with the "App permissions" manager. By June 5, 2017, the distribution of Android version is as follows: Marshmallow (31.2%), Nougat 7.0 (8.9%), and Nougat 7.1 (0.8%) [3]. We excluded participants who have rooted versions of Android to avoid potential bias in our targeted population (e.g., being more tech savvy than the wider population of Android users).

[3] The data truck is provided by the Center for Behavioral and Decision Research at Carnegie Mellon: `https://cbdr.cmu.edu/`

5.1.9 Limitations

Our study client collected logs from the "App permissions" manager every five minutes. Thus, if a participant made more than one app setting adjustments within the five-minute interval, our client only captured the latest decision. However, it is reasonable to assume that multiple app setting adjustments within a five-minute interval are rare.

Another limitation is how the "App permissions" manager generates logs for apps that target Android software development kit (SDK) versions prior to 23 (i.e. target Android versions prior to introducing the "App permissions" manager in Android Marshmallow). For these apps, the logs do not show whether the app is currently granted or denied access to a permission. Rather, it only shows the last time a request to a permission was granted or rejected. Thus, if a participant granted an app access to a permission but the app had never requested that permission, the logs cannot help us figure out that the permission was actually granted. Similarly, if a participant denied an app access to a permission but the app had never requested that permission, we cannot figure out from the logs that the permission was actually denied.

Finally, we want to acknowledge the potential self-selection bias in our pool of participants.

5.2 Results

5.2.1 Demographics

A total of 122 participants completed our experiment. We excluded seven participants who stated in the exit survey that they misunderstood how the experiment was supposed to work. These six participants assumed that sharing location information with apps was a requirement to participate in our study and, therefore, whenever they received a nudge, they chose the same option: "Keep sharing location information with the apps" button. As such, we ended up with 115 participants, and we focus on these participants throughout the results section.

Of the 115 participants, 62 (53.9%) are females. The age range is 18–69 and the average is 28.3 (SD = 12.5, median = 23). Fifty-five (47.8%) participants have bachelor's degrees or above. Of all participants, 71 (61.74%) are still students (both graduate and undergraduate), 39 (33.91%) are employed/self-employed, and 5 (4.35%) are unemployed or retired. Although we recruited participants from around the city of Pittsburgh (PA), our sample is skewed toward young college students.

App Permission Tools in Android	Yes (%)	No (%)	Not sure (%)
App Permissions Manager	66 (57.4%)	34 (29.6%)	15 (13%)
Per-app Permission Setting	66 (57.4%)	42 (36.5%)	7 (6.1%)

Table 5.1: Participants self-reported whether they had ever used any of the app permission tools available on their Android-powered phones before enrolling in our study. The majority of our participants had used the app permission tools before participating in our experiment.

Android users can adjust their app permissions anytime either through the "App permissions" manager (as shown in Figure 5.3, left), or through the per-app permission setting (as shown in Figure 5.3, right). In the exist survey, we asked participants whether they have used either of the two approaches to manage their app permissions. As Table 5.1 shows, 66 (57.4%) participants had used the permission manager and the per-app permission setting, whereas 49 (42.6%) participants either had never used them or were unsure. In addition, we used the Internet Users' Information Privacy Concerns (IUIPC) scale to gauge the level of privacy concern among our participants. As Figure 5.4 shows, our participants seem to be skewed toward being concerned with their privacy.

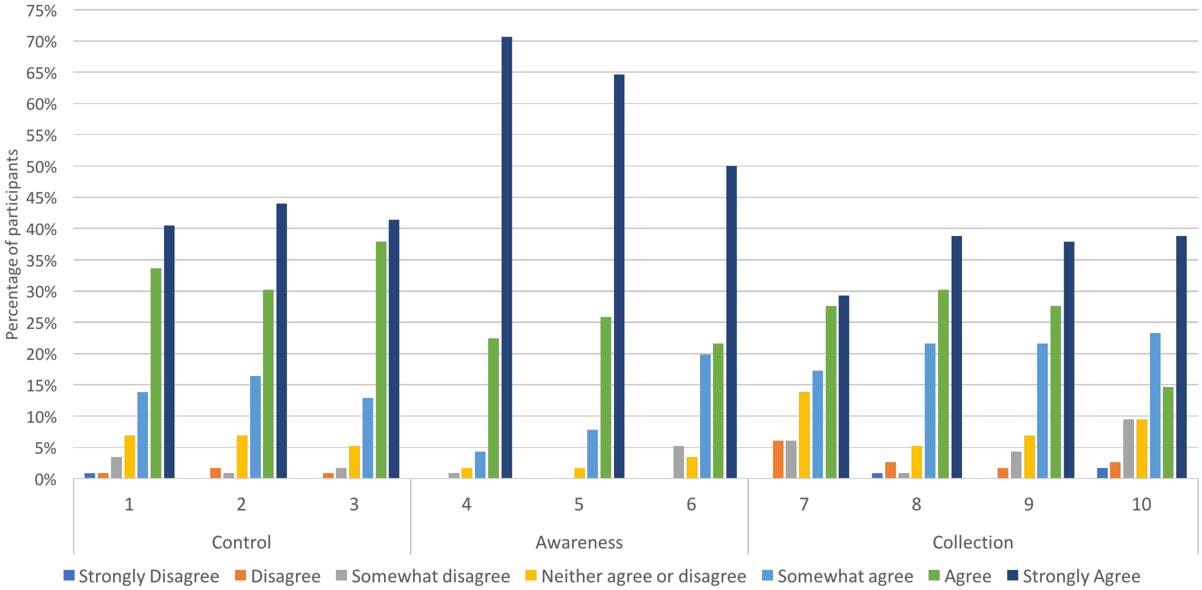

Figure 5.4: Participants' privacy concerns through the IUIPC privacy scale. Self-reported responses to the IUIPC scale questions show that participants tend to be skewed toward being concerned about their privacy.

Participants were randomly assigned to one of two conditions: updating the nudge's content (condition 1), and varying the nudge's content (condition 2). The number of participants in condition 1 is 57, and the number of participants in condition 2 is 58. Between the two conditions, the distribution of gender, age, education, occupation (student vs. non-student), IUIPC is not statistically significant (gender: $\tilde{\chi}^2$ = 2.51, df = 1, p-value = 0.11; age: t = 1.18, p-value = 0.24; education: $\tilde{\chi}^2$ = 0.008, df = 1, p-value = 0.92; occupation: $\tilde{\chi}^2$ = 2, df = 1, p-value = 0.16; IUIPC Control: t = 0.11, p-value = 0.92; IUIPC Awareness: t = -0.56, p-value = 0.57; IUIPC Collection: t = -0.68, p-value = 0.5).

5.2.2 Participants' Engagement with the Nudges

As Figure 5.1 shows, a participant has two options for interacting with the nudge: (1) the "Restrict which apps get my location" button (RestrictSettings) to redirect the participant to the "App permissions" manager to further review and possibly adjust app settings, and (2) the "Keep sharing my location with the apps" button (KeepSettings) to keep things as is and close the screen. We analyzed how participants engaged with nudges by looking at their intention to review/adjust app settings in response to nudges, the actual app setting adjustments that they made, the length of their interaction sessions with the nudges, and participants' reasoning to engage with nudges in certain ways through their responses in the exist survey.

Intention to Review/Adjust App Settings (IRAS)

Intention to review/adjust app settings (IRAS) represents whether a participant chose RestrictSettings when she received a nudge, regardless of whether the participant followed through and made any app setting adjustments after being redirected to the "App permissions" manager.

		Repeated Nudge		
		AdjustSettings	KeepSettings	
First Nudge	AdjustSettings	36 (31.3%)	24 (20.9%)	60 (52.2%)
	KeepSettings	10 (8.7%)	45 (39.1%)	55 (47.8%)
		46 (40%)	69 (60%)	

Table 5.2: How participants indicated their intention to review/adjust their app settings (IRAS) in response to the first nudge in comparison to the repeated nudge. Fewer participants indicated their IRAS in response to the repeated nudge than to the first nudge. Also, participants who indicated their IRAS in response to the first nudge were more likely to stick with the same decision when they received the repeated nudge.

As Table 5.2 shows, 70 (60.9%) participants indicated their IRAS in response to both or either of the nudges by choosing AdjustSettings, whereas 45 (39.1%) participants responded to both nudges by choosing KeepSettings.

Our main objective in this experiment is to explore how participants responded to the repeated nudge in comparison to the first nudge. We hypothesized that participants' engagement with nudges will decrease when they receive a repeated nudge. As Table 5.2 shows, 60 (52.2%) participants indicated their IRAS in response to the first nudge, whereas 46 (40%) did so in response to the repeated nudge. The difference is statistically significant (McNemar's $\tilde{\chi}^2$ = 4.97, df = 1, p-value = 0.026). Indeed, the results support our hypothesis and suggest that participants are less likely to indicate their IRAS in response to a repeated nudge than to a first nudge.

		Repeated Nudge		
		AdjustSettings	KeepSettings	
First Nudge	AdjustSettings	20 (35.1%)	10 (17.5%)	30 (52.6%)
	KeepSettings	3 (5.3%)	24 (42.1%)	27 (47.4%)
		23 (40.4%)	34 (59.6%)	

Table 5.3: How participants in Condition 1 (updating the nudge's content) indicated their intention to review/adjust their app settings (IRAS) in response to the first nudge in comparison to the repeated nudge. Although fewer participants in condition 1 indicated their IRAS in response to the repeated nudge than to the first nudge, the difference is not statistically significant.

We further explored whether the same results hold when analyzing how participants in each condition responded to the repeated nudge in comparison to the first nudge.

Table 5.3 shows results for participants in condition 1 (updating content). Thirty (52.6%) participants in condition 1 indicated their IRAS in response to the first nudge, whereas 23 (40.4%) participants did so in response to the repeated nudge. Although the percentages are similar to the ones for all participants combined, the difference is not statistically significant (McNemar's $\tilde{\chi}^2$ = 2.77, df = 1, p-value = 0.096). Similarly, 30 (51.7%) participants in condition 2 (varying content) indicated their IRAS in response to the first nudge, whereas 23 (39.7%) participants did so in response to the repeated nudge as Table 5.4 shows. Again, the difference is not statistically significant although the percentages are similar to the ones for all participants combined (McNemar's $\tilde{\chi}^2$ = 1.71, df = 1, p-value = 0.19). The results within each condition do not support our hypothesis. That is, the results do not show that participants within each condition are less likely to engage with the repeated nudge than the first nudge. Alternatively, given that the percentages within each condition are similar to the percentages for all participants combined (i.e. regardless of their assigned condition), it is possible that the absence of significant difference within each condition is due to the lack of statistical power as the number of participants within each condition is relatively small (e.g., see the numbers of participants in some cells in Table 5.3 and Table 5.4).

We also looked at how participants in condition 1 (updating content) indicated their IRAS in comparison to participants in condition 2 (varying content). Twenty-three (40.4%) participants in condition 1 indicated their IRAS in response to the repeated nudge, whereas 23 (39.7) participants in condition 2 did so in response to the repeated nudge. The difference between the two conditions is not statistically significant ($\tilde{\chi}^2 < 0.001$, df = 1, p-value = 1). It could be the case that varying nudge contents does not influence users' decisions when receiving repeated nudges as our statistical analysis suggests. Alternatively, the difference between nudge contents may have not been conspicuous enough to trigger different responses at least for some participants due to the nearly identical look and feel of the two nudge contents as Figure 5.1 shows. As a result, varying the nudge's content did not show its expected impact. In the exist survey, some participants pointed out that they did not notice the difference between the two nudge contents: "I wanted to change the app settings. I don't really care that ads know my location, but it does surprise me how many do. I didn't even realize the report messages were different until now."

		Repeated Nudge		
		AdjustSettings	KeepSettings	
First Nudge	AdjustSettings	16 (27.6%)	14 (24.1%)	30 (51.7%)
	KeepSettings	7 (12.1%)	21 (36.2)	28 (48.3%)
		23 (39.7%)	35 (60.3%)	

Table 5.4: How participants in Condition 2 (varying the nudge's content) indicated their intention to review/adjust their app settings (IRAS) in response to the first nudge in comparison to the repeated nudge. Although fewer participants in condition 2 indicated their IRAS in response to the repeated nudge than to the first nudge, the difference is not statistically significant.

Adjusting App Settings (AAS)

In Android Marshmallow (Android OS version 6) and above, the permission model uses an opt-in policy. That is, apps are by default denied access to certain sensitives permissions (see Table 5.5). To get access to these sensitive permissions, an app requests access to a sensitive permission and then the user is prompted to allow or deny such a request. The user can revert the decision anytime through the permission manager or through the per-app setting as Figure 5.3 shows.

Adjusting app settings (AAS) represents whether a participant actually adjusted app settings in response to a nudge after being redirected to the permission manager. Our focus here is on *restrictive adjustments* (i.e., restricting an app's access to a permission after the app was permitted access to a permission). If a participant adjusted app settings within 20 minutes of

Permission Group	Individual Permissions
CALENDAR	android.permission.READ_CALENDAR android.permission.WRITE_CALENDAR
Camera	android.permission.CAMERA
CONTACTS	android.permission.READ_CONTACTS android.permission.WRITE_CONTACTS android.permission.GET_ACCOUNTS
LOCATION	android.permission.ACCESS_FINE_LOCATION android.permission.ACCESS_COARSE_LOCATION
MICROPHONE	android.permission.RECORD_AUDIO
PHONE	android.permission.READ_PHONE_STATE android.permission.READ_PHONE_NUMBERS android.permission.CALL_PHONE android.permission.ANSWER_PHONE_CALLS android.permission.READ_CALL_LOG android.permission.WRITE_CALL_LOG android.permission.ADD_VOICEMAIL android.permission.USE_SIP android.permission.PROCESS_OUTGOING_CALLS
SENSORS	android.permission.BODY_SENSORS
SMS	android.permission.SEND_SMS android.permission.RECEIVE_SMS android.permission.READ_SMS android.permission.RECEIVE_WAP_PUSH android.permission.RECEIVE_MMS
STORAGE	android.permission.READ_EXTERNAL_STORAGE android.permission.WRITE_EXTERNAL_STORAGE

Table 5.5: We used permission groups introduced by Android Marshmallow (OS version 6) and above to count the number of setting adjustments made by participants in response to app privacy nudges. For example, if a participant denies an app access to location information, both ACCESS_COARSE_LOCATION and ACCESS_FINE_LOCATION permissions will be denied. However, we count this two permission adjustment as one LOCATION adjustment. This table is a reproduction of the table in Android API guides [8].

choosing AdjustSettings in response to a nudge, we deem these setting adjustments to be directly motivated by the nudge. We chose a relatively short time frame for setting adjustments (i.e. 20 minutes) to exclude any app setting adjustments that were not directly motivated by interacting with the nudge.

		Repeated Nudge		
		Adjusted App Settings	Did Not Adjust App Settings	
First Nudge	Adjusted App Settings	9 (7.8%)	10 (8.7%)	19 (16.5%)
	Did Not Adjust App Settings	13 (11.3%)	83 (72.2%)	96 (83.5%)
		22 (19.1%)	93 (80.9%)	

Table 5.6: Number of participants who adjusted (or did not adjust) their app settings (AAS) in response to the first nudge in comparison to the repeated nudge. The difference is not statistically significant between the number of participants who AAS in response to the first nudge in comparison to the repeated nudge.

As Table 5.6 shows, 32 (27.8%) participants adjusted their app settings in response to both or one of the two nudges that they received during the experiment. Again, we hypothesized that participants engagement with nudges will decrease when they receive a repeated nudge. The number of participants who adjusted their app settings in response to the first nudge is 19 (16.5%), whereas 22 (19.1%) participants adjusted their app settings in response to the repeated nudge. The difference is not statistically significant (McNemar's $\tilde{\chi}^2$ = 0.17, df = 1, p-value = 0.68). The results of app settings adjustments suggest that user engagement with repeated nudges did not decrease when receiving a repeated nudge and, thus, we reject our hypothesis.

		Repeated Nudge		
		Adjusted App Settings	Did Not Adjust App Settings	
First Nudge	Adjusted App Settings	5 (8.8%)	5 (8.8%)	10 (17.5%)
	Did Not Adjust App Settings	5 (8.8%)	42 (73.6%)	47 (82.5%)
		10 (17.5%)	47 (82.5%)	

Table 5.7: Number of participants in condition 1 (updating the nudge's content) who adjusted (or did not adjust) their app settings (AAS) in response to the first nudge in comparison to the repeated nudge. The difference is not statistically between the number of participants in condition 1 who AAS in response to the first nudge in comparison to the repeated nudge.

We further explored how participants in condition 1 (updating content) and condition 2 (varying content) adjusted their app settings in response to the first and repeated nudge. As Table 5.7 shows, the number of participants in condition 1 who adjusted their app settings in response to both or either nudges is 15 (26.3%) participants. The number of participants in condition 1 who adjusted their app settings in response to the first nudge is 10 (17.5%), whereas the number of participants in condition 1 who did so in response to the repeated nudge is also 10 (17.5%). However, the difference is not statistically significant (McNemar's $\tilde{\chi}^2$ = 0, df = 1, p-value = 1). Similarly, the number of participants in condition 2 (varying content) who adjusted their app settings in response to both or either nudges is 17 (29.3%) as Table 5.8 shows. The number of participants in condition 2 who adjusted their app settings in response to the first nudge is 9 (15.5%), whereas the number of participants in condition 2 who did so in response to the repeated nudge is 12 (20.7%). However, the difference is not statistically significant (McNemar's $\tilde{\chi}^2$ = 0.31, df = 1, p-value = 0.57). Our results of app settings adjustments per condition suggest that user engagement with repeated nudges did not decrease when receiving a repeated nudge and, thus, we reject our hypothesis. Alternatively, it is possible that the absence of significant difference (if it indeed exists) within each condition is due to the lack of statistical power as the number of participants, who adjusted their app settings in response to the nudges, is relatively small.

Furthermore, we compared how participants in condition 1 (updating content) and participants in condition 2 (varying content) adjusted their app settings in response to the repeated nudge. The number of participants in condition 1 who adjusted their app settings in response to the repeated nudge is 10 (17.5%) participants, whereas the number of participants in condition 2 who did so is 12 (20.7%). However, the difference between the two conditions is not statistically significant ($\tilde{\chi}^2$ = 0.04, df = 1, p-value = 0.85). These results suggest that varying nudge contents did not affect how participants adjusted their app settings in response to repeated nudges. Therefore, we reject our hypothesis. Again, it is possible that the absence of significant difference (if it indeed exists) between the two conditions is due to the lack of statistical power as the number of participants, who adjusted their app settings in response to the nudges, is relatively small.

		Repeated Nudge		
		Adjusted App Settings	Did Not Adjust App Settings	
First Nudge	Adjusted App Settings	4 (6.9%)	5 (8.6%)	9 (15.5%)
	Did Not Adjust App Settings	8 (13.8%)	41 (70.7%)	49 (84.5%)
		12 (20.7%)	46 (79.3%)	

Table 5.8: Number of participants in condition 2 (varying the nudge's content) who adjusted (or did not adjust) their app settings (AAS) in response to the first nudge in comparison to the repeated nudge. The difference is not statistically between the number of participants in condition 2 who AAS in response to the first nudge in comparison to the repeated nudge.

Of the 70 participants who chose AdjustSettings in response to either nudges, 32 (45.7%) participants followed through and made 448 app setting adjustments as Table 5.9 shows. The most frequently adjusted permission group is LOCATION with 176 adjustments (39.3%) made by 31 participants. We expected LOCATION to be the most frequently adjusted permission group given that in our experiment we focused on nudging users about how apps on their phones access their location information. However, participants also adjusted app settings for other types of sensitive information (other than location information). As Table 5.9 shows, participants made adjustments to the other eight permission groups to prevent apps from accessing sensitive information or sensors such as phone contacts' list, camera, microphone, and external storage. We refer to this phenomenon as the "spill-over effect" of privacy nudges. In the exist survey, participants drew attention to this phenomenon. For example, a participant explained: "After receiving your first "nudge" in report 1, I had already done a full permissions audit for all my apps wherein I looked not just at location, but other data that my apps were accessing." Another participant provided a specific example of the spill-over effect: "I was interested in restricting what apps recieve my location and ended up turning off "microphone" on a number of apps." We have seen the same spill-over effect in Chapter 3 and Chapter 4.

Permission Group	1st Nudge (%)	# of Participants	Repeated Nudge	# of Participants
LOCATION	128 (42.7%)	18	48 (32.4%)	19
CONTACTS	44 (14.7%)	6	25 (16.9%)	4
STORAGE	62 (20.7%)	3	42 (28.4%)	6
PHONE	22 (7.3%)	4	4 (2.7%)	3
CALENDAR	9 (3%)	2	12 (8.1%)	5
CAMERA	6 (2%)	4	10 (6.8%)	3
SMS	22 (7.3%)	3	3 (2%)	3
MICROPHONE	5 (1.7%)	3	4 (2.7%)	2
SENSORS	2 (0.6%)	2	0 (0%)	0

Table 5.9: Number of app setting adjustments per permission group in response to the first and repeated nudge and the number of unique participants who made these adjustments. The permission groups are ordered by the total number of unique users who made adjustments and then by the total number of adjustments per a permission group. The total number of unique participants who made adjustments per a permission group is simply the sum of unique participants per nudge except for LOCATION. The total number of unique participants who made LOCATION adjustments is 31 participants (instead of 37) because 6 participants made LOCATION adjustments in response to both the first and repeated nudge. Overall, participants made more app setting adjustments in response to the first nudge than the repeated nudge. In addition, participants made more adjustments per permission groups in response to the first nudge than the repeated nudge except for CALENDAR and CAMERA.

Overall, participants made more app setting adjustments in response to the first nudge than the repeated nudge: 300 (67%) adjustments in response to the first nudge, and 148 (33%) adjustments in response to the repeated nudge. The difference is statistically significant (Exact binomial test p-value < 0.001). As Table 5.9 shows, making more adjustments in response to the first nudge than repeated nudge also holds when we break adjustments down by permission groups with two exceptions: Calendar and CAMERA. Participants made more adjustments to these two permission groups in response to the repeated nudge. As mentioned earlier, 32 (27.6%) participants adjusted their app settings in response to both or one of the two nudges that they received during the experiment. However, the difference between the number of app setting adjustments per participant in response to the first nudge and repeated nudge is not statistically significant (Wilcoxon signed rank test V = 276, p-value = 0.19).

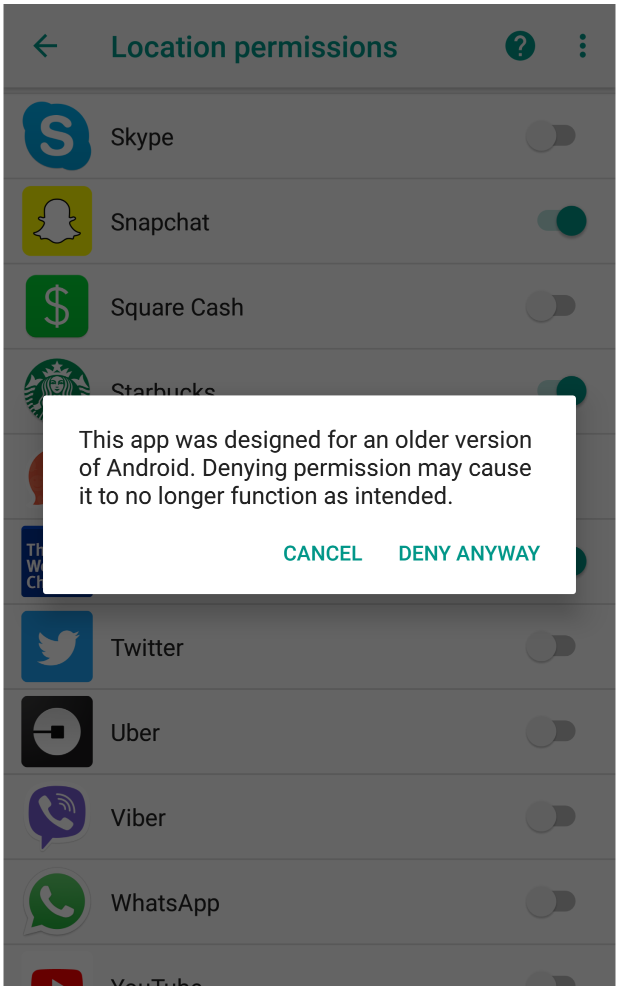

This app was designed for an older version of Android. Denying permission may cause it to no longer function as intended.

Figure 5.5: Android warns users that apps may not work properly if denied access to sensitive permissions if the app was designed to target SDK versions prior to 23 (Android OS version 6). This warning is shown for compatibility reasons because apps that target SDKS 22 or less were designed before Android introduced the current run-time permission model and, thus, these apps may crash if denied access to certain permissions. However, such a warning may influence users' decisions to allow or deny apps' access to their sensitive information.

As reported earlier, not all participants, who chose AdjustSettings in response to a nudge, followed through and made app setting adjustments. Of the 70 participants who chose AdjustSettings in response to nudges, only 32 (45.7%) participants followed through and made app settings adjustments. Participants' responses to the exist survey provide additional insights regarding why this happened.

Some participants used the opportunity of being redirected to the permission manager (after they chose AdjustSettings in response to a nudge) to double check that their settings aligned with their preferences without actually making any adjustments. For instance, a participant explained: "After getting report 2, I simply wanted to do (and did) a quick check to make sure that I hadn't missed anything in my first pass."

Some participants may have wanted to adjust app settings but decided not to because they were afraid to affect apps' functionality. For instance a participant explained: "I was surprised both times to see that the number of times my location was accessed was still so high. I was curious to see which apps were still accessing my location so I clicked on the first option again [Adjust-Settings]. However, once I saw those apps, I did not change any of the settings because I wasn't sure how much of a difference it would really make. So many of these apps claim to only work if there's location access." It is relevant to point out that Android (for compatibility reasons[4]) warns users that apps may not work properly if denied access to sensitive permissions if the app was designed to target SDK versions prior to 23 (Android OS version 6) as Figure 5.5 shows. Such warnings may have nudged participants in the other direction (i.e. not to adjust their app settings).

Other participants had a short attention span due to other more pressing priorities. A participant explained: "I got the message, I wanted to restrict, but I was in the middle of something else. When it prompted me to select the apps, I did not have time to go thru each one, so I stopped there." This short attention span for mobile users has be documented by earlier studies (e.g., [157]).

Some participants also used the nudge as a tool to discover permission management mechanisms in Android. A participant explained: "It was a quick short cut to the permissions page, though I didn't change anything." Another participant elaborated: "I just wanted the knowledge about how to restrict apps, not necessarily to actually do so."

[4] Apps that target SDKS 22 or less were designed before Android introduced the current run-time permission model. Therefore, these apps may not have taken into account the possibility of not having access to certain permissions and, thus, the apps crash if denied access to these permissions.

Length of Interaction Sessions with Nudges

We also analyzed the length of time participants interacted with the first and repeated nudges (interaction session). Specifically, the length of an interaction session is the difference between the time of presenting the nudge to a participant and the time in which the participant chooses an option (e.g., AdjustSettings). We removed 15 outliers[5] because they may not reflect the actual length of the interaction session (e.g., a participant may have left the nudge on the screen for a relatively long time without real interaction). The remaining 100 participants spent on average 16.89 seconds interacting with the first nudge (SD = 10.95, median = 15.31) and 14.76 seconds interacting with the repeated nudge (SD = 10.75, median = 12.77). However, this difference is not statistically significant (t = 1.57, p-value = 0.12).

We further explored whether the length of the interaction sessions differ within each condition. On the one hand, participants in condition 1 (updating content) spent on average 17.02 seconds (SD = 11.25, median = 15.1) interacting with the first nudge and 13.16 seconds (SD = 9.84, median = 10.98) interacting with the repeated nudge. The difference is marginally significant (t = 2, p-value = 0.05). On the other hand, participants in condition 2 (varying content) spent on average 16.76 seconds (SD = 10.75, median = 15.66) interacting with the first nudge and 16.37 seconds (SD = 11.46, median = 14.79) interacting with the repeated nudge. The difference is not statistically significant (t = 0.21, p-value = 0.84). These results may suggest that participants in condition 1 (updating content) were familiar with the content of the repeated nudge as it shows again the frequency of access by apps (albeit updated) and thus, participants spent shorter time interacting with the repeated nudge. However, participants in condition 2 (varying content) spent on average the same amount of time interacting with the repeated nudge as they did with the first nudge because they had not seen the content of the repeated nudge before. However, although participants in condition 2 (varying content) spent longer time interacting with the repeated nudge than participants in condition 1 (updating content), the difference is not statistically significant (F = 2.26, p-value = 0.14).

Why Participants Engaged with Repeated Nudges in Certain Ways

Our objective is to understand why participants responded to the repeated nudge in a similar or different way to/than the first nudge. To this end, the exist survey asked participants two consecutive open-ended questions about why they responded in certain ways to the first nudge and repeated nudge, respectively. For example, if the participant responded to the first nudge by

[5] Outliers are participants whose length of interaction sessions with the first or repeated nudge is greater than third quartile (Q3) plus 1.5 times interquartile range (IQR).

choosing AdjustSettings, we presented this response and asked the participant why she responded this way. We asked a similar question about the repeated nudge. Then, we asked a third open-ended question to understand why participants responded similarly or differently to both nudges. For example, if a participant responded to the first nudge by choosing AdjustSettings and to the repeated nudge by choosing KeepSettings, we presented both responses and asked why the participant responded differently. We used open coding [178] to analyze participants' responses to these open-ended questions. We focused on the third question in which participants explained why they made the same or different decisions in response to both nudges. We used participants responses to the first two questions (that asked about the first and the repeated nudge individually) to support and clarify our coding. Two researchers went iteratively through the responses and labeled sentences and phrases with concepts. Concepts were iteratively improved (e.g., new concepts were added or similar concepts were combined) and a code book was created. The two researchers independently went through all the responses and coded them using the code book. Later, they reviewed the coding together and disagreements were reviewed and resolved (by adopting coding by both researchers, agreeing on one coding after discussion, or keeping the disagreement). Next, related concepts within the code book were combined into more generic themes (i.e. categories). After applying the generic themes to the coding, the two researchers sampled the coding and reviewed the generic themes to ensure consistency. It is important to point out that a participant's response can be labeled by more than one theme.

We organized the results into four groups based on how participants responded to both nudges: (1) choosing AdjustSettings in response to both the first and repeated nudges, (2) choosing AdjustSettings in response to the first nudge only, (3) choosing AdjustSettings in response to the repeated nudge only, or (4) choosing KeepSettings in response to both the first and repeated nudges.

(1) Choosing AdjustSettings in Response to Both Nudges

Thirty-six (31.3%) participants chose AdjustSettings in response to both the first and repeated nudges.

Participants chose AdjutsSettings in response to both nudges because they noticed apps that they wanted to restrict from accessing their location information or other types of sensitive information. A participant explained: "They both showed apps that were receiving my location that I felt did not need that information." Another participant referred to other types of sensitive information (in addition to location information): "To restrict apps from using my location or any other permissions I didn't think we're necessary." Other participants pointed out that the repeated nudge showed new apps or apps that they did not notice when participants received the first nudge. A participant elaborated: “Other apps were getting my location that I did not notice the first time and I wanted to restrict them too.”

Some participants started (or intended to start) managing their app settings after receiving the first nudge and continued doing so after receiving the repeated nudge. A participant explained: “I felt like both reports were telling me the same information, so I utilized it in the same way. Once I realized what the first report did, I was just waiting for the second report so that I could block more applications.” Another participant described similar reasoning: “I did not cancel the application permissions from the first report as I got distracted with something else. With the second report, I was able to view the applications with this permission and made changes to some applications.”

Other participants responded to the repeated nudge, as they did to the first nudge, by choosing AdjustSettings to double check that their app settings are still aligned with their preferences. A participant explained: “After receiving your first ”nudge” in report 1, I had already done a full permissions audit for all my apps wherein I looked not just at location, but other data that my apps were accessing. After getting report 2, I simply wanted to do (and did) a quick check to make sure that I hadn’t missed anything in my first pass.” Others perceived the repeated nudge as a helpful reminder as they had not taken actions in response to the first nudge. A participant described: “After receiving report 1, I found it too complex to restrict location data. I never took action. (...) I was reminded that I needed to work on this when I received the report again.”

Participants also chose AdjustSettings in response to both the first and repeated nudges because the content of the repeated nudge was still concerning them. On the one hand, a participant in condition 1 (updating the nudge’s content), who received a frequency nudge in both times, explained: “I was surprised both times to see that the number of times my location was accessed was still so high.” On the other hand, a participant in condition 2 (varying the nudge’s content), who received a frequency nudge and then an implications nudge, described: “Essentially – report 2; with its additional details (track where you live, how much you make etc etc) made apps that would have been dismissed as innocuous in report 1 seem more suspicious/harmful.”

Participants also cited generic privacy concerns as reasons to respond to both nudge by choosing AdjutsSettings. A participant explained: “I want to control my privacy.” Another participant elaborated: “I do not want my information available to whoever wishes to purchase it.”

(2) Choosing AdjustSettings in Response to the First Nudge Only

Twenty-four (20.9%) participants responded to the nudges they received during the study by choosing AdjustSettings in response to the first nudge and KeepSettings in response to the repeated nudge.

Participants chose this option because they found the first nudge to be sufficient. A participant explained: “I responded differently because I had already taken action with the first report and did not feel the need to do so for the second report.” In particular, participants found the first nudge sufficient because the repeated nudge did not provide new information. A participant described: “In the first report, there were one or more apps I saw that I didn’t want to be able to access my location data, so I went and revoked it. In the second, there were no additional apps I wanted to revoke access for.” In addition, participants took advantage of the first nudge to discover permission management capabilities in Android and, thus, were less interested in repeated nudges. A participant explained: “Initially I only clicked restrict so that I could see how to restrict the behavior, in case I wanted to in the future.”

Participants also responded to the repeated nudge differently because app adjustments that they made in response to the first nudge may have caused their phones or apps to break. A participant explained: “When I decided to Restrict my apps in report 1, I was surprised at what my apps were doing. (...) I went on to remove several apps’s permissions, which made my phone act strangely for a day or 2. (...) When report 2 came about, I did not want to deal with the hassle again, and just allowed the apps to keep using my location.” Additionally, participants wanted to avoid the time-consuming task of carefully managing app settings. A participant elaborated: “But the second time I was afraid to disrupt the apps google maps or search apps. I wanted to look closely later at all the apps and then restrict, but I did not get time to do so.”

Some participants reconsidered their sharing preferences and, therefore, decided choosing KeepSettings in response to the repeated nudge. A participant described: “(...) However after some thought, I realized that even if I would be giving permission to share my location with all the apps again, there is not much one can do to prevent this from happening. Most, if not all, apps now track your location (if not more) and use that to market towards you. (...) I relegated myself to this unfortunate truth and allowed the apps to track my location.” Participants also reconsidered their earlier decisions due to the value of location-based services. A participant explained: “the first time I saw the message my instinct was to restrict access to location info as much as possible the second time I saw it, I thought there may be some utility in sharing some location info.” Other participants reconsidered their decisions because they wanted to explore available options for each nudge. A participant described: “I selected this because when this app randomly popped up on my screen, even though I was expecting it, I had no idea what the purpose was, so I wanted to explore the different options.”

(3) Choosing AdjustSettings in Response to the Repeated Nudge Only

Ten (8.7%) participants responded to the nudges they received during the study by choosing KeepSettings in response to the first nudge and AdjustSettings in response to the repeated nudge.

Participants chose a different decision because they received repeated nudges with different nudge contents. These are participants in condition 2 (varying content) in particular. A participant explained: "I responded to these 2 reports differently because they elicited different initial responses of emotion from me. The way that the content is phrased, the first report doesn't make me feel alarmed. However, the second report made me feel nervous and like my privacy was being invaded." Another participant pointed out, that the essential details the first nudge lack but the repeated nudge contained, caused different response: "Again, the 2nd report went more into detail regarding how location is used. I was initially apathetic in Report 1 because I believe that sharing location alone would not reveal much about my personal information. However, Report 2 revealed that location sharing is multi-faceted and thus more risky."

Participants also responded this way because the repeated nudge showed additional or new apps than the first nudge. A participant explained: "The second one had a different set of apps. The first one- most made sense why they'd need my location. But the second one there was definitely an app or two that really had no need to know my location."

Other participants reconsidered their sharing preferences when they received the repeated nudge. A participant explained: "Being asked the question twice made me wonder whether or not I should keep letting the apps see my location. I clicked restrict the second time mainly out of curiosity." Another participant reconsidered their sharing preferences because circumstances have changed: "i am testing out apps that claim to pay to know my location. since some apps did not pay me or were slow to pay, i limited how much information they got about me."

Participants also chose different decision because they wanted to discover permission management capabilities in Android. A participant explained: "I wanted to see what restriction power I had for the apps."

(4) Choosing KeepSettings in Response to Both Nudges

Forty-five (39.1%) participants chose KeepSettings in response to both the first and repeated nudge.

Participants chose KeepSettings in response to both nudges because they were not concerned about sharing their location information with apps. A participant explained: "I realize that some apps use location API's on my phone, and I generally don't have a problem sharing this data." Others did not particularly perceive potential privacy harms of sharing location information with apps. A participant described: "I did not wish to stop sharing my location with the apps, and I would not want to, unless I found out that the information was being used maliciously/with bad intent. Neither of the reports told me that, hence I responded the same way to both." Other participants believed privacy is already lost and thus felt no need to protect their information. A participant explained: "Because I really don't care if apps know where I am. Google already knows everything about me..why try to hide?"

Participants also chose KeepSettings in response to the first and repeated nudges because they previously allowed apps presented in the nudges to access participants' location information. A participant explained: "They [the nudges] both listed apps which I had previously given location access to." Others pointed out that the apps listed in the nudges match participants' expectations about which apps on their phones need/access location information. A participant described: "I responded this way because as I looked at the applications it all lined up with what I expected for location requests - nothing stood out or was an outlier that I noticed. I saw no need to restrict any one app's permissions based on the frequency of requests it made."

Some participants seem to have already adopted a policy in managing their app settings and, thus, they did not see value in the nudges and decided to choose KeepSettings in response to both nudges. A participant explained: "I try not to download apps that I think it might be dangerous to share information with in the first place. I take this into consideration when first downloading apps." Another participant elaborated: "I determine which apps I think need to know my location when I install them. The new Android system (I have Samsung S7, Android system 7.0) that allows me to deny certain permissions instead of giving an app blanket permission like in previous systems. I stuck with the decision I made when I installed the app. (...) I should note, though, that I turn my overall location information off the majority of the time. I only toggle my location on when I'm going to use an app that needs it (like Google Maps or Snapchat)."

Other participants thought that the repeated nudge was not convincing enough and, thus, responded to the repeated nudge as they did to the first nudge by choosing KeepSettings. A participant explained: "Even knowing the additional information the App Behavior Report gave me, I still valued the apps and the functions that using my location provided." Additionally, the repeated nudge did not convince some participants because the nudge did not provide new information from the first nudge. A participant explained: "The reports were almost identical, so I saw no reason in changing my response."

Participants also chose KeepSettings in response to both nudges because they like the benefits of sharing location information with apps. A participant explained: “Even knowing the additional information the App Behavior Report gave me, I still valued the apps and the functions that using my location provided.” Others pointed out that apps may not work properly without location information. A participant elaborated: “I answered the same way because my desire for the apps to work most efficiently did not change and for those apps to work best, they need to share my location data I believe.” Other participants were torn between functionality and privacy trade-offs. A participant described: “My desire to continue my usual phone usage (using the same apps, eg. Snapchat and Google Maps) outweighs my desire for privacy from those same apps.”

Some participants chose KeepSettings in response to the first and repeated nudge because they were afraid that adjusting app settings may cause apps or phones to stop working. A participant explained: “I wasnt sure how restricting the apps would change my ability to use them.” Another participant elaborated: “Most of the apps required knowing my location to function properly, such as maps, transit, and weather. The others I did not know what they did, and was wary of turning them off in case it would adversely affect the performance of my phone.”

Participants also chose KeepSettings in response to both nudges because app setting adjustment is a time-consuming task. A participant explained: “I didn’t want to restrict my location with the apps. I only use the location feature for some apps but I worried it would take a lot of time to go through and select specific apps to restrict (or even what consequences that would have on the phone function).” Others’ reasoning was generic: “Both times I didn’t want to change any setting on my phone.”

5.2.3 Factors that Influence Participants’ Responses to the Repeated Nudge

One of the main objectives of this experiment is to identify factors that may influence participants to respond to repeated nudges. To this end, we ran a logistic regression model which focuses on how participants responded to the repeated nudge. The model has one binary dependent variable that corresponds to the two options available to participants when interacting with the nudges: (1) RestrictSettings, and (2) KeepSettings. The logistic regression model has 17 independent variables as shown in Table 5.10. Two variables focus on the relations between the first and the repeated nudge. Specifically, the variables are (1) how the participant responded to the first nudge, (2) whether the participant adjusted her app settings in response to the first nudge and (3) whether the repeated nudge has additional apps than in the first nudge. Additionally, two other variables focus on how condition assignment affected participants’ decisions. In particular, these variables are (4) the condition that the participant was assigned to, and (5) whether the nudge’s content in the condition the participant was assigned to resonated with the participant (more than the nudge's content in the other condition) based on participants' responses in the exit survey. Four other variables focus on (6 & 8) awareness of and (7 & 9) concern about apps' practices presented in the nudges. Seven variables focus on demographics factors, namely, (10) gender, (11) age, (12-14) IUIPC privacy scale (control, awareness, collection), (15) education (a bachelor's degree and above vs. less than a bachelor's degree), and (16) occupation (student vs. non-student). The last independent variable focus on (17) the usage of permission tools in Android before participating in our experiment.

Coefficients	Estimate	Standard Error	z	p
Intercept	-4.2	3.2	1.31	0.19
Condition (base Condition 1)	-0.63	0.6	-1.1	0.29
Response to the first nudge (base KeepSettings)	1.44	0.65	2.22	0.027
Repeated nudge has more apps (base No)	1.71	0.64	2.67	0.008
Adjusted app settings in response to the 1st nudge (base No)	1.27	0.8	1.59	0.11
Awareness of frequency of accessing location information by apps	-0.08	0.19	-0.39	0.69
Concern about frequency of accessing location information by apps	0.37	0.3	1.2	0.22
Awareness of potential consequences of sharing location information with apps	-0.11	0.09	-1.27	0.2
Concern about potential consequences of sharing location information with apps	-0.16	0.16	-1	0.32
Your assigned nudge's content motivate you (base No)	1.83	0.68	2.68	0.007
IUIPC (Control)	-0.14	0.13	-1.11	0.27
IUIPC (Awareness)	-0.11	0.16	-0.71	0.48
IUIPC (Collection)	0.1	0.09	1.14	0.25
Gender (base Male)	0.35	0.53	0.66	0.51
Age	0.11	0.04	2.83	0.005
Education level (base Less than a bachelor's degree)	-1.28	0.63	-2	0.044
Occupation (base Non-student)	2.15	0.84	2.58	0.01
Usage of permission tools in Android	-0.01	0.33	-0.03	0.97

Table 5.10: Results of the logistic regression model. The results suggest that six factors may have influenced participants' decisions to respond to repeated nudges. The six factors are highlighted in green.

As Table 5.10 demonstrates, the logistic regression model shows statistically significant results for six factors: (1) how the participant responded to the first nudge, and (2) whether the repeated nudge has additional apps than in the first nudge, (4) whether the nudge's content in the condition the participant was assigned to resonated with the participant (more than the nudge's content in the other condition), (8) age, (11) occupation (student vs. non-student), and (15) education (a bachelor's degree and above vs. less than a bachelor's degree). Next, we elaborate on how these factors may have influenced participants' decisions.

Response to the First Nudge

Participants' responses to the repeated nudge are correlated with how participants responded to the first nudge as Table 5.10 suggests. Participants who indicated their IRAS in response to the first nudge were more likely to stick with the same decision when they received the repeated nudge. Similarly, participants who indicated that they wanted to keep app settings unaltered in response to the first nudge were more likely to stick with the same decision when they received the repeated nudge. Table 5.2 shows that 81 (70.4%) participants decided to stick with the same decision when they received the repeated nudge.

In the exist survey, participants highlighted their inclination toward sticking with the same decision (i.e. their response to the first nudge). On the one hand, a participant who chose Adjust-Settings in response to the repeated nudge explained: "Regardless, in both cases I did not want so many apps to view my location as it was a negative towards battery life and privacy." On the other hand, a participant who chose KeepSettings in response to the repeated nudge described: "I didn't see any difference between report 1 and report 2. I hadn't changed my mind about how I use the apps, so I made the same decision as previous." Participants responded to both nudges in the same way because their feeling toward the nudges stayed the same.

Repeated Nudges with Additional Apps

Each nudge states the total number of apps that have accessed the participant's location information and also presents a list of these individual apps (their names and icons). Thirty eight (33%) participants received a repeated nudge with more apps than the first nudge, whereas 32 (28%) and 45 (39%) received repeated nudges with same or fewer number of apps than the first nudge, respectively. We hypothesized that if the total number of apps in the repeated nudge exceeds the number of apps in the first nudge, then the participant might be motivated to interact with the repeated nudge because it has new information. We coded the existence of additional apps in the repeated nudge as a binary code. As shown in Table 5.10, the results of the logistic regression model suggest that participants are more likely to indicate their IRAS in response to the repeated nudge if it contains additional apps than the first nudge.

Participants' responses to the exit survey indicated similar reasoning. A participant who chose AdjustSettings in response to the repeated nudge explained: "I noticed a specific app with dubious requirements for location access the second time, and so I blocked it. The first time I had not noticed this app using my location." Similarly, a participant who chose KeepSettings in response to the repeated nudge (but chose AdjustSettings in response to the first nudge) pointed out: "In the first report, there were one or more apps I saw that I didn't want to be able to access my location data, so I went and revoked it. In the second, there were no additional apps I wanted to revoke access for."

Nudge Contents that Motivate Participants

	Condition 1: Updating the nudge's content (%)	Condition 2: Varying the nudge's content (%)
Assigned nudge's content	8 (14%)	34 (58.6%)
Another nudge's content	33 (57.9%)	11 (19%)
Both nudge contents	9 (15.8%)	8 (13.8%)
Neither nudge contents	7 (12.3%)	5 (8.6%)

Table 5.11: Participants self-reported whether they would be motivated by the nudge's content that they were assigned to, another nudge's content, both, or neither. Participants in condition 1 (updating the nudge's content) would have been motivated more to choose AdjustSettings (in response to repeated nudges) had they been assigned to condition 2 (varying the nudge's content).

In the exist survey, we presented the two nudge contents as shown in Figure 5.1 and asked participants to identify which of the two nudges would motivate them to choose AdjustSettings. Participants were given four options: the first nudge's content (i.e. the nudge content that shows frequency of access by apps), the second nudge's content (i.e. the nudge content that shows potential implications of sharing location information with apps), both nudge contents, or neither. We coded participants' responses to identify whether the nudge content that they received in the repeated nudge indeed motivated them. That is, if the participant indicated that she is motivated by the nudge that she received or both nudge contents, then we coded the responses as a match. Otherwise, the response is coded as mismatch. In other words, we coded whether the participants would have been better off[6] receiving a different nudge's content than the one the participant actually received. Then, we plugged the coding into the logistic regression model to check whether

[6] A participant is better off by choosing an option that matches her preferences.

participants' interaction with the repeated nudge would be influenced by a match/mismatch between a nudge's content and the participants' being motivated by the nudge's content that she received.

Fifty-nine (51.3%) participants self-reported that they would be motivated by the repeated nudge's content that they received, whereas 56 (48.7%) would be motivated either by another nudge's content or by neither one of the nudge contents in our experiment. Table 5.11 shows participants' responses grouped by conditions. It is notable that only 17 (29%) of participants in condition 1 (updating content) indicated that they would be motivated by the nudge's content they were assigned to, whereas 42 (72%) of participants in condition 2 (varying content) indicated they would be motivated by their assigned nudge's content. The difference between the two conditions is statistically significant ($\tilde{\chi}^2$ = 27.48, df = 3, p-value < 0.001) and pairwise comparison with FDR correction [75] shows that "Assigned nudge's content" and "Another nudge's content" differ significantly between the two conditions (adjusted p-value < 0.001). The results suggest a potentially missed opportunity for some participants, primarily in condition 1. These participants might have been motivated to review and possibly adjust their app settings had they been assigned to a different nudge's content that resonated more with them.

Demographic Factors

Our analysis suggest that two demographic factors correlate with participants' likelihood to choose AdjustSettings: Age and occupation.

The average age for participants who chose AdjustSettings is 31.24 (SD = 15.29), whereas the average age for participants who chose KeepSettings is 26.26 (SD = 9.83). The results of the logistic regression suggest that the likelihood of choosing AdjustSettings (in response to the repeated nudge) increases with the age of participants.

	AdjustSettings	KeepSettings
Less than a bachelor's degree	27 (45%)	33 (55%)
A bachelor's degree and above	19 (34.5%)	36 (65.5%)

Table 5.12: How participants who have bachelor's degrees and above vs. participants who have less than bachelor's degrees responded to repeated nudges. Participants who have higher education are more likely to choose KeepSettings.

We coded education as a binary variable: a bachelor's degree and above vs. less than a bachelor's degree. Table 5.12 shows how participants who have bachelor's degrees and above vs. participants who have less than bachelor's degrees responded to repeated nudges. Twenty-seven

(45%) of participants who have less than a bachelor's degree responded to the repeated nudges by choosing AdjustSettings, whereas 19 (34.5%) of participants who have a bachelor's degree or more responded to the repeated nudge by choosing AdjustSettings. As shown in Table 5.10, the results of the logistic regression model suggest that participants who have higher education level are more likely to choose KeepSettings.

	AdjustSettings	KeepSettings
Student	30 (42.3%)	41 (57.7%)
Non-student	16 (36.4%)	28 (63.6%)

Table 5.13: How students and non-student participants responded to the repeated nudge. Students are more likely to choose AdjustSettings in response to the repeated nudge.

We coded occupation as a binary variable: student vs non-student which includes employed, self-employed, retired, and unemployed. Table 5.13 shows how student and non-student participants responded to the repeated nudge. Thirty (42.3%) student participants chose AdjustSettings in response to the repeated nudge, whereas 16 (36.4%) non-student participants made the same decision. As shown in Table 5.10, the logistic regression model suggests a correlation between being a student and the likelihood of choosing AdjustSettings in response to the repeated nudge.

5.2.4 Repeated Nudges & Participants' Attitudes

We asked participants to indicate, through a 5-point likert scale question, whether the repeated nudge was a helpful reminder or was bothersome. As Figure 5.6 shows, 74 (64.3%) participants self-reported that the repeated nudge was a helpful reminder. Similarly, 76 (66.1%) participants self-reported that the repeated nudge was not bothersome. In other words, the majority of participants had a positive attitude toward the repeated nudge.

We also asked participants to indicate, through a 5-point likert scale question, whether the repeated nudge was necessary or unnecessary. Participants' self-reported answers show diverse views as depicted in Figure 5.7. Fifty-one (44.3%) participants considered the repeated nudge necessary, whereas 35 (30.4%) indicated that the repeated nudge was unnecessary. Although more participants indicated that the repeated nudge was necessary, a sizable number of participants considered the repeated nudge unnecessary.

Additionally, we asked participants about their attitudes toward the number of nudges they received during the study (i.e. two nudges). As Figure 5.8 shows, the majority of participants (68.7%) indicated that receiving two nudges within a week was reasonable.

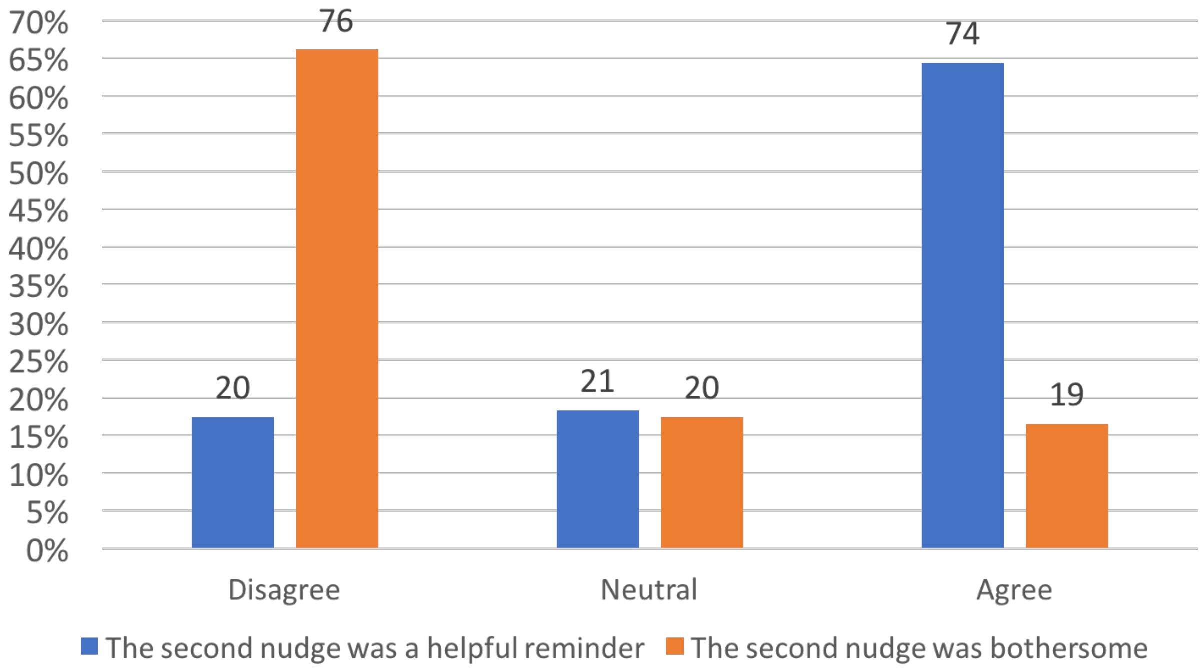

Figure 5.6: The majority of participants had a positive attitude toward the repeated nudge.

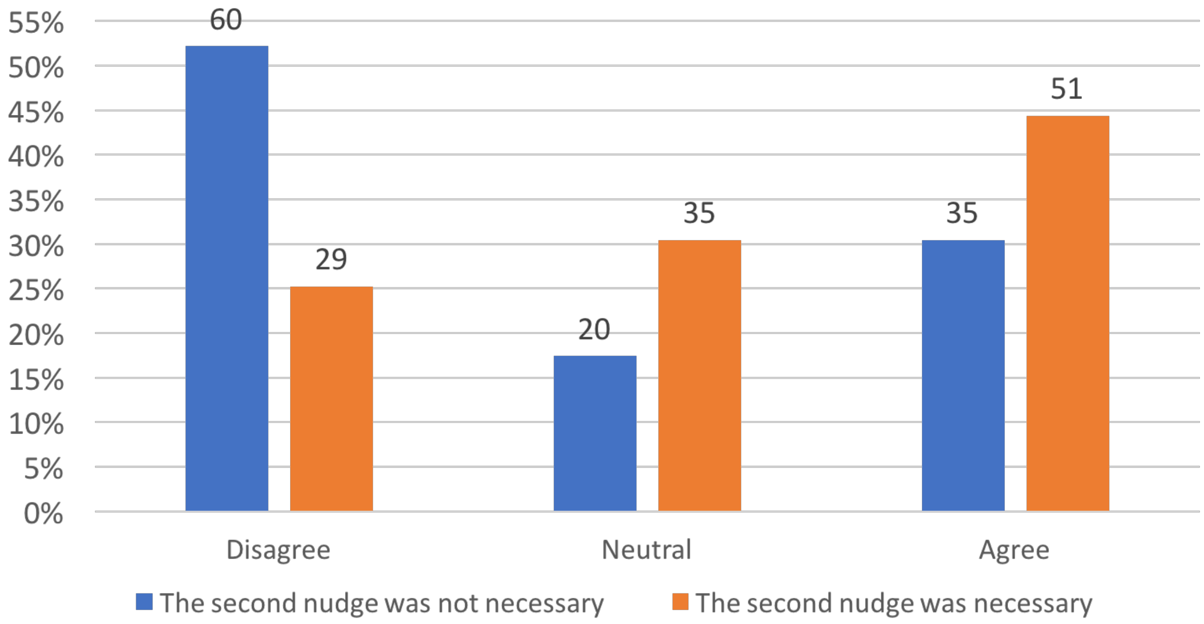

Figure 5.7: Participants showed diverse views regarding whether the repeated nudge was necessary or unnecessary. Although more participants indicated that the repeated nudge was necessary, a sizable number of participants considered the repeated nudge unnecessary.

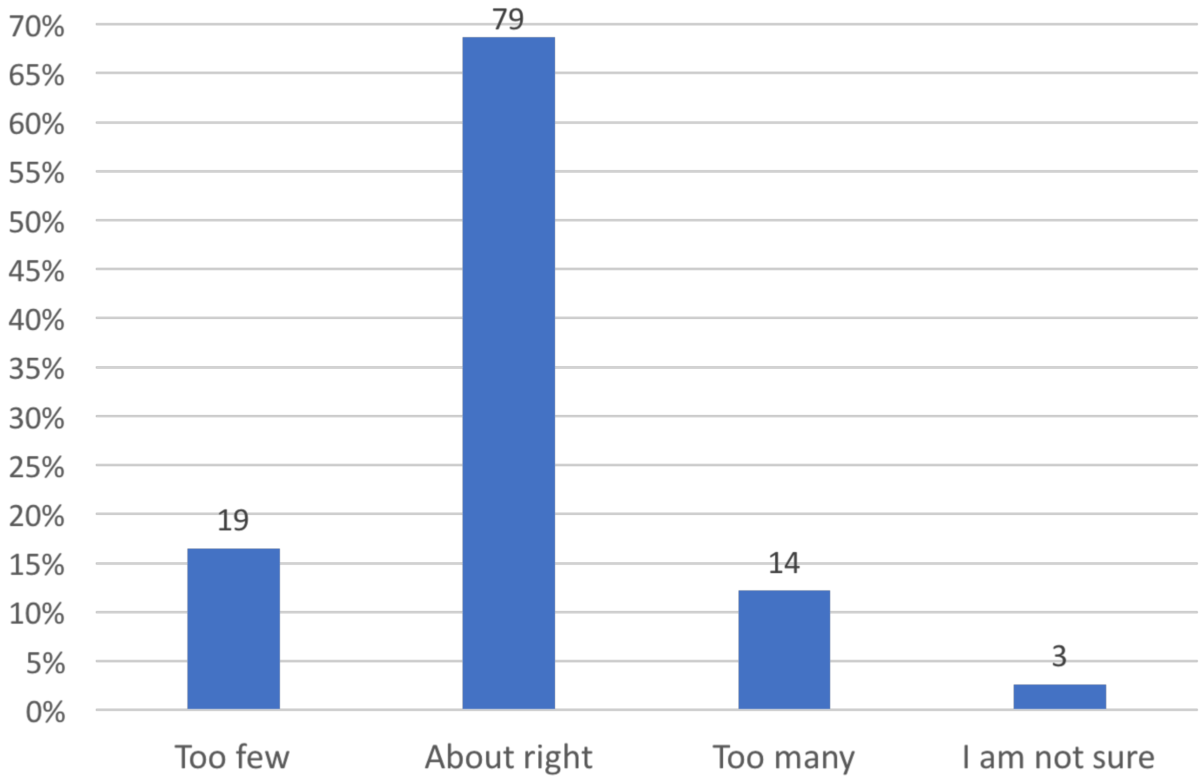

Figure 5.8: The majority of participants self-reported that receiving two nudges within a week was not too many.

5.3 Discussion

5.3.1 When Should We Send Users Repeated Nudges?

Our results suggest that user engagement with app privacy nudges decreases when users receive repeated nudges. More participants chose AdjustSettings in response to the first nudge than the repeated nudge. In addition, the number of app setting adjustments made in response to the repeated nudge was fewer than the adjustments made in response to the first nudge. These results are aligned with previous work on stimuli which suggest that user engagement with stimuli decreases with repeated exposure [64, 66, 67, 80, 82, 83, 84, 97, 179] and that such a decrease in engagement is observable from a second exposure [66, 83, 179, 183].

Nonetheless, our analysis suggests a number of factors that nudge designers may take into account when sending users repeated nudges. First, how users respond to the first nudge might be a good indicator of how they see a utility in the app privacy nudges. As our analysis showed, when receiving repeated nudges, participants tend to stick with their earlier decisions (i.e how they responded to the first nudge). Second, a user should receive a repeated nudge only when there is new information. In our experiment, participants were more likely to review and adjust their settings in response to a repeated nudge if it had additional apps than first nudge. Providing new information in repeated nudges can be taken a step further. Unlike our experiment in which the repeated nudge listed all apps including new ones, a nudge might be even more effective by only highlighting the new apps, thus making new information more salient. At the same time, the repeated nudge may give the user the ability to see the full list of apps (i.e. both new and old apps.). Future work may evaluate the effectiveness of this design decision. Third, our results suggest that users who receive nudge contents that resonate with them are more likely to engage with repeated nudges. Our results suggest that more participants in condition 1 (updating the nudge's content) might have been motivated to engage with the repeated nudge had they received the implication nudge content (i.e. the nudge's content that participants in condition 2 received). Such results may suggest a potential personalization opportunity/need. Prior work has shown that users' privacy preferences, attitudes, and behaviors are diverse (e.g., [59, 143, 169, 176, 196]). Future work may explore finding profiles of users that match profiles of nudge contents [96, 101, 143, 144, 145, 150, 168, 195]. In sum, these factors may increase the likelihood of engagement with nudges even when users receive them repeatedly.

5.3.2 Does Varying Nudge Contents Motivate Users to Engage with Repeated Nudges?

We hypothesized that varying nudge contents may influence participants' decisions to respond to repeated nudges. However, the logistic regression model shows no correlation between the condition a participant is assigned to and their likelihood of responding to the repeated nudge in a certain way. In addition, the difference between conditions is not statistically significant regarding how participants responded to the repeated nudge. Nonetheless, participants' responses in the exist survey show some influence of varying nudge contents. Simultaneously, some participants did not notice the difference between the nudge contents, most likely due to the nearly identical look and feel of the two nudges.

It could be the case that varying nudge contents does little to influence users' decisions when receiving repeated nudges as our analysis seem to suggest. However, the difference between nudge contents was not conspicuous enough at least for some participants. Thus, future work may explore the effect of varying nudge contents using two dissimilar nudge designs in contrast to the similar designs we used in our experiment.

5.3.3 Generalizability & Long-Term Effect of Repeated Nudges

When we analyzed participants' behavior within each of the two conditions in our experiment, we noticed signs of lacking statistical power. Lacking statistical power might be attributed to the relatively small number of participants per condition (i.e. 57 participants in condition 1 and 58 participants in condition 2) given the number of possible permutations of how participants respond to the first and repeated nudges (i.e. four possible responses: AdjustSettings vs. KeepSettings per nudge). As such, conducting a follow-up study with a larger number of participants may further validate/refine our findings.

In this experiment, we limited ourselves to two nudges, a first nudge and a repeated nudge, and to a relatively short duration (nudges were received three days apart within a six day experiment). However, a remaining open and interesting question to explore is how users engage with more than one repeated nudge over an extended period of time. The results of our experiment provide a basis for a future long-term experiment by narrowing the space and identifying some factors that affect how users engage with repeated nudges. For example, given that our experiment shows that new information (e.g., additional apps) increases the likelihood of engaging with repeated nudges, a future study might test how repeated nudges (more than one) with only new information affects users. In addition, a future experiment might focus on nudges that are received within a relatively longer period of time (e.g., a weekly nudge vs. monthly nudge) and how this influences users' decisions to interact with repeated nudges.

Chapter 6

Lessons Learned & Future Research Directions

The main outcomes of this dissertation can be summarized by the following thesis.

> *It is possible to design runtime mobile app privacy nudges that increase users' awareness of privacy risks, motivate them to review privacy settings, and cause them to adjust their app privacy settings to reflect their privacy preferences. Effective nudge contents highlight unexpected contexts, purposes, and potential implications of accessing sensitive information by apps. User engagement with repeated nudges can be maintained by limiting them to users who engaged with earlier nudges, to repeated nudges that contain new information, or to repeated nudge contents that resonate with users.*

In this concluding chapter, we summarize the main lessons from the three studies that form this dissertation. In addition, we highlight research directions that may extend the results of this dissertation within and beyond the area of mobile privacy research.

6.1 Lessons Learned

We have learned valuable lessons from designing, deploying, and evaluating app privacy nudges in three users studies. Here, we summarize these lessons.

6.1.1 Users Care About Privacy

The importance of privacy is sometimes downplayed [17, 28]. Additionally, it has been argued that users do not really care about privacy [37, 38] given the privacy paradox [154], which refers to the mismatch between users' stated privacy preferences and their actual behavior. However, our results suggest otherwise.

Throughout this dissertation, when users were given the opportunity to exercise control and protect their privacy, the majority did so. When Android users were given access to a permission manager (which was inaccessible by default) in the first study (see Chapter 3), they actively used the permission manager. Twenty-two out of 23 participants reviewed their app permissions at least once; half of them did so multiple times. Additionally, the permission manager led more than half of the study's participants to exercise control over the personal information apps were allowed to access, and participants modified permissions of both popular apps and pre-installed apps. In the second study (see Chapter 4), the majority of iOS and Android participants self-reported their willingness to review and adjust their app permissions in response to all app privacy nudges including the least effective nudges. In sum, our results suggest that the majority of users care about privacy. Our results are consistent with previous work which indicated that users care and value their privacy. In the context of mobile privacy, a recent Pew survey [135] showed that the majority (90%) of users want to know how much of their personal information is being collected by apps and how it may be used. In addition, 60% of mobile users refrained from installing an app due to privacy concerns [135].

6.1.2 Privacy Decision Making is Hard but Nudges can Help

Privacy decision making is difficult because "it is unrealistic to expect individual rationality" when users make decisions in the context of online privacy [60]. Privacy decision making is often subject to decision hurdles that may lead to suboptimal privacy decisions [58, 59]. Decision hurdles include incomplete information, bounded rationality, and cognitive and behavioral biases [58, 59, 60, 181]. Such difficulty in privacy decision making may provide an explanation for the privacy paradox [58, 60]. One approach that has been gaining traction to assist users when making privacy decisions is nudging [58, 181].

In this dissertation, we have designed mobile app privacy nudges that specifically address one hurdle in privacy decision making: incomplete information, which refers to the disparity between users' and service providers' knowledge of collection, use, sharing practices, potential consequences, and available protections concerning users' personal information [61]. We designed app privacy nudges to increase users' awareness of privacy risks associated with apps installed on their phones. As our studies have shown, users are mostly unaware of such risks and the nudges help bring these risks to users' attention.

6.1.3 Nudges Switch Users' Attention to Privacy Management

Privacy is typically a secondary task for users [79, 148, 198] which may explain (among other factors) why privacy enhancing tools (e.g., the permission manager) might be underutilized (see Chapter 3). Because privacy is a secondary task, most users are not motivated to switch their attention from their day-to-day life activities to privacy management. However, results from our experiments have shown an important benefit of privacy nudges: temporarily making privacy the primary task for users (see Chapter 3 & Chapter 4). Our results have shown that nudges can motivate users to stop their primary task and switch their focus to managing their app privacy settings. As a result, nudges increase the utilization of privacy enhancing tools (e.g., apps' permission manager) and help users make privacy decisions that are aligned with their privacy preferences.

As shown in Chapter 3, Chapter 4, and Chapter 5, when users received a nudge about a particular type of sensitive information (e.g., location information), they also adjusted settings for other types of sensitive information (e.g., camera or microphone). Furthermore, when users received a nudge about a list of apps, they reported intention to review which other apps were also accessing sensitive information. In other words, the nudges switched users' attention to privacy management which led users to review their app privacy settings beyond the limited content of the nudges. We refer to this phenomenon as the "spill-over effect" of privacy nudges.

6.1.4 Not All Nudge Contents May Resonate with Users

Both Android and iOS, the two most popular mobile operating systems, provide their users with some sort of app privacy nudge, which informs users that their location information has recently been accessed by apps (see Chapter 4, Figure 4.1). This "recent access" nudge in current mobile operating systems has been shown to be the least effective nudge content in Chapter 4. Worse, while we tested the "recent access" nudge content in an active setting in our experiment, the "recent access" nudge is passive in both Android and iOS. That is, users need to take a number of steps to find out the nudge and the information it provides. Previous work has shown that passive security and privacy warnings are less effective than their active counterparts [97, 197]. Thus, we expect the passive "recent access" nudge in current mobile operating systems to be even less effective in motivating users to review their app privacy settings than what our experiment showed.

Our experiment in Chapter 4 demonstrated that it is possible to design nudge contents that are more effective than the “recent access” nudge content which is used in current mobile operating systems. In particular, we showed that nudge contents which inform users about contexts (e.g., access users’ location information although users have not used the app), purposes (e.g., accessing users’ location information for advertising purposes), and potential implications of accessing their personal information (e.g., e.g., reusing users’ location information to infer additional unknown characteristics of users) are more effective than nudge contents which only inform users that their sensitive information has recently been accessed by apps on their phone.

6.1.5 User Engagement with Repeated Nudges

Prior work has shown that receiving a stimulus (e.g., a security warning) repeatedly decreases user engagement [64, 66, 67, 80, 82, 83, 84, 97, 179]. Results from Chapter 5 are aligned with previous work. User engagement with nudges decreases as users receive them repeatedly. Specifically, fewer users indicated their willingness to review their app settings in response to repeated nudges and a fewer number of app setting adjustments were made in response to repeated nudges. Nonetheless, Results from Chapter 5 identified factors that may motivate users to engage with repeated nudges. Users were more likely to engage with repeated nudges (1) if users have engaged with previous nudges, (2) if repeated nudges contain new information (e.g., additional apps, not shown in earlier nudges, that accessed sensitive resources), or (3) if the nudge contents of repeated nudges resonate with users.

6.1.6 Permission Managers Are Useful Tools

When we started working on this dissertation, the status of permission managers in mobile operating systems was drastically different than it is now. When we conducted our first study (see Chapter 3), only iOS provided its users with a permission manager to selectively grant/deny apps’ access to users’ personal information at run-time (i.e. after apps have been installed). In contrast, Android only supported an install-time approach, in which a user is presented with all the permissions requested by an app and the user needs to approve all requested permissions to install the app. Once installed, the user has very little control over how apps can access sensitive information.

When we made the permission manager available to Android users (which was inaccessible by default) in our first field study, participants actively reviewed and changed app permissions. All but one of our participants (22 out of 23 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.

We believe that our results from Chapter 3 and previous work by other researchers (e.g., [105, 133]), which have pointed out drawbacks of the install-time permission approach, as well as continuous demand from users and privacy advocates, may have prompted Google to bring permission management functionality to Android. As of Android Marshmallow (Android OS version 6), Android adopted a run-time approach to managing app permissions. Now, Android users have greater control over how individual apps access their sensitive information.

6.2 Future Research Directions

6.2.1 Longitudinal Evaluation of App Privacy Nudges

In Chapter 5, we designed and ran a controlled experiment in situ to evaluate how users respond to repeated nudges. However, we limited ourselves to two nudges, a first nudge and a repeated nudge, and to a relatively short duration (nudges were received three days apart within a six day experiment). A remaining open and interesting question to explore is how users engage with more than one repeated nudge over an extended period of time. In addition, future work may explore how users respond to repeated app privacy nudges that are only received sporadically (e.g., once a week or once a month). The results of our experiment in Chapter 5 may inform future studies in this direction. For example, given that our experiment shows that new information (e.g., additional apps or recently installed apps) increases the likelihood of engaging with repeated nudges, a future study might test how repeated nudges (more than one) with only new information may affect users.

6.2.2 Designing Effective Privacy Nudges

The nudges we designed and evaluated throughout this dissertation focused primarily on helping users overcome one decision making hurdle: incomplete information. Our nudges were designed to increase users' awareness of unknown or unexpected practices of mobile apps that may pose privacy risks. However, there are other hurdles that might affect privacy decision making. Thus, one future research direction is designing nudge contents that address one or more of these decision making hurdles. One example of a relevant hurdle to privacy decision making is hyperbolic discounting which refers to favoring an immediate benefit (e.g., using a game app) and discounting future cost or risk (e.g., future potential privacy consequences of sharing personal information with the app and advertising companies). Prior work has shown how hyperbolic discounting may affect users when making privacy decisions [59, 60, 140]. Designing nudges that address hyperbolic discounting may include bringing users' attention to privacy risks that are otherwise distant. Researchers can take a step further to evaluate how such a nudge may be effective for users who indeed exhibit hyperbolic discounting in comparison to users who are less prone to this particular psychological characteristic. Identifying such users might be achieved by using "hypothetical time discounting questions" [70].

In Chapter 4, we showed that different nudge contents matter and not all nudge contents may equally motivate users to review and possibly adjust their app privacy settings. However, we believe that our work is only a first step. Future work may include reevaluating the nudges in Chapter 4 but in situ (instead of a hypothetical setting). Furthermore, future work could explore the effectiveness of nudge contents when they convey real privacy risks in comparison to potential privacy risks. For example, it would be interesting to evaluate how users respond to nudges that highlight potential price or offer discrimination in comparison to nudges that actually show examples of these discrimination in which the user has been subjected to [147, 189] or that user might be subjected to (based on users' actual information, e.g., [118]). In addition, in this dissertation we only looked at nudge contents specific to location information. Future work may look at other types of information which users perceived as sensitive [104].

Another line of research is to evaluate the effectiveness of nudges that highlight privacy risks that pertain to a specific type of sensitive information (e.g., location information) in comparison to nudges that are app specific (e.g., how one app is accessing various types of sensitive information). Such a research question may identify scenarios and contexts in which one approach is more effective than the other. Recently, Norton Mobile Security app[1] has started using app-specific nudges as Figure 6.1 shows.

6.2.3 Informing the Design of Privacy Nudges in Emerging Domains

Although the primary focus of this thesis is designing and evaluating runtime mobile app privacy nudges, its results may have a broader impact on the design of privacy nudges for other emerging

[1] Norton Security and Antivirus `https://my.norton.com/mobile/home`

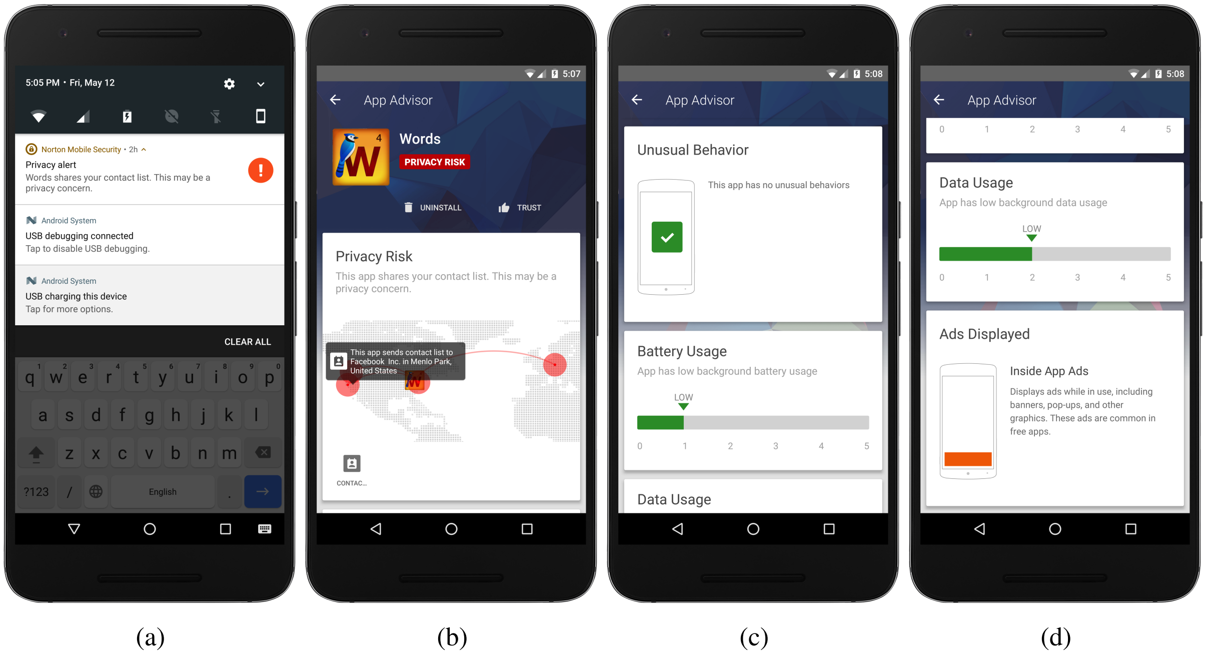

Norton Mobile Security • 2h

Privacy alert

Words shares your contact list. This may be a privacy concern.

Android System

USB debugging connected

Tap to disable USB debugging.

Android System

USB charging this device

Tap for more options.

Words

PRIVACY RISK

Privacy Risk

This app shares your contact list. This may be a privacy concern.

This app sends contact list to Facebook Inc. in Menlo Park, United States

Unusual Behavior

This app has no unusual behaviors

Battery Usage

App has low background battery usage

Data Usage

Data Usage

App has low background data usage

Ads Displayed

Inside App Ads

Displays ads while in use, including banners, pop-ups, and other graphics. These ads are common in free apps.

Figure 6.1: Norton Mobile Security app has recently introduced a new feature: "App Advisor." When the Norton app detects a privacy risk associated with another app installed on the user's phone, the app sends a privacy nudge in the form of a notification (a). When the user clicks on the nudge, the user is redirected to a detailed report showing the privacy risks highlighted in the nudge as well as other privacy risks associated with the app (b, c, and d).

platforms. For instance, it might be possible to extend the lessons learned from this research to an emerging research area of the Internet of Things (IoT).

The IoT refers to making physical objects network-enabled (e.g., smart watches, thermostats, or wristbands) so they can connect to the Internet and communicate with other entities (e.g., a server or a cellphone) [116]. Smartphones are expected to encourage the adoption and facilitate the success of the IoT [25, 42, 53, 54, 56, 193] because (a) network-enabled devices are typically configured and managed through accompanied mobile apps (e.g., Pebble watches[2], Fitbit wristbands[3] or Nest home products[4]) [193], and (b) smartphones will act as "proxies" to facilitate pulling information from the cloud that is relevant to nearby physical devices [193]. Although the IoT promises a wide range of innovative applications and products (e.g., health, home automation, or automobile) [23, 24], the IoT also brings privacy concerns and challenges [92, 158, 165, 193]. Given how smartphones are expected to play a central role in the IoT context, it seems reasonable to extend runtime mobile app privacy nudges to also include network-enabled devices. Such nudges will inform users of how network-enabled device access/collect their sensitive in-

[2] https://www.pebble.com/apps
[3] https://www.fitbit.com/app
[4] https://nest.com/blog/2015/06/17/one-home-one-app/

formation and inform users of privacy risks associated with such behavior. As we have shown in this thesis in the context of mobile apps, we would also expect privacy nudges to motivate users to review and possibly adjust privacy settings of their network-enabled devices. However, the scenarios, challenges, and privacy risks in the IoT context are different than the mobile apps context. Therefore, the content, form, timing, and frequency of privacy nudges in IoT context are compelling future research directions.

Bibliography

[1] App downloads up 15 percent in 2016, revenue up 40 percent thanks to China. https://techcrunch.com/2017/01/17/app-downloads-up-15-percent-in-2016-revenue-up-40-percent-thanks-to-china/,. Published: 2017-01-17, Accessed: 2017-07-09. 1

[2] Android Statistics - Number of Android applications. https://www.appbrain.com/stats/number-of-android-apps, . Updated: 2017-07-09, Accessed: 2017-07-09. 1

[3] Android - Dashboards. https://developer.android.com/about/dashboards/index.html,. Accessed: 2017-06-17. 5.1.8

[4] Android M Dev Preview delivers permission controls, fingerprint API, and more. http://goo.gl/OvJy6f,. Published: 2015-05-28, Accessed: 2015-06-07. 1

[5] Requesting Permissions at Run Time. https://developer.android.com/training/permissions/requesting.html, . Accessed: 2017-06-17. 2.1.1, 5.1.4

[6] System Permissions. https://developer.android.com/guide/topics/security/permissions.html,. Accessed: 2015-05-07. 2.1.1

[7] Android source code: RecentLocationApps.java. https://android.googlesource.com/platform/frameworks/base.git/+/android-cts-7.1_r5/packages/SettingsLib/src/com/android/settingslib/location/RecentLocationApps.java,. Accessed: 2017-05-14. 4.1.2

[8] System Permissions: Requesting Permissions. https://developer.android.com/guide/topics/permissions/requesting.html,. Accessed: 2017-05-07. (document), 5.5

[9] App Components - Services. https://developer.android.com/guide/components/services.html,. Accessed: 2017-05-14. 4.1.2

[10] Manifest.permission - SYSTEM_ALERT_WINDOW . https://developer.android.com/reference/android/Manifest.permission.html#SYSTEM_ALERT_WINDOW, . Accessed: 2017-06-17. 2

[11] UsageStatsManager . https://developer.android.com/reference/android/app/usage/UsageStatsManager.html, . Accessed: 2017-06-17. 1

[12] AppOpsManager.java. https://android.googlesource.com/platform/frameworks/base/+/android-7.0.0_r1/core/java/android/app/AppOpsManager.java. Accessed: 2017-05-14. 4.3.2, 13

[13] Optimizing for Doze and App Standby. https://developer.android.com/training/monitoring-device-state/doze-standby.html, . Accessed: 2017-06-17. 5.1.7

[14] Manage your battery - Keep battery optimization on . https://support.google.com/nexus/answer/7015477?hl=en, . Accessed: 2017-06-17. 5.1.7

[15] Complying with COPPA: Frequently Asked Questions. https://www.ftc.gov/tips-advice/business-center/guidance/complying-coppa-frequently-asked-questions. Published: 2015-03-20, Accessed: 2016-05-01. 4.1.2

[16] Awesome Privacy Tools in Android 4.3+. https://www.eff.org/deeplinks/2013/11/awesome-privacy-features-android-43. Published: 2013-12-11, Accessed: 2015-06-07. 2.1.1

[17] Google CEO Eric Schmidt Dismisses the Importance of Privacy. https://www.eff.org/deeplinks/2009/12/google-ceo-eric-schmidt-dismisses-privacy. Published:2009-12-10, Accessed: 2017-05-14. 6.1.1

[18] Privacy Online: A Report To Congress. Federal Trade Commission (FTC). https://www.ftc.gov/sites/default/files/documents/reports/privacy-online-report-congress/priv-23a.pdf. Published: 1998-06-01, Accessed: 2016-05-01. 4.1.2

[19] Android Flashlight App Developer Settles FTC Charges It Deceived Consumers. https://goo.gl/Zf18jI. Published: 2013-12-05, Accessed: 2015-06-07. 1

[20] Flurry Personas: Reach Your Target Audience with Flurry Personas. http://www.flurry.com/sites/default/files/resources/FlurryPersonasJune2014_3.pdf. Published: 2014-03-01, Accessed: 2016-03-09. 4.1.2

[21] Mobile Benchmarks - Q3 2015. https://www.adjust.com/assets/

downloads/mobile-benchmarks-q3-2015.pdf. Published: 2015-10-29, Accessed: 2016-02-15. 4.1.3

[22] Futuresight: User perspectives on mobile privacy. http://www.gsma.com/publicpolicy/wp-content/uploads/2012/03/futuresightuserperspectivesonuserprivacy.pdf. Published: 2011-11-01, Accessed: 2015-06-07. 1

[23] Gartner Says the Internet of Things Installed Base Will Grow to 26 Billion Units By 2020. http://www.gartner.com/newsroom/id/2636073, . Published:2013-12-12, Accessed: 2015-10-12. 6.2.3

[24] Gartner Says 6.4 Billion Connected "Things" Will Be in Use in 2016, Up 30 Percent From 2015. http://www.gartner.com/newsroom/id/3165317, . Published:2015-10-10, Accessed: 2015-10-12. 6.2.3

[25] Gartner Identifies the Top 10 Strategic Technology Trends for 2015. http://www.gartner.com/newsroom/id/2867917, . Published: 2014-10-08, Accessed: 2015-10-02. 6.2.3

[26] It Takes Google 'Now' Three Days To Figure Out Where You Live. https://www.forbes.com/sites/kashmirhill/2012/07/09/it-takes-google-now-three-days-to-figure-out-where-you-live/. Published: 2012-07-09, Accessed: 2016-05-01. 4.1.2

[27] Groupon - Shop Deals & Coupons. https://play.google.com/store/apps/details?id=com.groupon. Accessed: 2016-12-01. 4.1.3

[28] Privacy no longer a social norm, says Facebook founder. https://www.theguardian.com/technology/2010/jan/11/facebook-privacy. Published: 2010-01-10, Accessed: 2017-06-17. 6.1.1

[29] 2013 OECD Privacy Guidelines. http://www.oecd.org/internet/ieconomy/privacy-guidelines.htm. Published: 2013, Accessed: 2016-05-01. 4.1.2

[30] Pew Research Center: Public Perceptions of Privacy and Security in the Post-Snowden Era. http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/. Published: 2014-11-12. 1

[31] PrivacyGrade: Grading The Privacy Of Smartphone Apps. http://privacygrade.org/home. Accessed: 2015-2-17. 5.1.3

[32] Quiz: Can we guess your age and income, based solely on the apps on your phone? https://www.scientificamerican.com/article/tweets-

reveal-a-twitter-user-s-income-bracket/. Published: 2016-03-03, Accessed: 2016-05-01. 4.1.2

[33] An Update On Privacy At Uber. http://newsroom.uber.com/2015/05/an-update-on-privacy-at-uber/. Published: 2015-05-28, Accessed: 2015-06-07. 1

[34] On Orbitz, Mac Users Steered to Pricier Hotels. https://www.wsj.com/articles/SB10001424052702304458604577488822667325882, . Published: 2012-08-23, Accessed: 2016-12-01. 4.1.2

[35] Why You Can't Trust You're Getting the Best Deal Online. http://www.wsj.com/articles/why-you-cant-trust-youre-getting-the-best-deal-online-1414036862,. Published: 2014-10-23, Accessed: 2015-06-07. 4.1.2

[36] Quiz: Can we guess your age and income, based solely on the apps on your phone? https://www.washingtonpost.com/news/the-intersect/wp/2016/03/03/quiz-can-we-guess-your-age-and-income-based-solely-on-the-apps-on-your-phone/?utm_term=.ddbbd66aca81. Published: 2016-03-03, Accessed: 2016-05-01. 4.1.2

[37] Do we really care about our online privacy? https://thenextweb.com/insider/2011/09/13/do-we-really-care-about-our-online-privacy/,. Published: 2011-09-13, Accessed: 2017-06-17. 6.1.1

[38] Let's Face It, We Don't Really Care About Privacy. https://www.forbes.com/sites/gregsatell/2014/12/01/lets-face-it-we-dont-really-care-about-privacy/,. Published: 2014-12-01, Accessed: 2017-06-17. 6.1.1

[39] Weather - The Weather Channel. https://play.google.com/store/apps/details?id=com.weather.Weather. Accessed: 2016-12-01. 4.1.3

[40] Words With Friends Classic. https://play.google.com/store/apps/details?id=com.zynga.words,. Accessed: 2016-12-01. 4.1.3

[41] Words With Friends Classic. http://privacygrade.org/apps/com.zynga.words.html,. Accessed: 2017-02-10. 4.1.3

[42] Accenture Technology Vision 2015 - Digital Business Era: Stretch Your Boundaries. http://goo.gl/cZBdTV. Accessed: 2015-10-02. 6.2.3

[43] Android M Dev Preview delivers permission controls, fingerprint API, and more. http://goo.gl/NdmOx1. Published: 2015-05-28, Accessed: 2015-06-07. 2.1.1, 5.1.4

[44] comScore Reports January 2015 U.S. Smartphone Subscriber Market Share. http://

goo.gl/UUBPgb. Published: 2015-03-04, Accessed: 2015-03-23. 1

[45] iOS 11's blue bar will shame apps that overzealously access your location . https://techcrunch.com/2017/06/26/ios-11s-blue-bar-will-shame-apps-that-overzealously-access-your-location/, . Published: 2017-06-26, Accessed: 2017-07-03. (document), 4.3.2, 4.9

[46] iOS 4: Understanding Location Services. https://support.apple.com/en-lb/HT201674, . Accessed: 2015-06-07. 2.1.2

[47] About privacy and Location Services using iOS 8 on iPhone, iPad, and iPod touch. https://support.apple.com/en-us/HT203033, . Accessed: 2015-06-07. 12

[48] Apple denies Chinese report that location data are a security risk. http://on.ft.com/VXpZKR, . Published: 2014-6-12, Accessed: 2014-9-14. 1, 2.1.2, 3.4.2, 4.1.2, 4.3.2

[49] Apple's App Store hits 2M apps, 130B downloads, $50B paid to developers. https://techcrunch.com/2016/06/13/apples-app-store-hits-2m-apps-130b-downloads-50b-paid-to-developers/, . Published: 2016-06-13, Accessed: 2017-07-09. 1

[50] Apple's App Store just had the most successful month of sales ever. https://www.theverge.com/2017/1/5/14173328/apple-december-2016-app-store-record-phil-schiller, . Published: 2017-01-05, Accessed: 2017-07-09. 1

[51] Turn Location Services and GPS on or off on your iPhone, iPad, or iPod touch. https://support.apple.com/en-us/HT207092, . Accessed: 2016-05-31. 4.1.8

[52] Understanding privacy and Location Services on iPhone, iPad, and iPod touch with iOS 8. http://support.apple.com/en-us/HT203033, . Published: 2014-11-14, Accessed: 2014-12-08. (document), 1, 2.1.2, 4.1.2, 4.8, 4.3.2

[53] IoT And The Looming Mobile Tidal Wave . http://goo.gl/SBSPfp, . Published: 2015-04-28, Accessed: 2015-10-02. 6.2.3

[54] Skip the hub; your mobile device is your IoT gateway. http://goo.gl/uTdNzF, . Published: 2014-10-21, Accessed: 2015-10-02. 6.2.3

[55] Path official blog: We are sorry. http://blog.path.com/post/17274932484/we-are-sorry. Published: 2012-2-8, Accessed: 2014-8-4. 1, 2.3

[56] Mobile Phones Will Serve as Central Hub to "Internet of Things". http://goo.gl/ia9ylt. Published: 2011-02-16, Accessed: 2015-10-02. 6.2.3

[57] Jagdish Prasad Achara, Gergely Acs, and Claude Castelluccia. On the unicity of smartphone applications. In *Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society*, pages 27–36. ACM, 2015. 4.1.2

[58] A Acquisti. Nudging privacy: The behavioral economics of personal information. *IEEE Security & Privacy*, 7(6):82–85, 2009. ISSN 1540-7993. doi: 10.1109/MSP.2009.163. 1, 2.3, 2.3, 6.1.2

[59] A Acquisti and J Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, 2005. 1, 2.3, 5.3.1, 6.1.2, 6.2.2

[60] Alessandro Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29. ACM, 2004. 6.1.2, 6.2.2

[61] Alessandro Acquisti and Jens Grossklags. *What can behavioral economics teach us about privacy?* Taylor & Franics, 2007. 1, 2.3, 2.3, 6.1.2

[62] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proc. SOUPS*, 2013. 2.3

[63] Yuvraj Agarwal and Malcolm Hall. ProtectMyPrivacy: detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proc. MobiSys*, 2013. 2.2.2, 2.4.1, 2.5

[64] Devdatta Akhawe and Adrienne Porter Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Proc. Usenix Security*, pages 257–272, 2013. 2.4.2, 5.1.2, 5.3.1, 6.1.5

[65] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwa. Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proc. CHI*. ACM, 2015. 1, 2.3, 2.5

[66] Bonnie Brinton Anderson, Anthony Vance, C. Brock Kirwan, David Eargle, and Seth Howard. Users aren't (necessarily) lazy: Using neurois to explain habituation to security warnings. In *Proceeding of 35th International Conference on Information Systems (ICIS 2014)*. Association for Information Systems (AIS), 2014. 2.4.2, 5, 5.1.2, 5.3.1, 6.1.5

[67] Bonnie Brinton Anderson, C. Brock Kirwan, Jeffrey L. Jenkins, David Eargle, Seth Howard, and Anthony Vance. How polymorphic warnings reduce habituation in the brain—insights from an fmri study. In *Proc. CHI*. ACM, 2015. 5, 5.1.2, 5.3.1, 6.1.5

[68] Apple. App review. `https://developer.apple.com/app-store/review/`, . Published: 2014-08-18, Accessed: 2014-12-03. 2.1.2

[69] Apple. App store review guidelines. `https://developer.apple.com/app-store/review/guidelines`, . Published: 2014-08-18, Accessed: 2014-12-03. 2.1.2

[70] Nava Ashraf, Dean Karlan, and Wesley Yin. Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *The Quarterly Journal of Economics*, 121(2):635–672, 2006. 6.2.2

[71] Gökhan Bal, Kai Rannenberg, and Jason Hong. Styx: Design and evaluation of a new privacy risk communication method for smartphones. In *ICT Systems Security and Privacy Protection*, pages 113–126. Springer, 2014. 2.4.1, 2.5, 4.1.2

[72] Rebecca Balebako, Pedro G Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Nudging users towards privacy on mobile devices. In *Proc. CHI-PINC*, 2011. 2.3

[73] 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. 1, 2.3, 2.4.1, 2.5, 3.1.1, 3.1.2, 4.1.2

[74] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. The impact of timing on the salience of smartphone app privacy notices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, pages 63–74. ACM, 2015. 1, 2.5

[75] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300, 1995. 4.2.2, 5.2.3

[76] Alastair R Beresford, Andrew Rice, Nicholas Skehin, and Ripduman Sohan. Mockdroid: trading privacy for application functionality on smartphones. In *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications*, pages 49–54. ACM, 2011. 2.2.1

[77] Andrew Besmer, Jason Watson, and Heather Richter Lipford. The impact of social navigation on privacy policy configuration. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, page 7. ACM, 2010. 2.4.2

[78] Bin Bi, Milad Shokouhi, Michal Kosinski, and Thore Graepel. Inferring the demographics of search users: Social data meets search queries. In *Proceedings of the 22nd international conference on World Wide Web*, pages 131–140. ACM, 2013. 4.1.2

[79] Rainer Böhme and Jens Grossklags. The security cost of cheap user interaction. In *Proceedings of the 2011 workshop on New security paradigms workshop*, pages 67–82. ACM, 2011. 1, 4, 4.1.8, 6.1.3

[80] Rainer Böhme and Stefan Köpsell. Trained to accept?: a field experiment on consent dialogs. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2403–2406. ACM, 2010. 5.1.2, 5.3.1, 6.1.5

[81] Jan Lauren Boyles, Aaron Smith, and Mary Madden. Privacy and Data Management on Mobile Devices. `http://www.pewinternet.org/files/old-media//Files/Reports/2012/PIP_MobilePrivacyManagement.pdf`. Published:2012-09-05, Accessed: 2016-05-20. 4.1.3, 5.1.5

[82] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. Your attention please: designing security-decision uis to make genuine risks harder to ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, page 6. ACM, 2013. 2.4.2, 5.1.2, 5.3.1, 6.1.5

[83] Cristian Bravo-Lillo, Lorrie Cranor, Saranga Komanduri, Stuart Schechter, and Manya Sleeper. Harder to ignore? In *Symposium on Usable Privacy and Security (SOUPS'14)*. USENIX, 2014. 1, 2.4.2, 5, 5.1.2, 5.3.1, 6.1.5

[84] José Carlos Brustoloni and Ricardo Villamarín-Salomón. Improving security decisions with polymorphic and audited dialogs. In *Proceedings of the 3rd symposium on Usable privacy and security*, pages 76–85. ACM, 2007. 2.4.2, 5.1.2, 5.3.1, 6.1.5

[85] 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. 1, 2.2.1, 2.4.1, 2.5

[86] Sunny Consolvo, Ian E Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 81–90. ACM, 2005. 4.1.2, 4.1.3

[87] Vincent C Conzola and Michael S Wogalter. A communication–human information processing (c–hip) approach to warning effectiveness in the workplace. *Journal of Risk Research*, 4(4):309–322, 2001. 5.1.2

[88] Lorrie Faith Cranor, Joseph Reagle, and Mark S Ackerman. Beyond concern: Understanding net users' attitudes about online privacy. *The Internet upheaval: raising questions, seeking answers in communications policy*, pages 47–70, 2000. 4.1.2

[89] Aron Culotta and Jennifer Cutler. Predicting the demographics of twitter users from website traffic data. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. AAAI, 2015. 4.1.2

[90] Sauvik Das, Adam DI Kramer, Laura A Dabbish, and Jason I Hong. Increasing security sensitivity with social proof: A large-scale experimental confirmation. In *Proceedings of*

the 2014 ACM SIGSAC Conference on Computer and Communications Security, pages 739–749. ACM, 2014. 2.4.2

[91] A. Datta, M. Tschantz, and A. Datta. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *Privacy Enhancing Technologies*, 2015(1): 92–112, 2015. ISSN 2299-0984. doi: 10.1515/popets-2015-0007. 4.1.2

[92] Roberto Di Pietro and Luigi V Mancini. Security and privacy issues of handheld and wearable wireless devices. *Communications of the ACM*, 46(9):74–79, 2003. 6.2.3

[93] Yuxiao Dong, Yang Yang, Jie Tang, Yang Yang, and Nitesh V Chawla. Inferring user demographics and social strategies in mobile social networks. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 15–24. ACM, 2014. 4.1.2

[94] 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. 2.1.1, 3.4.1

[95] Manuel Egele, Christopher Kruegely, Engin Kirdaz, and Giovanni Vigna. PiOS: Detecting privacy leaks in iOS applications. In *Proc. NDSS*, 2011. 2.2.2

[96] Serge Egelman and Eyal Peer. Predicting privacy and security attitudes. *SIGCAS Comput. Soc.*, 45(1):22–28, February 2015. ISSN 0095-2737. doi: 10.1145/2738210.2738215. URL `http://doi.acm.org/10.1145/2738210.2738215`. 5.3.1

[97] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proc. CHI*, pages 1065–1074. ACM, 2008. 4.1.8, 4.3.1, 5.1.2, 5.3.1, 6.1.4, 6.1.5

[98] Serge Egelman, Andreas Sotirakopoulos, Ildar Muslukhov, Konstantin Beznosov, and Cormac Herley. Does my password go up to eleven?: the impact of password meters on password selection. In *Proc. CHI*, pages 2379–2388. ACM, 2013. 2.4.2

[99] William Enck, Peter Gilbert, Seungyeop Han, Vasant Tendulkar, Byung-Gon Chun, Landon P Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems (TOCS)*, 32(2):5, 2014. 1, 2.2.1, 2.3, 4.1.3, 5.1.3

[100] Facebook. Making it easier to share with who you want. `http://newsroom.fb.com/news/2014/05/making-it-easier-to-share-with-who-you-want/`. Published: 2014-05-22, Accessed: 2014-12-08. 1, 2.4.2

[101] Lujun Fang and Kristen LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World wide web*, pages 351–360. ACM,

2010. 5.3.1

[102] Adrienne Porter Felt, Erika Chin, Steve Hanna, Dawn Song, and David Wagner. Android permissions demystified. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 627–638. ACM, 2011. 4.1.3, 5.1.5

[103] Adrienne Porter Felt, Kate Greenwood, and David Wagner. The effectiveness of application permissions. In *Proceedings of the 2nd USENIX conference on Web application development*, pages 7–7. USENIX Association, 2011. 2.2.1

[104] 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. 2.3, 3.2.3, 4.1.3, 4.3.1, 5.1.5, 6.2.2

[105] 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. ISBN 978-1-4503-1532-6. doi: 10.1145/2335356.2335360. URL `http://doi.acm.org/10.1145/2335356.2335360`. 2.2.1, 2.3, 6.1.6

[106] Adrienne Porter Felt, Robert W Reeder, Hazim Almuhimedi, and Sunny Consolvo. Experimenting at scale with google chrome's ssl warning. In *Proc. CHI*, pages 2667–2670. ACM, 2014. 2.4.2

[107] Drew Fisher, Leah Dorner, and David Wagner. Short paper: location privacy: user behavior in the field. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices (SPSM)*, pages 51–56. ACM, 2012. 1, 2.2.2

[108] Alain Forget, Sonia Chiasson, Paul C van Oorschot, and Robert Biddle. Improving text passwords through persuasion. In *Proceedings of the 4th symposium on Usable privacy and security*, pages 1–12. ACM, 2008. 2.4.2

[109] The Federal Trade Commission (FTC). Data brokers: A call for transparency and accountability. Technical report, 2014. 4.1.2

[110] Huiqing Fu, Yulong Yang, Nileema Shingte, Janne Lindqvist, and Marco Gruteser. A field study of run-time location access disclosures on android smartphones. 2014. 1, 2.2.1, 2.3, 2.4.1, 2.5, 3.1.1, 4.1.2

[111] Jeremy Goecks, W Keith Edwards, and Elizabeth D Mynatt. Challenges in supporting end-user privacy and security management with social navigation. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, page 5. ACM, 2009. 2.4.2

[112] Sharad Goel, Jake M Hofman, and M Irmak Sirer. Who does what on the web: A large-scale study of browsing behavior. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*. AAAI, 2012. 4.1.2

[113] Google. Appopsmanager code in android 4.3. `https://android.googlesource.com/platform/frameworks/base/+/android-4.3_r2.1/core/java/android/app/AppOpsManager.java`, . Published: 2013-12-12, Accessed: 2014-12-03. 2.1.1, 3.2.1

[114] Google. Appopsmanager code in android 4.3. `https://android.googlesource.com/platform/frameworks/base/+/android-4.4_r1/core/java/android/app/AppOpsManager.java`, . Published: 2013-12-12, Accessed: 2014-12-03. 2.1.1

[115] Google. Appopsmanager code in android 4.3. `https://android.googlesource.com/platform/frameworks/base/+/android-5.0.0_r7/core/java/android/app/AppOpsManager.java`, . Published: 2014-08-18, Accessed: 2014-12-03. 2.1.1

[116] Jayavardhana Gubbi, Rajkumar Buyya, Slaven Marusic, and Marimuthu Palaniswami. Internet of things (iot): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7):1645–1660, 2013. 6.2.3

[117] Aniko Hannak, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 305–318. ACM, 2014. 4.1.2

[118] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proc. CHI*, 2014. ISBN 978-1-4503-2473-1. doi: 10.1145/2556288.2556978. URL `http://doi.acm.org/10.1145/2556288.2556978`. 1, 2.2.1, 2.4.1, 2.5, 4.1.2, 4.3.1, 6.2.2

[119] James J Heckman. Sample selection bias as a specification error (with an application to the estimation of labor supply functions), 1977. 4.1.8

[120] Baik Hoh, Marco Gruteser, Hui Xiong, and Ansaf Alrabady. Enhancing security and privacy in traffic-monitoring systems. *IEEE Pervasive Computing*, 5(4):38–46, 2006. 4.1.2

[121] 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. 2.2.1

[122] https://www.imore.com/ios-4-review . iOS 4 review. `https://www.imore.com/ios-4-review`. Published: 2010-6-14, Accessed: 2017-6-7. 4.2.2

[123] Lukasz Jedrzejczyk, Blaine A Price, Arosha K Bandara, and Bashar Nuseibeh. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, page 14. ACM, 2010. 2.4.1

[124] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proc. CHI*, 2004. 2.3

[125] 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. 2.2.1

[126] Maritza L Johnson. *Toward Usable Access Control for End-users: A Case Study of Facebook Privacy Settings*. PhD thesis, Columbia University, 2012. 2.4.2

[127] Jaeyeon Jung, Seungyeop Han, and David Wetherall. Enhancing mobile application permissions with runtime feedback and constraints. In *Proc. SPSM*, 2012. 3.1.1, 4.1.2

[128] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. Privacy attitudes of mechanical turk workers and the us public. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014. 4.2.1

[129] Amy K Karlson, AJ Brush, and Stuart Schechter. Can i borrow your phone?: understanding concerns when sharing mobile phones. In *Proc. CHI*, pages 1647–1650. ACM, 2009. 4.3.1

[130] 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. 3.1.2, 4.1.3, 5.1.3

[131] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A nutrition label for privacy. In *Proc. SOUPS*, 2009. 2.3

[132] Patrick Gage Kelley, Michael Benisch, Lorrie Faith Cranor, and Norman Sadeh. When are users comfortable sharing locations with advertisers? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2449–2452. ACM, 2011. 4.1.2

[133] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. A conundrum of permissions: installing applications on an android smartphone. In *Financial Cryptography and Data Security*, pages 68–79. Springer, 2012. 2.2.1, 6.1.6

[134] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Privacy as part of the app decision-making process. In *Proc. CHI*, pages 3393–3402. ACM, 2013. 1, 2.2.1, 2.4.1, 2.5, 4.1.2

[135] Kenneth Olmstead, Michelle Atkinson. Apps Permissions in the Google Play Store. Pew Research. `http://www.pewinternet.org/2015/11/10/apps-permissions-in-the-google-play-store/`. Published:2015-10-01, Ac-

cessed: 2017-05-14. 4, 4.1.3, 5.1.5, 6.1.1

[136] Soyun Kim and Michael S Wogalter. Habituation, dishabituation, and recovery effects in visual warnings. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 53, pages 1612–1616. Sage Publications, 2009. 5.1.2

[137] Jennifer King. How come i'm allowing strangers to go through my phone? smartphones and privacy expectations. In *Proc. SOUPS*. 4.3.1

[138] John Krumm. Inference attacks on location tracks. *Pervasive computing*, pages 127–143, 2007. 4.1.2

[139] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research methods in human-computer interaction*. 4.1.5

[140] Pedro Giovanni Leon, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Saranga Komanduri, Florian Schaub, Manya Sleeper, Yang Wang, Shomir Wilson, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Nudges for Privacy Security: Understanding and Assisting Users' Choices Online. *Manuscript submitted for publication*, 2015. 2.4, 6.2.2

[141] Lin Liao, Dieter Fox, and Henry Kautz. Location-based activity recognition. In *Proceedings of the 18th International Conference on Neural Information Processing Systems*, pages 787–794. MIT Press, 2005. 4.1.2

[142] 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. 1, 2.2.1, 2.3, 2.4.1, 2.5, 4, 4.1.2, 4.3.1, 4.3.2

[143] 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. 2.4.1, 4.1.2, 4.1.5, 5.3.1

[144] 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. 2.4.1, 5.3.1

[145] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, SA Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Proc. SOUPS*, 2016. 5.3.1

[146] Eric Malmi and Ingmar Weber. You are what apps you use: Demographic prediction based on user's apps. *arXiv preprint arXiv:1603.00059*, 2016. 4.1.2

[147] Dana Mattioli. On orbitz, mac users steered to pricier hotels. *Wall Street Journal*, 23: 2012, 2012. 1, 4.1.2, 6.2.2

[148] Gary McGraw and Edward W Felten. *Securing Java: getting down to business with mobile code*. John Wiley & Sons, Inc., 1999. 1, 4, 4.1.8, 6.1.3

[149] Kristopher Micinski, Daniel Votipka, Rock Stevens, Nikolaos Kofinas, Michelle L Mazurek, and Jeffrey S Foster. User interactions and permission use on android. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 362–373. ACM, 2017. 4.1.2

[150] Jonathan Mugan, Tarun Sharma, and Norman Sadeh. Understandable learning of privacy preferences through default personas and suggestions. Technical report, Institute for Software Research - Carnegie Mellon University, 2009. URL `http://reports-archive.adm.cs.cmu.edu/anon/isr2011/CMU-ISR-11-112.pdf`. 5.3.1

[151] Dan Murray and Kevan Durrell. Inferring demographic attributes of anonymous internet users. In *Web Usage Analysis and User Profiling*, pages 7–20. Springer, 2000. 4.1.2

[152] Mohammad Nauman, Sohail Khan, and Xinwen Zhang. Apex: Extending Android Permission Model and Enforcement with User-defined Runtime Constraints. In *Proc. CCS*, 2010. 2.2.1

[153] Helen Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004. 4.3.2

[154] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *J. Consum. Affairs*, 41(1):100–126, 2007. ISSN 1745-6606. doi: 10.1111/j.1745-6606.2006.00070.x. URL `http://dx.doi.org/10.1111/j.1745-6606.2006.00070.x`. 6.1.1

[155] US Department of Health, Human Services, et al. Summary of the hipaa privacy rule. *Washington, DC: Department of Health and Human Services*, 2003. 4.1.2

[156] Tadashi Okoshi, Julian Ramos, Hiroki Nozaki, Jin Nakazawa, Anind K Dey, and Hideyuki Tokuda. Attelia: Reducing User's Cognitive Load due to Interruptive Notifications on Smart Phones. In *Proc. PerCom*, 2015. 3.3.4

[157] Antti Oulasvirta, Sakari Tamminen, Virpi Roto, and Jaana Kuorelahti. Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile hci. In *Proc. CHI (CHI)*, pages 919–928. ACM, 2005. 5.2.2

[158] Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio, and Petri Myllymäki. Long-term effects of ubiquitous surveillance in the home. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 41–50. ACM, 2012. 6.2.3

[159] Paddy Underwood - Product Manager. Privacy checkup is now rolling out. `http://newsroom.fb.com/news/2014/09/privacy-checkup-is-now-rolling-out/`. Published: 2014-11-04, Accessed: 2014-12-08. 1

[160] Sameer Patil, Xinru Page, and Alfred Kobsa. With a little help from my friends: can social navigation inform interpersonal privacy preferences? In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 391–394. ACM, 2011. 2.4.2

[161] Sameer Patil, Roman Schlegel, Apu Kapadia, and Adam J Lee. Reflection or action?: how feedback and control affect location sharing decisions. In *Proc. CHI*, pages 101–110. ACM, 2014. 2.4.1

[162] Anand Paturi, Patrick Gage Kelley, and Subhasish Mazumdar. Introducing privacy threats from ad libraries to android users through privacy granules. In *Proceedings of NDSS Workshop on Usable Security (USEC'15)*. Internet Society, 2015. 1, 2.2.1, 2.4.1, 2.5, 4.1.2, 4.1.3, 5.1.3

[163] Joseph Phelps, Glen Nowak, and Elizabeth Ferrell. Privacy concerns and consumer willingness to provide personal information. *Journal of Public Policy & Marketing*, 19(1): 27–41, 2000. 4.1.2

[164] Daniel Preoţiuc-Pietro, Svitlana Volkova, Vasileios Lampos, Yoram Bachrach, and Nikolaos Aletras. Studying user income through language, behaviour and affect in social media. *PloS one*, 10(9):e0138717, 2015. 4.1.2

[165] Andrew Raij, Animikh Ghosh, Santosh Kumar, and Mani Srivastava. Privacy risks emerging from the adoption of innocuous wearable sensors in the mobile environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 11–20. ACM, 2011. 6.2.3

[166] Catharine H Rankin, Thomas Abrams, Robert J Barry, Seema Bhatnagar, David F Clayton, John Colombo, Gianluca Coppola, Mark A Geyer, David L Glanzman, Stephen Marsland, et al. Habituation revisited: an updated and revised description of the behavioral characteristics of habituation. *Neurobiology of learning and memory*, 92(2):135–138, 2009. 5.1.2

[167] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Symposium on Usable Privacy and Security (SOUPS)*, 2016. 4.3.2

[168] Ramprasad Ravichandran, Michael Benisch, Patrick Gage Kelley, and Norman M Sadeh. Capturing social networking privacy preferences: Can default policies help alleviate tradeoffs between expressiveness and user burden? In *Privacy Enhancing Technologies*, pages 1–18. Springer, 2009. 5.3.1

[169] Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, and Jinghai Rao. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal Ubiquitous Comput.*, 13(6):401–412, August 2009. ISSN 1617-4909. doi: 10.1007/s00779-008-0214-3. URL `http://dx.doi.org/10.1007/s00779-008-0214-3`. 5.3.1

[170] Stuart E Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. The emperor's new security indicators. In *Security and Privacy, 2007. SP'07. IEEE Symposium on*, pages 51–65. IEEE, 2007. 2.4.2

[171] Howard J Seltman. Experimental design and analysis. Technical report, 2015. 5

[172] Suranga Seneviratne, Aruna Seneviratne, Prasant Mohapatra, and Anirban Mahanti. Predicting user traits from a snapshot of apps installed on a smartphone. *ACM SIGMOBILE Mobile Computing and Communications Review*, 18(2):1–8, 2014. 4.1.2

[173] Fuming Shih, Ilaria Liccardi, and Daniel J. Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proc. CHI*. ACM, 2015. 4, 4.1.2, 4.3.1

[174] 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. ISBN 978-1-4503-2473-1. doi: 10.1145/2556288.2557421. URL `http://doi.acm.org/10.1145/2556288.2557421`. 1, 2.3, 2.5, 4.3.1

[175] Mario Silic and Dianne Cyr. Colour arousal effect on users' decision-making processes in the warning message context. In *International Conference on HCI in Business, Government and Organizations*, pages 99–109. Springer, 2016. 4.1.3, 5.1.3

[176] Sarah Spiekermann, Jens Grossklags, and Bettina Berendt. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proc. EC*, 2001. 2.3, 5.3.1

[177] Ryan Stevens, Clint Gibler, Jon Crussell, Jeremy Erickson, and Hao Chen. Investigating user privacy in android ad libraries. In *Workshop on Mobile Security Technologies (MoST)*, 2012. 4.1.3

[178] Anselm Strauss, Juliet Corbin, et al. *Basics of Qualitative Research*, volume 15. Newbury Park, CA: Sage, 1990. 4.2.3, 5.2.2

[179] Joshua Sunshine, Serge Egelman, Hazim Almuhimedi, Neha Atri, and Lorrie Faith Cranor. Crying wolf: An empirical study of ssl warning effectiveness. In *USENIX Security Symposium*, pages 399–416, 2009. 1, 2.4.2, 5, 5.1.2, 5.3.1, 6.1.5

[180] Joshua Tan, Khanh Nguyen, Michael Theodorides, Heidi Negrón-Arroyo, Christopher Thompson, Serge Egelman, and David Wagner. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proc. CHI*, pages 91–100.

ACM, 2014. 2.2.2, 4, 4.1.2

[181] Richard H Thaler and Cass R Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, 2008. 1, 2.3, 2.4, 6.1.2

[182] Christopher Thompson, Maritza Johnson, Serge Egelman, David Wagner, and Jennifer 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*, page 1. ACM, 2013. 2.4.1, 2.5

[183] Paula Thorley, Elizabeth Hellier, and Judy Edworthy. Habituation effects in visual warnings. *Contemporary ergonomics*, pages 223–230, 2001. 1, 5.1.2, 5.1.5, 5.3.1

[184] Paula Thorley, Elizabeth Hellier, Judy Edworthy, and Dave Stephenson. Orienting response reinstatement in text and pictorial warnings. In *Contemporary Ergonomics 2002*, pages 447–451. CRC Press, 2002. 5.1.2

[185] Scott Thurm and Yukari Iwatani Kane. Your apps are watching you. *The Wall Street Journal*, 17, 2010. 1, 2.3

[186] Janice Y Tsai, Patrick Kelley, Paul Drielsma, Lorrie Faith Cranor, Jason Hong, and Norman Sadeh. Who's viewed you?: the impact of feedback in a mobile location-sharing application. In *Proc. CHI*, pages 2003–2012. ACM, 2009. 2.4.1

[187] Janice Y Tsai, Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Location-sharing technologies: Privacy risks and controls. *Journal of Law and Policy for the Information Society (ISJLP)*, 6:119, 2010. 4.1.2, 4.1.3, 4.2.2, 4.3.2

[188] Blase Ur, Patrick Gage Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, and Lorrie Cranor. How does your password measure up? the effect of strength meters on password creation. In *Proc. USENIX Security Symposium*. USENIX, 2012. 2.4.2

[189] Jennifer Valentino-Devries, Jeremy Singer-Vine, and Ashkan Soltani. Websites vary prices, deals based on users' information. *Wall Street Journal*, 24, 2012. 4.1.2, 6.2.2

[190] Wall Street Journal. Apple Bows to iPhone Privacy Pressures. `http://on.wsj.com/160kjhv`. Published: 2012-2-16, Accessed: 2014-9-14. 1, 2.1.2, 3.4.1, 4.2.2

[191] Pengfei Wang, Jiafeng Guo, Yanyan Lan, Jun Xu, and Xueqi Cheng. Your cart tells you: Inferring demographic attributes from purchase data. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, pages 173–182. ACM, 2016. 4.1.2

[192] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain For-

get, and Norman Sadeh. A field trial of privacy nudges for facebook. In *Proc. CHI*, 2014. 2.4.2

[193] Roy Want, Bill N Schilit, and Scott Jenson. Enabling the internet of things. *Computer*, (1):28–35, 2015. 6.2.3

[194] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. Android permissions remystified: A field study on contextual integrity. In *Proc. USENIX Security 15 - To appear.*, 2015. 4.1.2, 4.3.1, 4.3.2

[195] Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor, Jay Springfield, Sae Young Jeong, and Arun Balasubramanian. Privacy manipulation and acclimation in a location sharing application. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 549–558. ACM, 2013. 2.4.1, 4.1.5, 5.3.1

[196] Allison Woodruff, Vasyl Pihur, Sunny Consolvo, Lauren Schmidt, Laura Brandimarte, and Alessandro Acquisti. Would a privacy fundamentalist sell their dna for $1000... if nothing bad happened as a result? the westin categories, behavioral intentions, and consequences. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014. 5.3.1

[197] Min Wu, Robert C Miller, and Simson L Garfinkel. Do security toolbars actually prevent phishing attacks? In *Proc. CHI*, pages 601–610. ACM, 2006. 4.1.8, 4.3.1, 6.1.4

[198] Ka-Ping Yee. Guidelines and strategies for secure interaction design. *Security and Usability: Designing Secure Systems That People Can Use*, pages 247–273, 2005. 1, 4, 4.1.8, 4.3.1, 6.1.3

[199] Yuan Zhong, Nicholas Jing Yuan, Wen Zhong, Fuzheng Zhang, and Xing Xie. You are where you go: Inferring demographic attributes from location check-ins. In *Proceedings of the eighth ACM international conference on web search and data mining*, pages 295–304. ACM, 2015. 4.1.2

[200] Yajin Zhou, Xinwen Zhang, Xuxian Jiang, and Vincent W Freeh. Taming information-stealing smartphone applications (on android). In *Trust and Trustworthy Computing*, pages 93–107. Springer, 2011. 2.2.1

Appendix A

Additional Statistical Results from Study 2

Condition	$\tilde{\chi}^2$	df	p
Baseline	1.44	1	0.23
Frequency	2.51	1	0.11
Background	0.008	1	0.93
Purposes	2.81	1	0.09
Purposes+Example	<0.0001	1	1
Inferences	1.84	1	0.18
Inferences+Example	<0.0001	1	1
Predictions	0.0002	1	0.99
Predictions+Implications	0.88	1	0.35

Table A.1: The difference between how iOS and Android respondents in corresponding conditions interacted with nudges is not statistically significant. This table shows the results of the statistical tests.

```
lm(formula = Age ~ OS + Decisions + OS * Decisions,data = decisions_age_os)
Coefficients:
                                   Estimate   Std. Error   t value   Pr(>|t|)
(Intercept)                        35.3136    0.4902       72.044    < 0.001 ***
OSiPhone                           -1.6864    0.8063       -2.091    0.0368 *
DecisionsCloseScreen               -4.5560    1.7528       -2.599    0.0095 **
DecisionsDontKnow                  -2.9136    4.3512       -0.670    0.5033
DecisionsKeepSettings              -1.7525    0.9553       -1.834    0.0669 .
OSiPhone:DecisionsCloseScreen      3.8177     3.7239       1.025     0.3056
OSiPhone:DecisionsDontKnow         6.7864     6.5352       1.038     0.2994
OSiPhone:DecisionsKeepSettings     -0.9858    1.7474       -0.564    0.5728

Signif. codes:0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
Residual standard error: 9.668 on 853 degrees of freedom
Multiple R-squared: 0.02044 Adjusted R-squared: 0.0124
F-statistic: 2.543 on 7 and 853 DF, p-value: 0.01356
```

Table A.2: This table shows the results of a statistical test between age, operating system, options chosen by respondents, and the interaction between them. The results shows no significant correlation between options chosen by participants and age when blocking by the operation system. In other words, changing the operating system does not change how age correlates with options chosen by respondents.

Moderately aware vs. Extremely aware	Adjusted p-value
Baseline : Frequency	< 0.001
Baseline : Background	0.009
Baseline : Purposes	< 0.001
Baseline : Purposes+Example	0.01
Baseline : Inferences	< 0.001
Baseline : Inferences+Example	0.004
Baseline : Predictions	< 0.001
Baseline : Predictions+Implications	< 0.001

Table A.3: This table shows the pairwise comparison with FDR correction between two awareness levels (Moderately aware vs. Extremely aware) in the baseline in comparison to treatment conditions.

Likelihood	Adjusted p-value
Frequency : Inferences	0.15
Frequency : Inferences+Example	0.17
Frequency : Predictions	0.26
Frequency : Predictions+Implications	0.44
Inferences : Inferences+Example	0.64
Inferences : Predictions	0.58
Inferences : Predictions+Implications	0.55
Inferences+Example : Predictions	0.95
Inferences+Example : Predictions+Implications	0.81
Predictions : Predictions+Implications	0.95

Table A.4: This table shows the pairwise comparison with FDR correction between five treatment conditions for the likelihood question. These five conditions were shown to be singificantly different from the baseline.

Awareness	Adjusted p-value
Frequency : Background	0.005
Frequency : Purposes+Example	0.01
Frequency : Inferences	< 0.001
Frequency : Inferences+Example	0.012
Frequency : Predictions	0.002
Predictions+Implications : Background	0.001
Predictions+Implications : Purposes	0.02
Predictions+Implications : Purposes+Example	0.002
Predictions+Implications : Inferences	< 0.001
Predictions+Implications : Inferences+Example	< 0.001
Predictions+Implications : Predictions	0.005

Table A.5: This table shows the pairwise comparison with FDR correction between treatment conditions for the awareness question.