Carnegie Mellon University

CARNEGIE INSTITUTE OF TECHNOLOGY

THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF Doctor of Philosophy

TITLE **Mitigating the Risks of Smartphone Data Sharing: Identifying Opportunities and Evaluating Notice**

PRESENTED BY **Rebecca Balebako**

ACCEPTED BY THE DEPARTMENT OF

Engineering and Public Policy

Lorrie Cranor	September 26, 2014
ADVISOR, MAJOR PROFESSOR	DATE
Douglas Sicker	September 29, 2014
DEPARTMENT HEAD	DATE

APPROVED BY THE COLLEGE COUNCIL

Vijayakumar Bhagavatula	October 2, 2014
DEAN	DATE

Mitigating the Risks of Smartphone Data Sharing: Identifying Opportunities and Evaluating Notice

Submitted in partial fulfillment of the requirements for
the degree of
Doctor of Philosophy
in
Engineering and Public Policy

Rebecca Balebako
B.S., Math and Asian Studies, Mount Holyoke College
A.L.M., Information Technology, Harvard University

Carnegie Mellon University
Pittsburgh, PA
December, 2014

Keywords: privacy, mobile, smartphone, usability, user experience, notifications, data sharing, risk communication

Abstract

As smartphones become more ubiquitous, increasing amounts of information about smartphone users are created, collected, and shared. This information may pose privacy and security risks to the smartphone user. The risks may vary from government surveillance to theft of financial information.

Previous work in the area of smartphone privacy and security has both identified specific security flaws and examined users' expectations and behaviors. However, there has not been a broad examination of the smartphone ecosystem to determine the risks to users from smartphone data sharing and the possible mitigations. Two of the five studies in this work examine the smartphone data sharing ecosystem to identify risks and mitigations. The first study uses multi-stakeholder expert interviews to identify risks to users and the mitigations. A second study examines app developers in order to quantify the risky behaviors and identify opportunities to improve security and privacy.

In the remaining three of five studies discussed in this work, we examine one specific risk mitigation that has been popular with policy-makers: privacy notices for consumers. If done well, privacy notices should inform smartphone users about the risks and allow them to make informed decisions about data collection. Unfortunately, previous research has found that existing privacy notices do not help smartphone users, as they are neither noticed nor understood. Through user studies, we evaluate options to improve notices. We identify opportunities to capture the attention of users and improve understanding by examining the timing and content of notices.

Overall, this work attempts to inform public policy around smartphone privacy and security. We find novel opportunities to mitigate risks by understanding app developers' work and behaviors. Also, recognizing the current focus on privacy notices, we attempt to frame the debate by examining how users' attention to and comprehension of notices can be improved through content and timing.

I dedicate this my family. To my husband – it wouldn't have been possible without you. To Meghan, my qualifiers baby, and Dissi, my thesis baby – it might have been possible without you, but it wouldn't have been as much fun.

Acknowledgments

First and foremost, thanks to my committee: Lorrie Cranor, Alessandro Acquisti, Jaeyeon Jung, and Jon Peha. This research would certainly not have been possible without your guidance and encouragement.

Many thanks to all my co-authors in these and other papers (in no particular order): Pedro Leon, Rich Shay, Cristian Bravo-Lillo, Blase Ur, Manya Sleeper, Laura Brandimarte, Eyal Pe'er, Abigail Marsh, Jason Hong, Jialiu Lin, Jaeyeon Jung, Wei Lu, Carolyn Nguyen, Alessandro Acquisti, Yang Wang, Idris Adjerid, Florian Schaub, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. Additional thanks to the lab-mates and partners in crime in CUPS and the Nudge group.

For advice, perspective. and for making an internship at Microsoft fascinating, I thank Carolyn Nguyen.

I acknowledge and thank the Global Communication Center for providing feedback on the readability of early versions of the introduction and conclusion.

Thanks to the anonymous reviewers whose insights made my work better.

Thanks for the John and Clair Bertucci Fellowship, the Google Anita Borg scholarship, and an NSF IGERT grant for partially funding my studies. Additional acknowledgements for individual chapters are provided at the end of each chapter.

Thanks to my advisor, Lorrie Faith Cranor, who has gone to bat for me on numerous occasions. Thank you for supporting me as a scholar and a person.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Smartphones allow for increasing amounts of data to be created about the smartphone users. This data can be collected and shared with a variety entities, including application (app) developers, smartphone platform developers, telecommunications providers, the government, and even malicious attackers. Data sharing with a variety of players can create security risks or privacy concerns for smartphone users. These risks and concerns include government surveillance, the theft of financial information, and even just a creepy feeling that someone knows too much.

Previous research has identified security flaws in the systems running the smartphones or apps, or have identified new technical options to reduce the flaws or identify malware. In contrast, the work presented here looks at human decision-making about privacy and security. In other research around human decision-making, users' expectations and behaviors of smartphone privacy have been examined. By understanding expectations and behaviors, researchers have identified which types data smartphone users are most concerned about, and what they are doing, or not, to protect their phones. We, however, evaluate the harms and risks to users first through stakeholder expert interviews. From there, we examine opportunities to mitigate privacy and security harms. Furthermore, recognizing that notice is a necessary part of mitigating privacy harms, we examine users' understanding and memory of privacy notices on smartphones.

1.1 Thesis Statement

I identify the privacy risks to smartphone users from data sharing, including privacy and security risks related to data sharing by apps, and the mitigations of those risks. These mitigations may be made by smartphone users, app developers, or platform developers. We find novel opportunities to mitigate risks by interviewing and surveying app developers' work and behaviors. The opportunities include usable software development tools that encourage privacy protective behaviors. As current public policy efforts are focusing on notice about application data sharing, I examine how users' attention to and comprehension of notices can be improved through clear content and timing. We find that users do not understand the data sharing ecosystem that is largely hidden from them. Our results also show that privacy notices should be showing during app usage in order to improve recall.

1.2 Thesis Overview

This chapter provides an overview of this thesis. The background and related work section motivates and defends our focus on smartphone privacy notices. It provides information relevant to all the following chapters about the current state of data collection on smartphones, as well as what users' concerns are and how public policy is addressing those concerns.

Previous research has not provided a broad view of the smartphone ecosystem to determine the risks to users and the possible mitigations to the risks. We do so in Chapters 2 and 3.

In Chapter 2, we identify the risks to users from smartphone data sharing. We do so by interviewing privacy and security experts from multiple stakeholder groups. We find that there are a variety of risks, and transparency in the form of a privacy notice is one of several mitigations that can reduce the risk. Notices can inform users about data collection, and this may allow them make appropriate individual decisions to protect themselves based on their own context and preferences. Other risks cannot be mitigated by individuals and require different interventions, such as improved security and data minimization.

In Chapter 3, we examine the privacy and security behaviors of app developers. App developers make many decisions about what data to collect about their users, how it is transmitted and stored, and with whom to share data. By understanding the behaviors of app developers, we are able to quantify the risky behaviors and identify opportunities to improve security and privacy. We interviewed and surveyed app developers to understand why they are not able to fully implement best privacy and security practices. Finding that lack of resources and time particularly constrain small app development companies, we propose that usable and inexpensive privacy tools be integrated into the app software development process.

In Chapters 4, 5, and 6, we examine one specific risk mitigation that has been popular with policy-makers: privacy notices for consumers. If done well, privacy notices should inform smartphone users about the risks and allow them to make informed decisions about data collection. Unfortunately, previous research has found that existing privacy notices do not help smartphone users, as they are often neither noticed nor understood [107, 45, 76]. Therefore, we evaluate options to improve notices.

The study described in Chapter 4 on user understanding of terms is specifically designed to inform a public policy effort to create a national code of conduct on smartphone short-form notices. We examine the terms developed by a multi-stakeholder group moderated by the National Telecommunication and Information Administration (NTIA). In an on-line survey, we find that that some terms were particularly confusing. These confusing terms required additional knowledge, either about the phone operating system or about how data is bought and sold. Also, this study demonstrates the need for user testing when privacy policies are developed. We use this experience as case study to help usable privacy practitioners understand their role in public policy-making.

In Chapter 5, we examine whether users notice or and understand privacy notices if the content and timing of the notices is modified. We designed privacy notices that inform users about what data was requested by apps, who requested the data, and how often it was collected. In a lab study, participants played two popular app games with our privacy notices. We find that users were surprised by the frequency and destination of data collection. Our study suggests that users gain a better understanding of the data collection ecosystem when shown more information during

the app usage. Additionally, participants value the new information provided in the notices.

Smartphone platforms differ on whether apps show privacy notices during installation or during app usage. In Chapter 6 we describe an online and field study which addresses the question of when privacy notices should be shown. We examined whether the timing of notices impacted users' memory of the content of the notice. In a field study, users were shown a prototype of the NTIA privacy notice, and asked to recall 24 hours later what data was collected and with whom it was shared. We find participants were most likely to recall notices shown during app use. In contrast, notices displayed in the app store are not well remembered.

In Chapter 7, we conclude this document with policy recommendations. Based on our research, we recommend that policy focus on putting some burden of improved privacy and security on platforms, app stores, and other stakeholders in the data sharing and app development ecosystem. In particular, we believe policy should focus on data minimization and best security practices. We recommend that policy makers, platform developers and app developers continue to integrate user-studies into the development of privacy nudges and privacy notices. Finally, we discuss some particular concerns about relying on notice to inform users about the risks of data sharing. We discuss a large hurdle to improved notices: smartphone users are not aware of the data sharing ecosystem.

1.3 Thesis Contributions

The research presented here draws on results and methods from the fields of human computer interaction, behavioral economics, risk communication, and human decision making. Through this interdisciplinary approach, this work attempts to inform public policy around smartphone privacy concerns from data sharing. The findings presented in this work can help policy-makers in the area of smartphone privacy in two ways; one is identifying risks and opportunities to mitigate privacy risks. This includes improving app developer privacy behaviors through understanding their decisions about data collection. Second, recognizing the current focus on privacy notices, we attempt to inform efforts to improve notices by examining how users' attention to and comprehension of notices can be improved through better content and timing. Furthermore, we identify gaps in user understanding of data sharing.

Our results may also inform platform-developers and app developers who wish to design privacy notices that users will remember and understand. Finally, our work builds on and contributes to research on usable privacy.

1.4 Background and Related Work

This section provides background on smartphones and privacy and security notices. First, we describe the characteristics of smartphones that make their users vulnerable to privacy and security concerns. Then, we briefly describe the smartphone ecosystem, including the major platforms and how they address some of the privacy and security concerns. In the following section, we summarize previous work exploring users' concerns about smartphone data collection.

The work presented in this thesis aims to inform and improve public policy around smartphone notices. Therefore, we give an overview of public policy regarding privacy, and then we focus on smartphone privacy efforts by policy-makers.

1.4.1 Privacy and Security on Smartphones

In this work, we define smartphones as mobile devices that serve as phones, have internet connectivity, and allow the owner to download and install additional software applications, popularly known as "apps."

Smartphones have characteristics that distinguish them from personal computers, and these characteristics create additional privacy harms and concerns. Smartphones are smaller than PCs, and users tend to carry them wherever they go. The portability and size allows for a greater chance of loss and theft. Smartphone sensors (e.g. microphone or GPS) permit increased data collection. This allows inferences about the users' behavior [89, 144, 163]. The increased data from sensors also increases the possibilities of eavesdropping [112]. Smartphones are a relatively new technology, and many of the security and privacy techniques users have learned for PCs don't apply to smartphones [52]. There are several vulnerabilities specific to smartphones. Smartphones can be susceptible to particular malware attacks that use smartphone's ability to call premium-rate numbers. Other attacks may include direct access to financial information as mobile money usage increases [39]. Additionally, the smaller batteries and memory capabilities reduce the capacity of security solutions [112, 39]. Smartphones also have reduced screen sizes, which limits the ability to communicate complicated ideas or show security icons such as SSL indicators [28].

There are a number of players in the smartphone data sharing ecosystem. Several companies and groups of people make decisions about smartphone data sharing and collection. These include the platform developers, the carriers that provide the data transfer and calls, the app developers that build apps, and the advertising and analytics companies that collect data on users for various services. This thesis will focus on users and app developers, but to gain a broader understanding of the ecosystem, we provide a broad strokes description of the major platforms and the app developer market.

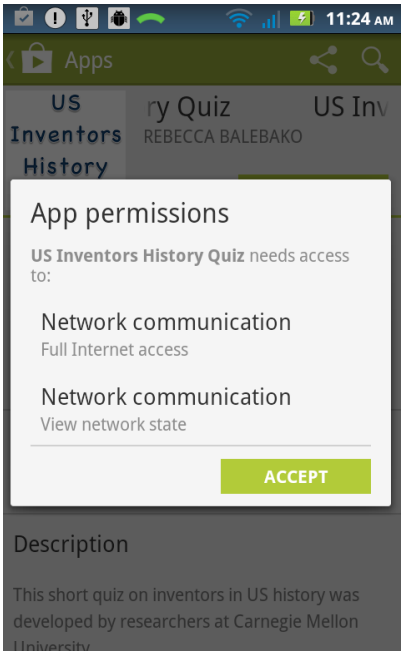

Figure 1.1: Android permission notice

The major smartphone platforms in the United States are currently Google's Android, and Apple's iPhone. The platform companies develop and manage several important components: the operating system, the stores to purchase or download apps, and software development kits so that app developers can create new apps. Google and Apple rely on the carriers to provide system and security updates to their users' phones.

Much of this thesis focuses on the data collection initiated by apps. The smartphone app industry is growing, with both Android Play Store and Apple iTunes offering over one million apps each [104, 21]. The app development workforce is highly fragmented; there is no dominant app development company holding a large share of the market. Apps are often initially created by independent developers or small startups, as opposed to large, established corporations [23, 74]. This fragmented market means a large number of independent developers, with fewer personnel and resources, are making decisions about their users' sensitive data.

Both Google and Apple provide some mechanisms to police apps in their stores for security issues. Apps for Android are available in Google's app market. Since 2012, new apps added to the market are automatically scanned for malware by Google.[1] Apple requires developers to register for a developer ID before submitting apps to their app store. Apple reviews apps that are submitted for many factors, including possible security concerns.[2]

On Android phones, the standard privacy notice for apps consists of a notification when an app is installed about which permissions the app can access. An example is shown in Figure 1.1. This notification lists which of 130 possible permissions, including location and network communication, that the app has requested.[3] Users may either accept all the permissions and install the app, or they may choose to stop the install. Research on Android permissions finds that the this install system is not effective in informing users about permissions, due to lack of user attention and comprehension [76, 107]. In June of 2014, Google simplified the list of app permissions by grouping them into the 13 most important group of permissions, and hiding some of the most common permissions.[4]

The iPhone system uses notifications the first time certain data, such as location, is accessed by an app. An example is shown in Figure 1.2. iPhone also includes several additional privacy settings. One setting allows users to control whether each app can have access to a short list of permissions, such as location and contacts. Another setting allows users to "limit ad tracking," which stops sending the phone's unique id and prevents tracking across apps [13].

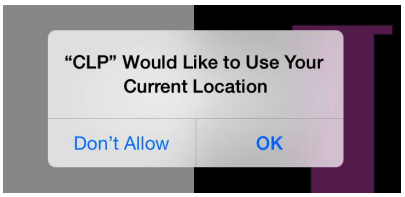

Figure 1.2: iOS permission notice

1.4.2 User Concerns about Privacy and Security

A wealth of previous work has examined users' perceptions and desires for smartphone privacy and security. A Pew Internet Study found that smartphone users have concerns about sharing personal information, reporting that "57% of all app users have either uninstalled an app over concerns about having to share their personal information, or declined to install an app in the first place for similar reasons." [53]. In this section we describe some of the concerns users have about data privacy and security on smartphones.

Several categories of smartphone data raise privacy concerns. Biometric data can serve as a unique identifier for linking to a user's other activities [58]. These unique identifiers can cause particular privacy concerns as they often cannot be revoked or changed, even when stolen [142]. Users' concerns about the collection of their browsing history have been documented a number of times [131, 156, 116]. Additional privacy issues inherent in the collection of metadata, such as logs of browsing, phone calls, or text messages, have been publicized in the wake of revelations about the U.S. National Security Agency's PRISM program. Phone usage data and metadata can

[1] **http://googlemobile.blogspot.com/2012/02/android-and-security.html**
[2] **http://www.apple.com/osx/what-is/security.html**
[3] **http://developer.android.com/reference/android/Manifest.permission.html**
[4] **https://support.google.com/googleplay/answer/6014972?p=app_permissions&rd=1**

be collected by apps and used to infer hobbies, medical conditions, and beliefs [151]. Users' beliefs and activities can often be inferred from the people with whom they associate [151].

Smartphone phone users may wish to constrain the collection of their contact information [157]. The collection of users' contacts has led to privacy outrage in the past, such as when Facebook's smartphone app was discovered uploading the names and phone numbers from users' address books to Facebook's servers without providing notice [32]. The metadata from users' emails alone can be used to infer their real-life social network and associations [148, 98]. Furthermore, the fact that data is collected can have a chilling effect on individuals' free speech [150], and most individuals would likely be unaware when their data and metadata could reveal them to be violating the law [126].

Sensitive information may exacerbate privacy concerns. Financial information can cause privacy issues both because individuals might be loath to disclose information about their earnings, as well be concerned about the potential of price discrimination [158]. Similarly, privacy is fundamental to a doctor-patient relationship, and disclosure of health information could cause financial harm if used by a health-insurance company to deny coverage to a patient [30]. Information collected by fitness apps may include sensitive information that could be sold to insurance companies [17].

Location data can also arouse privacy concerns [135], particularly when the location is not visited by many people [154] or when the location information is highly granular [44]. Users also believe their files, such as photos and videos, to be sensitive [135]. Furthermore, nearly all participants in a study by Felt et al. would have been upset if the text messages and emails stored on their phone were shared publicly [75].

Users also have concerns about the security of their phones, including the physical theft of their data, malware, and wireless network attacks, or unauthorized calls to 1-900 numbers [59, 75].

In addition to the type of data, users are concerned about with whom the data is shared. Social networks, government, and advertisers may all be of particular concern. A Pew Research Study found that 63% of Americans would feel their privacy had been violated if they knew the government had collected information about their calls and online communication [153]. In addition, social networks may be a concern due to the accidental leakage of private information (willingly provided by the user) to unanticipated parties [110, 87]. A global study found that vast majority of respondents did not like people knowing information about them or their habits unless they had themselves shared it - a specific concern for smartphone sensors that may collect data passively without the user's awareness [22].

1.4.3 Public Policy

We provide an overview of three major attempts to define consumer privacy principles before discussing the policies specific to mobile devices and smartphones.

More than 30 years ago, the Organization for Economic Cooperation and Development (OECD) published the Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. Most of the OECD's eight principles focus on the data collector's responsibility, including limited collection of data, and collection of data limited to specified purposes [2].

The Federal Trade Commission (FTC) may be the most powerful and influential body in privacy jurisdiction, and that notice and choice is one of the "most central aspects" of the jurisprudence [149]. In 1998, the FTC released the Fair Information Practice Principles (FIPPs) that focus

on the consumer's role in managing their data [3]. The principles as defined by the FTC are summarized here.

1. Notice/Awareness: This prerequisite for other rights says notices should inform consumers about data collection.
2. Choice/Consent: Consumers should have options about data collection.
3. Access/Participation: Consumers should be able to view data about themself and ensure data are accurate.
4. Integrity/Security: Collectors must take steps to maintain accurate data and secure it from unauthorized access.
5. Enforcement/Redress: There must be a means to enforce the above rights.

Recently, the FTC issued a report on Mobile Privacy Disclosures, which included specific guidelines that app developers, smartphone platforms, and advertising networks could use to improve notice to smartphone users about data collection. The FTC has also published advice to app developers recommending that they prioritize security and minimize data collection [1]. The FTC has also endorsed "Do Not Track," a simplified mechanism allowing consumers to indicate if they wish to receive targeted ads [6]. The California Attorney General issued advice on protecting mobile privacy to app developers, platforms, ad networks, and carriers. The advice to app developers included understanding their data collection, developing a privacy policies, and limiting data collection [91].

In 2012, the White House issued a Consumer Privacy Bill of Rights, which advanced seven rights for consumers over their electronic data [11] . These rights are:

1. Control
2. Transparency
3. Respect for Context
4. Security
5. Access and Accuracy
6. Focused Collection
7. Accountability

These overlap with the FTC FIPPs; all five of the FTC FIPPs are represented in the Consumer Privacy Bill of Rights [11]. Transparency is similar to notice and awareness, and accountability is similar to enforcement and redress. The Consumer Privacy Bill of Rights adds two elements. One is "Respect for Context," defined as: "Consumers have a right to expect that companies will collect, use, and disclose personal data in ways that are consistent with the context in which consumers provide the data." The second is Focused Collection defined as: "Consumers have a right to reasonable limits on the personal data that companies collect and retain." Like the OECD principles, both of these imply that there should be some limit to the amount of data collected, based on context and the specific users of the data.

Previous work has pointed out why notice and choice cannot be the only mitigation used to protect users. Privacy law's current focus on notice and choice puts much of the decision-making burden on the user. Solove argues that this "self-management" is problematic, due to users' cognitive and structural limitations. The structural limitations include the scale of data collection and data aggregation, which makes it difficult for users to assess harm [152]. While there have been

fifteen years of notice and choice initiatives for online privacy, they have not been sufficient to protect users. This is due to lack of incentives to participate and the lack of enforcement [63]. Mandated disclosures in general (not just privacy) are ignored or misunderstood [41]. Calo recognizes the problems of typical privacy notices and argues for trying visceral notices instead of the standard textual privacy notices [57]. We use visceral notices in one condition in Chapter 5.

At least ten different documents for app developers on describing how they can protect their users' privacy have been developed and released by U.S. government agencies, trade associations, and advocacy groups. For example, the Federal Trade Commission released a staff report on Mobile Privacy Disclosures [19]. The California Attorney General provided recommendations for privacy in the mobile ecosystem [91]. To synthesize these different documents, the Information Association of Privacy Professionals (IAPP) created a web tool allowing readers to access, search, and compare ten different privacy guidelines for app developers from the US, Australia, and Europe [78].

We examined the ten guidelines in the IAPP Mobile App Privacy Tool. The guidelines often overlap, and making it feasible to comply with most or all guidelines. We summarize four pieces of advice offered by most of these guidelines.

1. App developers should minimize the amount of data collected. Minimizing data limits the app developers' liability and protects the user from unexpected and surprising data collection.

2. Developers should avoid retaining old data by defining retention periods for data and deleting old data that is not needed.

3. Developers should communicate to their users through a privacy policy.

4. Developers should encrypt sensitive data. Many guidelines suggested encrypting all data that is transmitted or stored.

This thesis was particularly informed by the Department of Commerce's National Telecommunications and Information Administration (NTIA)'s multi-stakeholder initiative on Mobile Application Transparency [20]. This group created a code of conduct for app developers that included a standardized short-form privacy notice for mobile devices. We examine this code of conduct for user understanding in Chapter 4. In the research described in Chapter 3 we asked app developers about their awareness of the code of conduct. In Chapter 6, we use a prototype of the short-form notice to test the timing of privacy notifications.

1.4.4 Human Decision-Making and Privacy Notices

Informing users about privacy and security issues is an important step in involving humans in security and privacy decisions [26, 62]. However, privacy and security are usually not the user's primary task. Furthermore, feedback or repercussions of decisions are not immediate (a hack or intended disclosure may occur days or weeks after the notice was shown) and users may not associate any consequences with the notice. The Computer Human Information Processing (C-HIP) framework discusses the stages in which humans notice and process warnings [162], and has been expanded for security warnings [62]. The C-HIP framework models the stages and variables a human may go through when presented with a security warning. First, if the warning

is conspicuous or noticeable, the user may switch attention from their task to the notice, and then maintain attention sufficiently to encode the information, i.e., perceive its content. Ideally, the user will enter the stage of memory and comprehension, depending on whether the user has sufficient prior knowledge to process and understand the notice's content. Finally, attitudes, beliefs and motivation will impact an individual's reaction to the warning notice.

In our work, we largely focus on studying participants' memory and understanding of privacy notices. However, the C-HIP model emphasizes that many aspects of notifications must work in concert to influence the behavior. For example, work on privacy notices in online social networks successfully increased attention by reorganizing a notice's format. Yet, participants did not significantly change their behavior, despite spending more time looking at the notice [82].

The design of notices and warnings needs to conserve user attention; decisions should be easy to make and unnecessary interruptions should be avoided [76]. Privacy decision-making may be overwhelming: the cognitive costs associated with considering potential ramifications of sharing data may hamper decision-making [24, 25]. Another issue is that when notices are shown too frequently, users may become habituated. Habituation may lead to users disregarding warnings, often without reading or comprehending the notice [47]. One example in which computer users demonstrate habituation is the ubiquitous End User License Agreements (EULAs). EULAs are typically complex, legal documents that do not provide users with real options, but rather ask them to accept the stated conditions before they can use the software. Even security-concerned users react to EULA-like dialogs by clicking 'accept,' instead of exercising real choice [48]. Other work has identified the effects of habituation on security dialogs and proposed some options that reduced habituation by requiring user interaction [55]. To reduce habituation, Felt et al. identified specific app permissions that should be de-emphasized [75].

While no work, to our knowledge, has isolated the impact of timing of mobile privacy notices on attention, some work has studied timing effects of privacy notices on purchases in Web and desktop contexts. The timing of privacy indicators can impact users' willingness to pay a premium when shopping on on-line websites if they are not examining multiple websites. In these cases, users paid more attention to privacy warnings that were shown in a search result before selecting, as opposed to seeing the privacy warning with the website content [70].

Experimental work supports the hypothesis that timing of privacy notices is an important factor in getting people to pay attention to and act on the privacy notice. Individuals who spend more time looking at an install-time security warning were less likely to install the software. However, people who are engrossed in the installation process may fail to pay attention to an install-time notice [85]. We observed similar results for mobile privacy notices. Other work has shown that while privacy notices that directly preceded user privacy decision-making were effective in influencing user behavior, introducing only a 15 second delay between the presentation of privacy notices and privacy relevant choices was enough to render notices ineffective at driving user behavior [27].

1.4.5 Designing Usable Privacy Notifications

Research on improving privacy policies is much needed. The privacy policies available on most websites have been deemed unusable, due to their location, form, and focus on legal content [101]. The time it would take consumers to read the privacy policy of every website they visit makes this task unfeasible [130]. In this section, we discuss some efforts to design usable privacy notices that

take into account the issues discussed in the previous section. We first discuss privacy notices for browsers and wi-fi, before we discuss work on privacy notices for smartphone apps.

A pioneering study on usability issues of privacy notification describes the development of a browser agent called Privacy Bird. Privacy Bird notifies users with sounds and icons when a website's privacy policies do not match the user's preferences. This study finds that users appreciate short summaries, meaningful terminology, and the appropriate granularity of information [64]. We attempted to integrate all of these into the design of the notifications in Chapter 5. We also examine the terminology of a privacy notice to evaluate whether it is meaningful to users in Chapter 4.

Standardized grids were found to be an effective way of presenting information about a website's use of data [106]. The notification in Chapter 5 also uses grids to visualize which information has been transmitted off the phone by applications.

Wi-Fi Privacy Ticker is a tool designed to improve users' awareness of personal information exposure over unencrypted Wi-Fi networks and provide control to prevent unwanted exposure [60]. It automatically drops a user's connection when a highly-sensitive term (as defined by the user) is being sent in the clear. A notification called the 'Ticker Display' and balloon tip provide instant notification about the data leakage. Participants used the ticker for 3 weeks and had a resulting change in awareness, as found by both open-text statements and responses to specific questions [60].

The smartphone app privacy notices currently being used by industry could be improved. An effective risk communication should focus on issues that people at risk need to know but currently ignore [133]. Current smartphone notices inform users about data sharing, but not about the consequences and actual risks of data sharing. For example, the Android permission "Internet is explained with the following text: "Allows applications to open network sockets. One warning for iOS is "[App name] would like to use your current location. This suggests that risk communication for smartphone data sharing can be improved. In this thesis, we begin with expert interviews in Chapter 2 to understand the risks to users.

Several studies demonstrate a lack of user understanding of privacy and security risks associated with installing smartphone applications. Android users find it difficult to understand the terms and wording of the Android permissions [107]. In another study, only 17% paid attention to the permissions (including ones which grant an application access to privacy-sensitive data) when installing an application. Furthermore, only tiny percentage of study participants demonstrated full comprehension of the permissions screen [76]. This was collaborated by further work on Android permissions, finding users did not understand the implications or risks associated with the permission requests [45]. Our study in Chapter 5 goes deeper into this lack of understanding and discusses users' misconceptions about data sharing with two popular game applications using a role-play technique, while our study in Chapter 6 examines whether users remember a privacy screen that is not based on permissions.

Some work has proposed providing users with a privacy score for the app. One study used crowd-sourcing to analyze users' expectations of app permissions and provide a privacy score [122]. Other work has proposed providing a "sensitivity score" or a "risk score" based on the number or type of sensitive permissions requested by an app [121, 83, 137]. This work is promising, but is not operating system agnostic; it relied heavily on an understanding of Android permissions. The notices used in Chapter 4 and Chapter 6 would be applicable to all smartphone platforms, and do not rely on Android-specific permissions.

Users may use privacy information to select between apps, if the information is provided in clear manner at a point where they can make a decision. When shown a 'privacy checklist' in the Google Play Store, and users would select the app that requested fewer permissions [108]. Additionally, when asked to compare similar apps with different permission requests, users demonstrated that they were willing to pay more for the apps with fewer permission requests [69]. A "risk" score generated by examining the permission was shown to be effective in the lab and in on-line survey in helping users choose an app that requested fewer permissions [84]. Another research prototype informed users about specific examples of data sharing, such as one of the user's photos or contacts, and found that users paid more attention to these modified notices [90]. However, users may follow different paths when selecting an app. They may not always be comparing two similar apps with different privacy impacts, but instead may select an app based on recommendations of a friend or a review, and may forgoe comparison shopping. Therefore, our work examines reactions to notices within the context of specific apps.

Chapter 2

Assessing and Mitigating the Risks of Smartphone Data Sharing

Consumers are increasingly taking advantage of the benefits of smartphones. A variety of entertainment and productivity apps, from calendars to restaurant reviews to GPS mapping, offer benefits that smartphone users find compelling. Over half of the US mobile market now is using smartphones [29]. At the same time, apps, platforms, and telecommunication carriers are collecting increasing amounts of data and transmitting it to other parties. While the benefits of data sharing may be clear to smartphone users, the potential risks and harms are not as clear.

Privacy advocates and security researchers have looked at which aspects of data collection are the most concerning to users [43, 59, 40, 75, 53, 135, 122, 107, 68]. Security experts have found security holes in smartphone platforms and have proposed solutions [86, 66, 72]. Yet, the literature lacks a holistic analysis of the harms, from tangible damages to privacy concerns, that can come to users as a result of smartphone data collection. This study tries to assess the real harms and the intervention points where policy can make a difference by mitigating these harms. In this work, we use expert interviews to evaluate the privacy and security harms that occur to users due to smartphone data sharing. Based on a series of interviews with 20 experts from 10 stakeholder groups, we enumerate the major risks to smartphone users from data sharing and the solutions proposed to mitigate these risks. We use the understandings gained through these interviews to evaluate current policy efforts.

Current policy efforts are focused on transparency and alerting smartphone users to data collection practices. Our work addresses the question of whether the notice approach is the right place to focus attention. We ask whether the status quo is sufficient, and conclude that while current efforts are useful, other areas such as security need more attention.

In the next section, we provide background on smartphone data sharing. In section III, we describe the methodology for performing the expert interviews and analyzing the results. In section IV, we itemize the harms and concerns identified by the experts, and in section V we discuss the interventions that can mitigate these harms. In section VI we discuss and synthesize our findings, offering suggestions about what risk mitigations public policy can address and how risk communications can be improved.

This chapter is largely based on paper [33] co-authored by Rebecca Balebako, Cristian Bravo-Lillo and Lorrie Faith Cranor.

2.1 Related Work

Much of the background for this chapter was provided in Chapter 1.4. In this chapter, we briefly discuss previous research on notice and choice. We address attempts to educate users about smartphone risks. Finally, we explain how expert elicitations and interviews are used to evaluate risks in policy-making, and elucidate how that method can be applied to smartphone privacy and security.

2.1.1 Notice and Choice

A brief background on policy in the United States regarding smartphone data collection has been provided in Section 1.4.3. In particular, we enumerated the Fair Information Practice Principles (FIPPs), which include notice, choice, access, security and enforcement. We pointed out that much policy has focused on notice and choice, but that there has also already been criticism of policy that relies too heavily on notice and choice.

2.1.2 Educating Users

In Section 1.4.1, we discussed the characteristics of smartphones that contribute to data collection risks. In this section, we discuss efforts to educate users about smartphone risks.

Non-profits, government agencies, and media outlets have attempted to educate smartphone users about smartphone privacy, and have offered advice to smartphone users on protecting their smartphone privacy and security. For example, the European Network and Information Agency's advice on smartphone security includes automatically locking the smartphone with a password to prevent unauthorized access, checking the reputation of apps or services to avoid malware, and clearing the phone's data ("reset and wipe") before disposing of the phone [93]. A Forbes.com article advising readers about smartphone privacy includes the above advice, as well as updating apps for security patches, using smartphone privacy settings to limit location tracking and access to other information, and closing apps when they aren't being used to limit access [127]. The Privacy Rights Clearinghouse, a non-profit consumer advocacy group, released a fact sheet titled "Privacy in the Age of the Smartphone" which informs readers that criminals, advertisers, and the government all want to "snoop" on their smartphone. In addition to protecting the phone with a password and researching apps before downloading, the privacy tips offered to consumers include contacting carriers to opt-out of data collection and advocacy such as writing to the congressional representatives for better laws [12].

We hope our research will help improve advice to consumers and other stakeholders by highlighting the harms and providing solutions to reduce the risk of the harms.

2.1.3 Expert Elicitations

Expert elicitations are one step used in creating effective risk communication for many areas of public policy [133]. The method focuses on one risk at a time, and includes five steps: creating an expert model, conducting open-ended mental model interviews, conducting structured confirmatory interviews, drafting an appropriate risk communication, and evaluating the communication. One should observe the frequency with which new concepts emerge with each interview. In most cases, after 20 or 30 interviews no new concepts will emerge [133]. Similarly, Meyer provides

guidelines for conducting expert elicitations [132], and Kynn et al. discuss how to counter some of the heuristics and biases in expert elicitations. For example, experts (and non-experts) tend to be overconfident in their estimations [111].

Expert elicitations have been used to inform public policy in a number of areas, particularly those where the risks are difficult to quantify, such as biological invasions by non-native species [124], the use of biofuels [79], and other areas investigated by the Environmental Protection Agency [8]. Expert interviews have also been used in computer security, which we describe in the next paragraphs.

Expert elicitations have been used to determine the risks of phishing to users, and to determine the capabilities and incentives of stakeholders to prevent phishing. In a method similar to ours, the researchers interviewed experts and then looked for recurring themes across the experts to synthesize high-level findings about risks and options to stakeholders. They identified several places where the stakeholder most able to fight phishing had little incentive to do so [145].

Expert elicitations have also been used in the field of computer security to compare expert and novice evaluations of security warnings. Bravo et al. interviewed both users with advanced security knowledge, and users with average or little knowledge about their reactions to security warnings. They found several key differences between advanced and novice users. For example, novice users usually decide to trust a site or software based on its look-and-feel, while advanced users would only use look-and-feel as a warning against trusting a site or software. The authors specifically mention the importance of mitigating risks where feasible, as opposed to relying on warnings [54].

Our study enumerates multiple risks and harms that may affect smartphone users. In order to find those risks that experts agree that lay users face, we used a large pool of experts (20) from different backgrounds, and report on those risks mentioned by five or more experts. While we believe this approach allows for some consistency in identified risks, further work is necessary to determine the statistical occurrence of these risks in the smartphone ecosystem.

2.2 Methodology

We interviewed 20 experts on privacy and security from different stakeholder groups about smartphone data sharing. The anonymous interviews were typically one hour long. Experts were not compensated for their time. The interviews were recorded, transcribed verbatim, and coded for themes regarding harms, risks and interventions.

2.2.1 Stakeholder Selection and Recruitment

In order to get a broad range of opinions and perspectives, we first identified nine stakeholder groups from which to select participants. All participants worked in privacy or security, and typically had experience with mobile or smartphone privacy or security. We classified experts based on the current or recent employment sectors as follows:

- *Academia* - Researchers and professors in university or research lab settings who conduct research on smartphone security or privacy.
- *Application Industry* - App developers or app industry representatives.

- *Platform providers* - Developers or managers in companies building smartphone operating systems or platforms.
- *Telecommunications providers* - Researchers or managers in companies providing telecommunication services.
- *Security Experts* - Developers or managers in company providing security solutions, or managing the security branches of IT companies.
- *Aggregator or advertiser* - Developers or managers in a company aggregating data or providing ads based on smartphone data.
- *Consumer advocates* - Non-profit agencies advocating for consumer privacy.
- *Industry or Industry Lawyers* - Representatives from online advertising and other stakeholder industry associations as well as attorneys who represent multiple industry stakeholders
- *Government* - Public policy specialists working for federal regulation agencies.
- *Privacy Industry* - Developers or managers in a company providing consumer privacy tools.

We looked for experts that had been involved in recent public policy efforts on mobile transparency such as the NTIA's Privacy Multistakeholder Process on Mobile Application Transparency (NTIA MSHP), who had published papers on mobile privacy and security, or who had been recommended by other experts. We recruited experts using a personalized email, asking them to volunteer one hour to participate in the interview on mobile privacy and security. Participants were told that they would be anonymous and were not expected to represent their employers. Overall, we interviewed 20 experts representing all of the stakeholders above. Some experts fell into two categories, due to the range of their experience. For example, an expert who worked in one field for a number of years and then recently switched employers, or an expert whose job includes multiple roles, could represent two stakeholder groups.

In addition to the twenty interviewed, ten experts were invited but did not agree to be interviewed, citing time constraints (4), constraints due to their employer or profession (2), did not provide a reason (1), or did not respond to multiple requests (3). These experts represented all of the stakeholder groups, except industry and government. Therefore, we feel that there was not a stakeholder selection bias in participation.

Table 2.1 gives an overview of the participants and the stakeholder they represented. The experts were typically well-seasoned: 13 experts had over 15 years of experience, and only two had 5 or fewer years of experience. Half (10) of the experts were participants in the NTIA MSHP.

If the expert agreed to the interview, there were asked to fill out an anonymous consent form, as required by Carnegie Mellon University Institutional Review Board. The researcher then contacted them by phone or in person for a one-hour interview. The experts were all advised that they would remain anonymous. In some cases, experts requested clarification on how they would be identified in the final report. In this case the researchers worked with the expert to help identify an anonymous description that they or their employer would be comfortable with. All experts were told they would be provided with the final report, but they were not given the option to modify or change the results. The interviews took place in the first quarter of 2013, before the eruption of news regarding government surveillance due to the Snowden leaks.

ID	Stakeholder
AC1	Academia
AC2	Academia
AC3	Academia
SE1	Security Expert & Academia
SE2	Security Expert & Platform Provider
AD1	Aggregator or Advertiser
AD2	Aggregator or Advertiser & Industry
AP1	Application Industry
AP2	Application Industry
CA1	Consumer advocate
CA2	Consumer advocate
L1	Industry or Industry Lawyer
L2	Industry or Industry Lawyer
G1	Government
G2	Government
PL1	Platform Provider
PL2	Platform Provider
PI1	Privacy Industry
TE1	Telecommunications Provider & Application Industry
TE2	Telecommunications Provider

Table 2.1: Participants who were interviewed, including stakeholder group. The numbers used in the IDs do not correspond to the order in which participants were interviewed.

2.2.2 Interview Design

The interviews were "standardized open-ended interviews" [155], also known as "semi-structured interviews." The interview script contained 10 open-ended questions regarding harms and risks of smartphone data sharing, the possibilities for reducing risks, future directions, and vulnerable populations. The researcher-interviewer asked clarifying questions or detailed questions as needed throughout the interview. The interview script is provided in Appendix A.

Great care must be taken in designing the questions for expert elicitations to ensure that that experts will be able to interpret them correctly. We conducted pilot tests with four graduate students involved in privacy and security research. An additional pilot test was conducted with a graduate student with experience running expert interviews on risks in a different domain (nuclear energy). Finally, we shared our interview script with an expert on expert elicitation for risk communication to gather feedback on the questions and coding methodology. These steps allowed us to refine the interview questions, both helping with the flow of questions, the wording of the questions, and the amount of time required to complete the interview.

We designed our questions to be neutral and open-ended. In our pilot tests, we found that interviewees were able to respond better to a specific scenario about what a user can do to avoid risk than a general question. Therefore, we framed the question about what the user can do to prevent harms and risks as, "My mother recently got a smartphone. What should she do to protect herself from the harms we discussed?" Furthermore, our pilot tests indicated that experts struggled to rank the harms in terms of likeliness or harmfulness. Therefore, we made the question less

precise and asked the experts to identify the "most" harmful and the "most" likely.

Despite our attempt to be neutral, some experts were concerned that we only asked about the risks or harms of data sharing from smartphones, instead of asking about the benefits as well. However, our goal was to identify all harms and concerns in a holistic manner, so that the appropriate mitigations can be considered, and smartphone users can continue to enjoy the benefits of smartphones.

The same researcher conducted all interviews from February to April, 2013. In some interviews a second researcher took notes. Five interviews were done in person, in private offices. All other interviews were conducted remotely. We recorded the audio of all interviews, except for two participants who declined to be recorded. The interviewer refrained from offering personal opinions or reacting emotionally to responses, and tried to take notes consistently throughout the interview. If the interviewer was unclear about a response, she tried to re-phrase it neutrally and give the interviewee a chance to respond and clarify.

2.2.3 Results Coding

To code the results, we used "emergent coding" to create a list of themes [113]. Two researchers independently reviewed the notes and transcripts of 15 interviews to create coding sheets for themes. Then they compared the two sets of codes to resolve differences and create a consolidated list of themes. A third researcher with experience coding expert interviews acted as a moderator to help define the major themes. One researcher then coded the transcripts using the themes identified in the above process. The transcripts were marked to identify salient quotes, frequency of comments, and also to identify which stakeholders discussed which themes.

2.2.4 Limitations

Qualitative interviews allow for in-depth analysis that cannot be obtained through quantitative surveys. However, the small sample size necessarily limits the conclusions that can be drawn. We found that very few new themes emerged after 15 interviews, regardless of the stakeholder. Therefore, our selection of 20 experts appears to be sufficient to get a broad representation of possible harms and interventions. However, it does not provide a large enough sample to evaluate differences between stakeholders. In addition, there were some themes that emerged in the second half of the interview process that inspired additional questions to subsequent participants. We do not know how the earlier participants would have responded if asked directly about those themes.

Furthermore, it is difficult to elicit probabilities associated with risk when the chances of harms are extremely small, or extremely dependent on context. Therefore, we avoided asking for or performing quantitative evaluations of risk.

2.3 Harms and Concerns

In this section we describe the risks and harms identified by the experts. First, we describe the major themes that were identified. We then describe which harms were considered either likely or harmful.

2.3.1 Definition of Data Sharing

As a warm-up question, and to make sure that experts were using similar definitions, we began the interview by asking the experts to define "data sharing" from smartphones. Experts typically defined the term as data that is sent from the phone to any other party, including app developers, phone carriers, the OS or platform providers, and any third-parties with whom data is further shared.

Harm	Examples	# of Experts
Social problems & embarrassment	embarrassment, problems with social relations, spamming friends, social boundaries crossed (employer sees something they shouldn't), sensitive data being viewed by others, cyber-bullying	17
Direct financial harm	malware, thieves discover house location, id theft, premium texting	16
Surveillance & monitoring	government surveillance, location monitoring (whether or not physical harm/stalking results), activity monitoring	13
Privacy concerns	strangers/enemies find location, sensitive data being viewed by someone else, identified based on biometrics	13
Financial discrimination	price discrimination, job discrimination, insurance discrimination, redlining	11
Physical harm/stalking	strangers find location, stalking, physical harm due to location being known, harm due to knowledge about physical vulnerability	9
Behavioral advertising	unwanted marketing	8
Resource usage	spam, downloading unwanted software, battery drain	8
Health discrimination	medical insurance discrimination, discrimination based on disability	5
Harm to society	phone converted to botnet, filter bubble	4

Table 2.2: Themes for Harms, Risks, and Privacy Concerns, ordered by the number of experts that mentioned them.

2.3.2 Identifying Harms and Concerns

The goal of the first part of the interview was to brainstorm all the possible harms or concerns that could occur to a smartphone user. We then used follow-up questions to identify whether they consider these issues likely or harmful. Some experts expressed reservations with the word "harm." They were particularly concerned about whether this included only things could be proven

harmful in a court of law. Our goal was to open the field so that all possible concerns could be aired. We asked experts to consider not just "harms," but also concerns. One expert mentioned that users were already protected by laws (such as against identity theft). We asked him to discuss what was possible, assuming that a lawsuit or other action was less desirable than preventing the harm.

Using emergent coding, we identified several major themes to the harms or concerns that could occur to smartphone users as a result of data sharing. These themes are listed in Table 2.2. The examples are those specifically given by experts.

Experts' responses included a range of high-level themes, such as those in the left column of Table 2.2, or very specific examples of harms and how they are caused. Some themes overlapped with other themes, but the examples given for each justified treating them separately. For example, physical harm and stalking could also be related to surveillance and monitoring, in that stalking implies monitoring. However, there were significantly different examples in each group. Being monitored was described as a harm in itself, whether or not it leads to a physical attack. Physical attacks were a significant concern that could come not just from stalking but through other ways in which the data was shared.

2.3.3 Evaluating Risks of Harms

We asked experts to tell us which of the harms they identified was the "most harmful," and which was the "most likely." Many experts did so, but several experts expressed that this was difficult as it may depend on the context, the user, or a specific scenario. G1 expressed his concern as follows: "It ends up not being super helpful to talk about what's most likely and what's most dangerous because you don't know anybody's individual situation, and I think there's a wide diversity of situations out there and contexts in which that calculus might change." SE2 expressed a concern about quantifying the level of harm or risk: "One of the biggest risks in smart phones and data and big data is that we don't fully understand the implications of the data, so the harms are unquantified." Due to the difficulty in identifying whether harms are likely, these experts often identified causes of harms (such as being infected by malware or unexpected data sharing) rather than harms to the user (such as financial theft) when identifying which was the most likely harm.

Likely Harms

When asked to describe likely harms, experts included both the harms and causes of harms. We coded their responses into the harm themes described above. The following issues were identified by five or more experts as being likely:

- Infection by malware (10 experts)
- Unexpected or excessive data sharing (8 experts)
- Social problems and embarrassment (5 experts)

Infection by malware was identified by the most experts (10) as being likely. For example, L2 described malware with the following examples: "I think that there is a real risk that cyber criminals will find ways just as they try to phish today. Or for that matter, that foreign governments may try to compromise mobile devices and turn them into bot nets. Certain malware-based mischief is

probably the biggest risk." Many experts identified malware as leading to financial harms to the user. Some thought malware could also result in other types of harms, such as harm to society caused by bot nets (AC1) or resource use caused by spam (PL2).

Two experts said that although malware was currently not that frequent, they expected malware to increase in the future due to financial incentives. SE1 said, "As far as malware goes on phones, it's still a pretty small problem, especially compared to the PC malware. However, mobile devices are increasingly ubiquitous. So I don't think anybody's questioning that it's going to be a big problem in the future."

Eight experts identified unexpected or excessive data sharing as likely. This includes data sharing with apps or with third-parties. PI1 explained the high probability of data sharing, "The immediate threat to look at is the opposite of data minimization by apps right now, in terms of are they collecting only what they need for their particular process, or are they taking more data they're trying to find a secondary use for later." However, some experts said that although this was likely, it did not necessarily lead to a direct harm. AP1 said, "The most likely is the data sharing with others, but not necessarily leading to your identity being stolen."

Social problems and embarrassment were described as occurring either because of poor user interfaces (UI), or the user not being aware of the possible use or re-use of their data. L1 described the poor UI problem: "I do think people are inadvertently posting, sharing, having trouble with the UI... I'm making decisions to share or not share with UIs that arent always well designed, and so I may be over-sharing, either because of social network or just because of posting, tweeting, contacting, messaging."

Harms that could cause the most damage

The most damaging concerns identified by five or more experts were:

- Financial (12 experts)
- Physical (5 experts)
- Social Harms (5 experts)

Financial harms, typically resulting from direct financial theft, phishing, identify theft or malware, were identified the most frequently as harmful. PL1 expressed this concern, "I believe that the one that is most harmful is the direct theft of financial data because that has a direct financial impact on the user."

Physical harm – from stalking or from location being known – was also identified as harmful. AP2 described it, "The most harmful would be stalking, leading to ultimate dire consequences." G2 said, "Stalking isn't that likely but the damages are so great."

Social harms covered a range of social issues, from divorce to loss of job. Embarrassment also fell into this category. G2 expressed that this could fall within a range of very harmful to not harmful, "Embarrassment sounds like it should be low on the list but people do lose their jobs from information that's found out, and marriages break up and things based on information getting out that people didn't intend to get out." Some research has been done on embarrassment and regret on social networks such as Facebook or Twitter [161, 146, 147].

2.4 Interventions

We were interested in what could be done to prevent the harms and concerns described by experts. We identified three groups who could help protect the user: smartphone users themselves, app developers, and platform or OS developers. Several experts also described what the government or regulation could to to mitigate harms, and we specifically asked the government stakeholders about the role of regulation in mitigating harms. We describe the mitigation themes that were mentioned most frequently by experts.

2.4.1 Interventions by users

Five or more experts mentioned each of the following four ways smartphone users themselves could mitigate harms.

- Education (17 experts)
- Lighted Streets (15 experts)
- Protect Phone (7 experts)
- Reduce Functionality (5 experts)

Education: Experts suggested that smartphone users need to become better educated about a variety of topics including privacy settings, how location works, or the data ecosystem. Some experts felt users should understand app origin and behavior. AP1 said the entire ecosystem needed to be better understood: "No amount of improvement in logical interface, better icons, better information flow will prevent the problem, which is that people lack context. Therefore even something that fully notifies them — unless they understand its implications or what it means — it's still pointless."

Some experts emphasized that smartphone users needed to understand the risks behind different phone features. AP1 described the issue: "The larger question is that if you share, but by sharing you put yourself at risk because you shared too broadly, you didn't understand the full complexity of what you're sharing or how it's being shared, There exists some risks there."

Experts stressed that it was the users' responsibility to educate themselves. PL1 said, "You [the smartphone user] have to be smart about it and you have to know where the app is coming from and try to know as much as possible about what the app is doing."

Previous research found evidence of the need for smartphone user education. Mylonas et al. investigated users awareness of smartphone security, and how it impacts their decision-making about app downloads. They found that users who are not security-savvy, or who are unaware of smartphone malware are more likely both to trust app repositories and to store personal data in their phones [136]. Thus, the need for education for less aware smartphone users becomes especially important.

Play on the Lighted Streets: experts frequently mentioned that users needed to download only trusted apps or use only trusted app stores. We borrow the title of this theme from AC3, who said "The best thing I can tell you is to play on the lighted streets, and by that I mean that for the most part, the popular applications are safer because they receive more scrutiny." PL2 says, "The first thing is ... have some notion of which apps are trusted."

This advice typically requires that users download apps only from well-known brands or manufacturers. PL1 explained, "One of the biggest things that I always look for and I always encourage

my friends to look for is a trusted vendor. If that vendor or manufacturer misuses my data, what do they stand to lose?" AP2 echoed that a user should rely on well-known brands, such as, "large brand companies with reputational risks attached to their name. Usually publicly-traded companies, which are traded on the Stock Exchange, or companies with a brand name, are more likely to be responsive to consumers and therefore, easier to trust because press accounts or journalistic inquiries about their practices are likely to create more scrutiny of the company."

Other advice in this category also included that users should only download apps from major app stores. G1 stated,"I think the number one thing that she should do, is she should only download apps from marketplaces, and really she should only download apps from whatever the ... relevant OS marketplace is for her phone."

Experts also advised that users only download popular apps that have been downloaded many times before. APL1 said, "So when you're installing apps, try not to be like the first one to install an app." AP1 advised reading the app store reviews, "So be cognizant of the reviews, what the number of stars are. Simple things like that can at least help to some degree. Is it a solution? No Does it mitigate? Yes."

Protect Phone: Protecting the phone involves installing protective software or physically safeguarding the phone, typically against the phone being stolen or physically intercepted. AC1 said, "The types of precautions is to guard physical control of the phone and to think of it as just as sensitive as your computer, and that means you put a password on it and you don't leave it lying around for your suspicious father to search through."

Suggested software protections include: a remote finder in case the phone is lost, setting a secure phone password, using a password manager with encryption, using secure VPN, setting up a remote wipe, and backing up the phones data. However, G1 cautioned that this type of protection was not sufficient, "PINS and passwords are not going to be guaranteed security against a really determined criminal or guaranteed security against law enforcement when they're trying to access your device, but theyre like door locks. They keep honest people honest, and keeping honest people honest can be really helpful when many privacy risks come from people you know."

Reduce Functionality: Experts mentioned specific functionality that should be turned off in order to reduce data sharing. These included turning off Bluetooth (AC2 and SE2), location (AC2), using airplane mode (L1), network settings (SE2), and avoiding public Wi-Fi (SE2, AP2). Turning off these functions may limit the usability of the phone for certain apps or usages, but it also limits the data being sent, or limits when or where it is transmitted.

SE2 described how a smartphone user could protect herself: "She should take a look at the network settings and disable anything she doesn't use... For example, if she has no intention of using a Wi-Fi network, turn off Wi-Fi. There's no reason to have it on." AP2 said, "She probably should not use a public Wi-Fi network when sending or transmitting any sensitive information."

One example of this is that brick-and-mortar stores are currently using public Wi-fi to track their shoppers movements indoors, often without their knowledge or permission. The advice the popular press has given to those who wish to avoid this is to, "turn their phone off and take the battery out" [99].

Nothing: Four experts expressed concern that there wasn't much the user could do to prevent the harms discussed. While this was not a frequent theme, we mention this issue as an important concern. G1 said, "In terms of mitigating the risk, reducing the risk, attempting to prevent the risk ... there's a bunch of stuff she can do. But in terms of actually outright preventing the risk, get rid

of the smartphone." CA2 expressed concerns about the data sharing ecosystem, "It's going to be very hard to escape the system. Almost impossible." All four of these experts did suggest at least one of the interventions mentioned above in addition to expressing skepticism that real protection was possible.

2.4.2 Interventions by App Developers

Five or more experts identified each of the following ways app developers could mitigate harms.

- Transparency (15 experts)
- Best Security Practices (8 experts)
- Priority (7 experts)
- Data Minimization (6 experts)
- Understand APIs (5 experts)
- Customer Relationship (5 experts)

Transparency: Fifteen experts mentioned that app developers needed to be more transparent about their data sharing practices through disclosures to users. This included disclosing the purpose of the data collection. G2 said app developers should disclose, "what data they're collecting for themselves and what data they're planning on collecting and sharing for other purposes." This was echoed by AP2 and PL10, who were concerned about what data was being collected, why it was being collected, and with whom it was being shared.

SE2 felt transparency would address privacy concerns by removing surprises, "I have a hunch that the majority of people's concern about privacy on the web and on smartphones has to do with the lack of transparency. They just don't know what's going on, so when they find out it's a surprise 'cause they assumed it wasn't."

The CA Attorney General has also been addressed this mitigation in the recommendation for app developers to, "Develop a privacy policy that is clear, accurate, and conspicuously accessible to users and potential users" [91].

Ten experts also mentioned concerns about the efficacy of transparency. They felt transparency would not be effective if users were not interested in nor had the time to learn about or read privacy policies. When discussing privacy notices, AD2 stated that many consumers would reject notices: "They don't want a bunch of disclosures and notices and stuff that either they don't understand or they don't particularly care about. I think there is a percentage of people who do care a lot about that [...] [b]ut I think the majority of people don't want to click through all that stuff." G1 had a similar statement: "there is always going to be either a majority or a super majority of these folks who simply aren't going to read the stuff and aren't going to take the time to compare anyway." AC2 stated a similar concern about balancing the right amount of information on small screens with consumers limited attention, "This is a tension for us, as user interface designers too, which is we can put in a lot of nuance in terms of what's going on, but how much will people actually read?" CA1 put it bluntly, "I don't know what [a link to a privacy policy] gets you because no one reads the damn things."

Best Security Practices: Eight experts said app developers should be following known best security practices. SE2 posited that secure code was the foundation for protecting users' privacy and security: "If you don't have a secure application you can't guarantee privacy at all. It's

impossible. If I write the most awesome privacy preserving software... there's a bug in my code and somebody can exploit my software and make it eavesdropable, ...that's completely useless. So I think building secure code is the foundation of all privacy and data control." Examples of secure code given by experts include using SSL and proper encryption of data.

There are resources for app developers on developing best security practices. These include guidelines on mobile web from a standards consortium [7], and guidelines for each platform from the platform developers [1].

AP1, an app developer, recommended that app developers create a privacy policy as part of best practices, stating, "In fact, the generation or creation of a privacy policy is something that often leads to more insight about your product. The developer might view the creation of a privacy policy as something for his customers, but in fact its real value is for himself and his developer team, or herself and her developer team."

Priority: Seven experts said that app developers may not make privacy and security a priority, but they should. Some explained this as a lack of resources. For example, PL1 said, "[App developers] are working with very few resources and they're trying to develop complex applications in very short time frames in order to try to make some money."

AC1 had a similar explanation, "I think often privacy and security is one of the things they plan to do. It's on a later day, and their primary or their first order of concern is to get a running app that does something valuable, and they're gonna think about things like privacy and security much later, perhaps when they make money or otherwise are more successful."

In Chapter 3, we discuss our findings that app developers did not prioritize privacy and security, and in particular that smaller app companies were less likely to exhibit privacy and security best practices.

Data Minimization: Five stakeholders, including both of the app developer stakeholders, discussed the need to minimize the data that was collected. SE2, a lawyer, put it succinctly, "Don't ask for privileges you don't need. It's a liability." AP1, an app developer, said, "If the developer or the application or the carrier isn't collecting the information, then no potential security risk exists if the information is leaked because there's nothing to leak. That often is referred to as data minimization. Data minimization is a way to use privacy to try to enhance security. If I don't have it, I can't leak it."

Understand Third-Party Libraries: Five experts discussed the fact that app developers often use third-party libraries, Application Programmer Interfaces (APIs), and code toolkits. App developers should reveal to users what data is collected by these third parties, but they often do not, in part because developers themselves may not know what information is being collected by this code. G1 said, "App developers really need to be aware. And some app developers weren't really good at this and some don't give it a second thought.... When they are using widgets or modular pieces of code or third-party services in order to provide portions of their app, that they need to pass on those disclosures concerning those which are modular pieces of code and third-party services to the users."

[1] Guidelines for iOS developers are available at from Apple at "Introduction To Secure Coding Guide" **https://developer.apple.com/library/ios/documentation/Security/Conceptual/SecureCodingGuide/Introduction.html**
Guidelines for Android developers are available at "Best Practices For Security & Privacy" **http://developer.android.com/training/best-security.html**

L2 felt that third-parties had a responsibility to disclose their data collection practices. "It would be good for third parties who collect a lot of information through apps to put up a kind of standardized notice so information can be sent along and populate the little short privacy notice that ideally the mobile apps will provide." Other experts thought that app developers maintained responsibility for understanding third-party code. In Chapter 3, we describe further evidence that app developers did not always read or understand the terms of service or privacy policies of the companies or tools they used.

Customer Relationship: Five experts mentioned that app developers need to understand the role of privacy and security in their customer relationships. Often, this was tied to the need for transparency. PL1 said, "App developers need to understand that users will partly choose to use their app or not use their app based on whether they trust them. And they need to make sure that they do the right things in order for users to trust them. So a lot of that, again, comes down to being upfront with the user in terms of what an application is doing and why." G2 stated that app developers, unlike the platform developers, have a more direct relationship with the user and therefore had increased responsibilities to be transparent, "because app developers have a direct relationship with users. They need to be able to utilize that direct relationship. So, especially when they're dealing on the sensitive data, financial apps, kids' apps."

AP1 emphasized that this is more an issue of trust than privacy: "Trustworthiness is a category under brand, and so having an educated populace that sees your product as more trustworthy because it provides either better control or limits stuff. You never ever, ever will ever make money selling privacy."

The research described in Chapter 3 was informed by the experts' suggested mitigations, in that we specifically examined opportunities for app developers to mitigate the risks, or the hurdles to implementing better privacy and security.

2.4.3 Interventions by Platform Developers

Several themes emerged when we asked experts what platform developers or OS providers should be doing to protect the smartphone user from harm. Typically, platform or OS developers are also app store providers.

- Transparency (17 experts)
- Improve UI Control (10 experts)
- Security Improvements (7 experts)
- Work with App Developers (6 experts)

Transparency: Seventeen experts argued for more transparency about data sharing from platform developers. Three experts were concerned about location sharing (AC1, L1 and G2). L1 suggested that location should have a notification every time it was shared. L1 was also concerned about sharing the phone's unique id.

Several experts emphasized that telling users what permissions were being used was not sufficient, but they also needed to know why, how often, and where data was being shared. AP2 said, "There should be a clear statement about what's collected, why it's collected, and who it's shared with... And I think there should be clear benefit statements about what benefit the consumer receives, paired up with the what, why, and with whom." AC2 said, "It'd be nice if it could just

sort of categorize, 'We're using your data for this reason.' It'd also be nice if it could say how often it's doing it too, like is it doing it automatically or is it doing it in the background every five minutes or so on? So these are things that right now there's no easy way of trying to determine." In Chapter 5, we examine a notification that tells users how often data is collected.

Suggestions for improving transparency in the user interface included the app store and the operating system itself. Two experts suggested just-in-time notifications (AP1 and CA1). AC3 also discussed the need to develop notifications for smaller mobile devices. "You can't just take what kind of works from the desktop, throw it on to the platform with totally different visual characteristics and pretend that it's going to work. It doesn't even work for expert users. What hope do regular users have?"

However, many experts expressed doubt about whether users want additional notices, and whether users would be willing to read or learn enough to understand them. P19 said, "It's probably a good practice to allow consumers access to that information when they want it but not in a way that undermines the experience of the app itself. People come to the app to use the app, not to read about a bunch of information practices that really won't impact them."

Improve UI Control: Ten experts suggested that platforms should improve the privacy and security controls in the user interface (UI). Two experts said that both a simple set of controls and a more fine-grained set of controls should be available. SE2 explained, "The OS vendors have to be really careful to provide for the people that don't wanna think about it by securing things as best they can by default. Then provide the cues for people who want to dig into it a little bit and maybe make a more informed risk decision."

Three experts said they would like to see a "do not track" setting implemented, or that they believed it would be implemented. One was concerned that do not track was still not defined.

Security improvements: Seven experts said that platforms should also implement security best practices. These best practices were considered to be well known, but some experts offered specific advice. SE2 said, "OS developers can protect their software that they're building and figure out what people want for security and privacy and do that by default." G1 suggested remote wiping:"The remote wiping is a use case that happens all the time, and it's something that third-party apps provide." AC1 mentioned the need to push out security updates: "So many Android phones are in an insecure state because the carriers aren't pushing out OS updates."

Eight experts mentioned that finding and removing malware from app stores was an important part of platforms' role in protecting users. For example, AC2 said, "One thing they could do is try to find this malware faster or do better testing and all, and I know they are trying to do that too." AC3 said, "All of the markets, the major markets, are pretty vigilant in keeping the absolute worst stuff out and they have strong financial incentives to do that."

Work with App Developers: Six experts said that platforms should provide app developers with tools and education to enable improved privacy and security. Stakeholders said that solutions could include more example code for security (AC3), making it easier to implement security features such as SSL (AC2), better toolkits (AC2), a security checklist (AC2), and enabling app developers to be transparent (SE2) by giving them tools to let users know about apps' data requirements.

The platform industry stakeholders agreed with this solution. PL1 said, "I think one thing that we can do as an industry is make it easier for developers to secure their applications and give them tools and libraries to do that, because if we expect developers to put in the time and the effort necessary in order to create their own security, they're often going to mess it up. Or in most

cases, honestly, they just won't do it at all because they don't have the time to and they don't have the incentives to." PL2 said, "Every app developer should be doing privacy by design. But I think realistically, the other players have more resources to raise privacy awareness... than the app developers."

AP2 said that responsibility needed to be shifted from just the app developers to the platforms as well, "Shared responsibility would lessen the burdens on apps, and actually would assign responsibility for developing private tools and notices, and helping educate consumers, and helping consumers achieve the goals that they set out to when they use the phone. [This is] better than leaving all the responsibilities for app developers themselves. And so I consider it a systemic failure that we're all experiencing right now."

Nothing: Only one participant thought platforms should do nothing more. AD2 said "I think actually in some ways that they're more restrictive than they could be with respect to data sharing in a way that undermines competition and probably limits offerings to consumers."

2.4.4 Role of Government

Seven experts said the government can aid in mitigating risks and harms, and five of the seven brought up government intervention without being explicitly asked about government's role. However, not all experts were asked about the role of the government; this theme emerged naturally through the interviews. The interviewer specifically asked the government stakeholders what they perceived as governments role in mitigating risk, but other interviewees were not asked.

The two government representatives, G1 and G2, thought that any policy should be sensitive to company needs and innovation. G2 said that the government should "promote ways for companies to work together to come up with good practices. It helps the marketplace in general and trying to convince companies that working together to do that with government is gonna be more successful than just hearing about these cases of the bad actors." G1 was concerned about how to set up regulation that did not stifle innovation, "So government has to walk a really fine line as it does in many areas in terms of technology between imposing responsibilities to ensure that consumers are protected, while at the same time promoting innovation in the space."

Suggestions from other stakeholders included new regulation, promoting best practices, intervening when companies do not meet security best practices, working internationally, and developing standardized notices. CA2, a consumer advocate and participant in the NTIA process, described the NTIA MSH goal: "To develop a code of conduct or mobile apps to cover so-called transparency. Which is a very limited approach and only is one." CA2 also felt that, "What's needed is [for] the FTC to promulgate regulations and legislation passed by Congress to empower users to opt in to all this data collection and use."

One stakeholder (TE1) expressed concern that the legislative process did not allow the time or communication needed to understand the technical details and create a quality standard. This stakeholder emphasized that self-regulation efforts within industry allowed the companies to "get technical input and to really get into the nitty-gritty of the words and what they mean in a way that's impossible in a legislative environment. You know, [in a legislative environment] you might have one meeting with the bill sponsor that you can maybe make one point. You can't wordsmith a document... So its just unlikely to be timely and effective when its done through legislation."

Several stakeholders felt that policy should not be focused on app developers but on other stakeholders, such as data brokers, platforms, and app stores. AP2 wanted to see limitations on

data and collection and use by data brokers and advertisers, saying, "while we spend a lot of time publicly debating what apps should be doing, we've spent almost no time discussing and debating limitations on those other entities and their data usage." AC1 specifically supported the California Attorney Generals approach in attempting to "police the elephants – the carriers – rather than all the little mice who are making these apps."

Many of the experts we interviewed were at the time of the interview participating in a government led process, the NTIA MSHP on mobile transparency. Therefore their views are probably somewhat reflective of their opinions on that process. There is likely a correlation between participating in a government process on mobile transparency and believing that government processes can be useful but difficult.

2.4.5 Vulnerable Populations

In order to understand whether specific interventions are needed for different groups, we asked experts what populations are most vulnerable to harms from smartphone data sharing. We also asked experts whether any of the harms discussed were different for children. Most agreed that harms were different for children. As PL1 said, "Adults are generally more aware of the long-term implications of their actions, whereas children don't necessarily have that same level of awareness." AC1 and AD1 said that teenagers might be more sensitive to social embarrassment or bullying. PL2 said that children might not be as vulnerable to financial exploitation since they have fewer financial resources. Obtaining parental consent was discussed as a difficulty for apps. Several experts mentioned that parents were also vulnerable to mistakes made by their children while using their parents phone.

We also asked if there were any vulnerable populations besides children. Thirteen experts cited the elderly, but some noted that not all elderly are vulnerable. Experts mentioned that some elderly may be less technologically savvy, may have trouble seeing small screens, or may have trouble manipulating small devices. Notice and choice interventions may need to take into account such needs.

Other vulnerable populations mentioned were: battered women, mentally or emotionally disabled, visually disabled, those living in countries without due process, members of the military who would be at greater risk if their location was revealed, those in financial situations where they can't purchase apps without advertising or technological protections, groups that were not previously exposed to PC technology, those with language barriers (e.g. non-English speaking in the US), and minorities who could be unfairly targeted for unhealthy or undesirable products.

2.4.6 Interventions and the Privacy Principles

We looked at how the interventions discussed by experts related to the FTC's Fair Information Privacy Practices [3]. Several interventions are related to the notice and choice principles. Security was also frequently mentioned as an intervention. The other two principles — participation and enforcement — were not frequently mentioned by experts. In addition to notice, choice, and security, data minimization was also mentioned. Data minimization is not part of the FTC's principles, although it is part of other sets of fair information principles, including the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data [2]. The security principle described by the FTC focuses on preventing unauthorized parties from accessing data, while data

minimization includes reducing the amount of data collected, even by authorized parties. Some experts stated that data minimization is a part of security best practices. We highlight data minimization as a separate category because experts found it to be an important intervention.

Principle	**Interventions**		
	by users	**by app developers**	**by platform developers**
Notice	education	transparency, customer relationship	transparency
Control			improve control
Security	protect phone, play in the lighted streets	best security practices, prioritize security and privacy, understand APIs	security, work with app developers
Data Minimization	reduce functionality	data minimization, understand APIs	

Table 2.3: Interventions grouped by who is responsible and whether it will lead to improved notice, control, security, or data minimization.

We show how each intervention relates to a privacy principle in Table 2.3. We further explain the categorizations below. Some inventions fall into several categories. For example, app developers making privacy and security a priority could lead to improved notice, control, and data minimization. Making privacy and security a priority is a prerequisite for most of the other interventions, and recognizing the importance of customer relationships is an incentive to improve privacy and security.

Notice: Notice about data sharing can be improved through transparency, education, or standard notices. Notice may lead to improved user decision-making. To educate themselves, some users pay attention to notices and other available information. Therefore, they can make more informed choices about the data they share. App developers and platforms have a role in improving transparency and providing better notices.

Notices should go beyond stating what data was shared, and should include purpose and secondary uses, and also take into account third-party libraries. Several experts said notices about data collection should include why data is collected and with whom it will be shared. Notices providing this information can help to reduce surprise from unexpected data sharing and embarrassment caused by data sharing, while also helping users understand why some data uses are necessary. Including notices about the data practices of third-party libraries will help insure that the notices provided by apps are complete. Platform providers should provide tools that will assist in conveying notices in a standardized format.

Notice can include user education about malware. The greatest risks come to users from malware. Some experts indicated that users should use only apps from well-known companies, but this may discriminate against legitimate but less well-known companies. Smaller app companies may desire a way to indicate trustworthiness. While the major platforms are taking steps to scan their app markets for malware, there may be room for a notice or indication that an app has been scanned and can be trusted.

Control: Control over data sharing can be improved by making existing controls more usable and by adding additional controls. While control does not necessarily imply notice – it is possible to add control through new interfaces that users don't see or don't understand – we assume that control options would be well implemented and usable, and that the user understands the control mechanisms.

Platforms could provide better control, allowing users to make decisions according to their privacy and security preferences. This could mitigate privacy concerns, stop location monitoring (which can lead to stalking or physical harm), allow users to turn off behavioral advertising, and control resource usage.

Security: Two user interventions may improve security: 1) users can protect the phone so that others will have less access to it or the information within, and 2) "playing on the lighted streets" could result in fewer malware downloads, and data shared with fewer malicious third-parties. Many of the app developer mitigations, including best security practices, data minimization, and understanding APIs, can improve security. This in turn can reduce the risks of data reaching unintended audiences, reducing harms such as physical harm, surveillance, financial harm, and social problems. Platforms can work with app developers to improve security practices. Security in the platform or app store could lead to less malware, fewer data breaches, less unencrypted data transmittal, and fewer coding mistakes that allow unintended transmission of data. Therefore, security can mitigate the most harmful risks.

Data Minimization: Data minimization requires sending, collecting, and storing the minimum amount of information that is needed. When users choose to reduce functionality, they will trade some usability or functionality in order to share less information, which minimizes the data shared. Data minimization also addresses many of the same harms as security. If there is less data to transmit and protect, there is less chance of unauthorized access. If the authorized data collector's purpose in collecting data is to profile and to make decisions about the consumer, financial and health discrimination may result. Reducing collected data may help prevent the discrimination, as there is less information to create profiles.

2.5 Linking the Mitigations to the Harms

In this section, we analyze the relationship of the risks and harms to each of the FIPPs discussed above. We classify how the harms identified by the experts can be mitigated through notice, control, security, or data minimization. As Table 2.4 shows, most harms are not mitigated through notice or control alone, but require security and data minimization. We explain these classifications and provide examples.

The harms that can be mitigated by notice alone – social problems and embarrassment, or privacy – tend to be highly personal. In these cases, users can reduce the risk of harm by changing their behaviors, such as not installing an app, or not posting information. In these cases, notice of the data sharing or collection may suffice and may be the only appropriate mechanism.

Other harms can be mitigated if users have both notice and control over the data collection or sharing. In these cases, notice is an important pre-condition for control, but the control itself allows the user to make the decision that mitigates the risk. For example, users may be able to specify with whom data is shared through a control. We illustrate this through the scenario of a woman who is concerned that an abusive ex-partner will stalk her if he has access to her location

Harm	Notice	Control	Security	Data Minimization
Social problems & embarrassment	yes	yes	yes	yes
Privacy concerns	yes	yes	yes	yes
Behavioral advertising	yes	yes		yes
Surveillance & Monitoring	yes	yes	yes	yes
Physical harm		yes	yes	yes
Stalking		yes	yes	yes
Harm to society			yes	yes
Direct financial harm			yes	yes
Financial discrimination				yes
Resource usage				yes
Health discrimination				yes

Table 2.4: Our proposal for whether all harms are addressed by User Notice, User Control, Security, or Data Minimization.

information. She could share her location with friends who might be concerned about her, but could disallow the abusive ex-partner from accessing the location information. In this example, control allows the user to apply their personal or situational information to mitigate harms, while still taking advantage of the benefits of information sharing.

Some risks cannot be addressed by notice and control. In these cases, there may be a malicious party, or a party motivated to act against the interest of the user. The malicious parties may circumvent notice and choice, deliberately hiding their access to information or preventing the user to control the data sharing. For example, a party that is interested in causing financial harm to the user will try to do so in a way that the user cannot control through notice or control. Similarly, discrimination is not likely to be an explicit option from which users can opt-out. In these cases, users are not able to mitigate the harm. As most app developers are not trained in security and privacy (discussed further in Chapter 3) and are thus unlikely to put much effort in implementing security or privacy features, platform developers are probably in a better position to protect users as long as this protection does not encourage users to switch to other platforms (as it may happen if users felt annoyed by added, unusable privacy or security interfaces). Hence, these harms should be mitigated through security (not allowing malicious users to get access to the information by protecting the data with encryption), or data minimization (not creating or storing data).

We classify discrimination as being mitigated only by data minimization. We assume the party collecting the data is motivated by interests that contradict the smartphone owners' best interests, resulting in a possible harm. Price discrimination may result in corporate profit but consumer loss. Security won't protect the user, as the party that causes harm may have unmitigated access to information, regardless of best security practices protecting the data from access from other parties. Price or product discrimination is likely to be perceived as harmful even by users who wish to receive targeted ads. For example, in September 2000 an Amazon's customer discovered that if he removed the cookies in his computer, he obtained consistently a lower price for a DVD that was offered to him.[2]

[2]CNN Law Center, "Web sites change prices based on customers habits", **http://edition.cnn.com/**

If users are informed about behavioral advertising, they can use various tools to opt-out of advertising.[3] In those cases in which users do not want to receive targeted ads and are not given a mechanism to opt out of data collection, data minimization is the remaining mitigation option.

To summarize, notice alone may help mitigate harms only in situations in which users have control. Notice and control can be helpful to mitigate harms in the cases when there is not a malicious party who is motivated to circumvent notice and choice.

2.6 Discussion

By interviewing experts from many stakeholder groups, we were able to get a holistic perspective of the harms and concerns to users from smartphone data sharing. A number of harms from smartphone data sharing were identified by the experts. These harms included tangible and direct harms such as financial harm and physical harms. They also included less direct harms such as behavioral advertising and embarrassment. In order for users to continue enjoying the benefits of smartphones, it is best to mitigate the risks of the harms. Experts identified a number of mitigations that users, platform developers, app developers, and regulators could implement. We classified these mitigations as providing either notice, control, security, or data minimization. We find that many interventions are related to improved security and data minimization.

The interventions that improve security and minimize data collection mitigate the most harmful risks. The damaging financial, physical, and social harms cannot be addressed by notice and transparency alone. Therefore, we encourage app developers, platforms, and policy-makers to enlarge their efforts to include best security practices and data minimization. Work is currently being done to improve notice and transparency about data collection by apps, for example, in the NTIA MSHP. Similar work could be done to create a code of conduct addressing security and data minimization for apps and platforms. While notice and user education is a precondition for better control and better decisions by users, the focus of future policy efforts should include improved security and data minimization.

We offer two suggestions for improving risk communication and notices to users, based on the harms and interventions discussed by experts. Both suggestions are also directions for future research.

First, notices should go beyond stating what data was shared. Notices should include purpose and secondary uses, and also inform about third-party libraries. Several experts said notices about data collection should include why data as collected and with whom it will be shared. Notices providing this information can help to reduce surprise from unexpected data sharing and embarrassment caused by data sharing, while also helping users understand why some data uses are necessary. Including notices about the data practices of third-party libraries will help insure that the notices provided by apps are complete. Platform providers should provide tools that will assist in conveying notices in a standardized format.

Our second suggestion for improved risk communication is to provide guidance to users about malware. The greatest risks come to users from malware. Some experts indicated that users should

2005/LAW/06/24/ramasastry.website.prices/.RetrievedonApril/3/2014

[3]Some work has found that these tools could be more usable and clear to users. See [115] Other work has investigated whether the existing notifications are effective, finding that the current AdChoices icon could be improved. See [117]

use only apps from well-known companies, but this may discriminate against legitimate but less well-known companies. Smaller app companies may desire a way to indicate trustworthiness. While the major platforms are taking steps to scan their app markets for malware, there may be room for a notice or indication that an app has been scanned and can be trusted.

Acknowledgements

We thank all the experts who participated in this project by volunteering their time and knowledge.

Chapter 3

The Privacy and Security Behaviors of Smartphone App Developers

While research has looked at smartphone users' perceptions and needs for privacy and security, there has been a dearth of work about the perspectives of app developers Apps are developed by a broad array of companies and individuals. As the space for innovation is huge, and the barrier to entry is low, many small to medium size app development companies have been able to publish apps. Over 200,000 active developers contribute to the Apple store [9]. There is no training or certification process for app development designed to protect the client. Furthermore, app developers may feel pressure to develop quickly and be the first to market. In the race to innovate, privacy and security might not be the top priority for time- and resource-constrained app developers.

In this paper, we examine the ways app developers make decisions and the steps they take to protect security and privacy. Through in-depth interviews with 13 developers, we explored the trade-offs app developers make, knowledge acquisition, and barriers to implementing privacy and security best practices. Informed by the results of these interviews, we formulated several hypothesis about the privacy and security behaviors of app developers. We ran an online survey of 228 app developers to examine factors that predict good privacy and security behaviors, such as encrypting data and providing privacy policies. Our two-step research process is similar to that used in other work examining human subjects' motivations [103].

We first begin by discussing previous work on smartphones and privacy. Then, we describe the interviews and the themes that emerged. In the following section, we describe the on-line survey and the results of testing specific hypotheses about privacy and security. We find that many developers lack awareness about privacy, and we dentify a number of barriers to improved privacy and security behaviors. These include the lack of resources in smaller companies and the difficulty of understanding third-party collection of user data. We identify where developers seek privacy and security advice, and discuss intervention points and improved tools to help developers.

This chapter is largely based on a paper [36] co-authored with Abigail Marsh., Jialiu Lin, Jason Hong, and Lorrie Cranor, L and on a column [34] co-authored with Lorrie Cranor

3.1 Related Work

We first describe the smartphone app ecosystem, including major platforms and how apps are submitted. We also discuss users' perceptions of smartphone privacy and security. We then describe public policy efforts to guide app developers when making privacy and security decisions and previous efforts to inform app developers about privacy and security.

3.1.1 App Development Ecosystem

The two most popular smartphone platforms are Apple's iOS and Google's Android, with Blackberry and Microsoft holding a smaller market share. Apple and Google both have app markets that allow independent developers to distribute or sell apps, which users can download from their devices. This has allowed many independent developers to sell smartphone software directly to users, and has resulted in a huge variety of apps, with over 800,000 apps on each of the iOS and Android platforms as of October 2013 [128].

Previous work has found a relationship between data collection and advertising as a revenue model. The ad-based revenue model, which often relies on targeted ads, is currently popular [118]. Apps may provide ads through third-party code, such as that provided by Flurry[1] or Google AdSense.[2] Targeted advertising requires collecting information about users, and therefore the targeted advertising revenue model may require more permissions and therefore be more privacy-invasive [51, 123]. Apps may also include third-party code for analytics, whose primary goal is to collect information about the users' interactions with the app.

Some previous work has examined app developer security behaviors, such as that by Egele et al. [67] and Fahl et al. [73], which found the significant portions of apps with security failures or substandard implementations of security code. Throughout our work, we explore app developers' perceptions of their work, including self-reported intentions. Several researchers have argued that privacy is not the only cost to users of ad-supported aps. Both network traffic and energy usage are increased by additional ads. Vallina-Rodriguez et al. examined the impact of ad networks on mobile apps. They found them to be prevalent across many apps on Android, but also on iOS as well. They also found that ad traffic was a significant fraction of traffic, and these ad communications consumed energy [159]. Zhang et al. looked at the amount of traffic generated by ads and analytic components of apps, and specifically compared free apps to their ad-supported versions. They found that, "ads contributed 87% of the total overhead traffic" [164].

3.1.2 User Concerns about Privacy and Security

A wealth of previous work has examined users' perceptions and desires for smartphone privacy and security [43, 59, 75, 136, 77]. Users are often surprised by what permissions are requested by apps [122], the frequency of data collection, and the data recipients (see Chapter 5). Furthermore, they often do not understand existing privacy notices, particularly in Android phones [107, 76]. While users are concerned about privacy and security, they are neither informed nor empowered to protect themselves. Therefore, the decisions made by app developers have great impact.

[1] `www.flurry.com`
[2] `www.google.com/adsense/`

Previous work has examined users' reactions to privacy policies. While privacy policies offer the illusion of notice to users, the reality is that the required time [130], reading level [101], and vague language [141] pose significant usability barriers. Our work indicates that app developers have similar troubles with privacy policies.

3.1.3 Public Policy and Tools

There have been several efforts to educate app developers about privacy and security. We reviewed five privacy guidelines for app developers: three were published by government agencies in Australia [18], Canada [16], and California [91]; one by an industry consortium in Europe [15]; and one collaboration by two consumer privacy advocacy groups [10]. These guidelines typically offered clear and readable advice and avoided "legalese." While they were lengthy (14-32 pages), some offered privacy and security checklists for developers. These guidelines often suggest that privacy policies can help developers think through their data collection practices in addition to notifying users.

There were five recommendations made by all of the above-cited guidelines, which we paraphrase as follows:

1. Someone must be responsible for privacy.
2. The app should have a clear and easy to find privacy policy.
3. The app should encrypt data during transmission.
4. The app should encrypt data it stores.
5. The app should limit data collection to what is needed.

These are the five main privacy and security behaviors we explored quantitatively in our online survey. We describe them in greater detail in Section 3.4.

Our work focuses on app developers in the United States, so we briefly discuss US public policy efforts to address smartphone privacy. In 2012, California passed the California Online Privacy Protection Act, requiring all mobile apps to have a privacy policy [134]. In 2013 the United States Federal Trade Commission published a report based on a workshop regarding mobile app transparency [19], recommending that app developers have privacy policies and notifications at the time of data sharing. In July of 2013, the National Telecommunications and Information Administration completed a multi-stakeholder process that released a voluntary Code of Conduct for mobile app privacy "short-forms" [20]. This Code of Conduct specifies standard elements that mobile apps should include in a short privacy notice. Some apps fall under additional privacy regulation, such as the Health Insurance Portability and Accountability Act (HIPAA), [3] which deals with health information, or the Children's Online Privacy Protection Act (COPPA), [4] which regulates the collection of information about children 12 and under.

Tools have been developed to help developers practice privacy and security behaviors. Many open-source databases, such as mySQL, allow encryption of stored data. Several free or low cost privacy policy generators[5] exist that allow developers to create a policy by answering questions about their app's behaviors. Our interviews examined whether developers were aware of or used these tools.

[3] `http://www.hhs.gov/ocr/privacy/index.html`
[4] `http://www.coppa.org/`
[5] `freeprivacypolicy.com`, `generateprivacypolicy.com`, `appprivacy.net`

3.2 Interview Method

We conducted semi-structured interviews with 13 smartphone app developers in August and September of 2013. Our research goals were to understand what decisions app developers make that they consider privacy and security related, and to better understand what resources they were aware of to help them make those decisions.

Interviewees represented a variety of app types and company sizes, as shown in Table I. We asked "What type of service does your app provide," and offered choices based on a taxonomy developed by Hyrynsalmi et al. [96]. Interviews lasted approximately one hour. The interviews were usually conducted remotely, with only one in-person interview. The audio was recorded for transcription, although participants had the option to refuse audio recording, as some said it made them uncomfortable or unlikely to be forthcoming. Interviewees received $20 as compensation. Our interviewees were overwhelmingly male, which is in-line with evidence that 94% of app developers are male [65].

ID	Company Size	Revenue Model	Service	State
P1	10-30	Advertising, Free trial, Subscription	Digital, Physical, Service, Contents	CA
P2	2-9	Advertising, Free trial, Other	Digital, Service, Contents, Advertisement, Personalized information, Other	CA
P3	2-9	Free trial, Other	Digital, Service	PA
P4	2-9	Pay-per-user	Physical, Service	WA
P5	2-9	Free trial	Digital	WA
P6	100+	Subscription	Other	PA
P7	1	None	Contents	TX
P8	10-30	Subscription	Digital, Service	CA
P9	2-9	Other	Service	CA
P10	1	None	Contents	PA
P11	2-9	Advertising, None	Physical, Personalized information, Other	IL
P12	2-9	None	Personalized information	PA
P13	100+	None	Physical	MI

Table 3.1: Interview participant mobile app and company demographics.

We recruited participants for interviews through a number of methods, including in-person recruiting at local meetups for smartphone app developers, online postings on sites such as Craigslist and Backpage, and through our social networks. Recruitment text said, "Participate in an interview to understand and improve smartphone app development." Security and privacy were not mentioned in the recruitment to avoid participant bias. We asked interested parties to first fill out a screening survey to see if they qualified. We included two technical questions to determine whether the applicant had credible knowledge of app development. Valid applicants were invited by email to set up an interview time with one of two researchers. We contacted 20 developers, and 13 completed the interview. Five of the invited developers who did not complete the interview

Service	Examples
Digital	games, MP3, Ebooks
Physical	selling books
Service	e-mail, banking, ticketing
Stock Information	stock prices
Contents	news, weather, entertainment
Personalized information	location information

Table 3.2: Service categories based on classifications by Hyrynsalmi et al. [96].

failed to respond to the email invitation, and two invitees were unable to find a suitable interview time.

We did not collect identifying information, such as given name or company name, from participants unless it was volunteered. The interviewed developers ranged from 26 to 58 years old, and were from six states. Most worked in groups of 2-9 developers, but company size ranged from 1 to 100+ employees. Most interviewees were programmers, but one was a product manager. Several interviewees played multiple roles in their company, such as CEO, manager, or quality assurance. Their apps represented a variety of business models and services, and were at various stages of maturity. Some apps were not yet released to the app market, and others had already had several versions on the app market.

Questions included, "What, if any, online resources do you use to help make privacy and security decisions?" and "Have you ever decided not to collect certain information from users due to privacy concerns?" While we generally followed a script, we iterated on the script as each interview informed the next. Participants were asked what subjects we should have addressed, which revealed gaps in our questions and allowed us to improve the interviews.

3.3 Interview Results

We describe the themes that emerged from our interviews. We discuss how app developers learn about privacy and security, whether they are aware of regulation and third-party data collection, and where they seek advice and resources for privacy and security decisions. We discuss developers' perceptions of privacy policies and the trade-offs that app developers confront when making privacy and security decisions.

3.3.1 Education and Advice about Privacy and Security

Only a few of the developers we interviewed had formal training on privacy and security, typically received through corporate training or certification. Other developers rely on online research to find answers to specific questions. They are not accessing the guidelines published by government agencies, and instead are more likely to rely on their social networks, or specialists within their companies for information.

Many participants did not have formal privacy and security training. This suggests that many developers learn about security and privacy when they are confronted with these issues in the course of their work, at which point they may seek out further education. The lack of education on security and privacy available at the introductory levels was not lost on developers. P3 stated, "Most classes in computer science...there isn't much of a focus on security. That could have a very big impact on how this stuff [implementation of secure code] happens." On the other hand, some participants were confident that they were learning what they needed to know, or had a good background. P13 said "I have no formal training with privacy and security, but I feel that I am a journeyman in privacy knowledge, and pretty expert at security knowledge." Similarly, P10 stated that his privacy and security learning, "is pretty much internal knowledge based on my experience in Web."

Some participants discussed receiving formal training from a variety of sources. Certain businesses have specific training or certification requirements. For example Payment Card Industry (PCI) has security standards for handling credit card information. P11 states, "When you work at E-Commerce, they want you to be what they call PCI compliant." In less regulated areas, participants reported education including certifications, previous work experience, and conferences such as the RSA Conference.

When asked about current and upcoming privacy and security regulations, participants showed little knowledge. While a few app developers brought up issues of the government requesting user data as a concern, none were aware of guidelines such as those discussed in Section 6.2. The exceptions were apps that were marketed to children under 13 or used health information; these developers were aware of the privacy laws specifically related to their cases.

Participants were asked to discuss what resources they used when they needed advice on security- and privacy-related decisions. We received a variety of responses, which could be grouped into a number of common themes, including searching online, consulting friends, and seeking legal or specialist advice.

One of the most common responses was that developers simply searched online when they were looking for advice. As P10 put it, "I would Google it, to be honest, and I would look for articles from developers who have focused on building secure systems and kind of start my research there." Developers consulted Hackernews, TreeHouse, StackExchange, Lynda.com, Google, Facebook's Terms of Use, and various smartphone developer forums to search for advice and examples from other developers.

Many developers also consulted their friends and social networks for advice: P7, a professional developer and part-time student, consulted a "Facebook group with... some 300 students," many of whom do mobile development. Others consulted with fellow developers in person, like P5 who said, "I go to a couple meetups, especially if I'm looking for a technical element, or I want to get more into usability." Participants also consulted with contacts who had experience in security or privacy: P10 stated, "I would also talk to my social network, if I knew anyone who has a background in security, about what they would recommend. I fortunately know one or two people."

Lawyers were also consulted when they were available to developers. Some participants worked for companies with dedicated legal staff, such as P13 who stated, "I try to raise [privacy concerns] up to my management level and let them interact with whatever back-end legal that needs to happen. I try to avoid directly communicating with the lawyers." P12 makes it clear that privacy awareness was the legal division's domain: "Ultimately the legal staff is responsible

for making sure that we get the right and accurate information." Generally, the interviews suggest that developers who had access to legal teams seemed to be less personally involved in the understanding of privacy and security regulations.

Some developers relied on terms of service documents provided by the app markets, with P4 stating, "I would expect that those guidelines fall into the realm of what is legally expected in the United States." P8 depended on lawyers to understand regulations that affect app development, leading to less personal knowledge: "The only times we had to change anything, lawyers are on top of it. The reason I didn't bother to know [is that I] depend on a lawyer." As P3 observed, "Unfortunately, I very rarely have time to actually sift through [privacy and security regulations] and try to digest everything that's going on, so I primarily rely on other people to let me know."

3.3.2 Security Tools Used More than Privacy Tools

App developers seemed to use and rely on off-the-shelf or third-party tools for security, but did not have as many tools for privacy. The use of third-party tools could also introduce additional privacy concerns, as these tools may collect information that the app developer was unaware of.

Some developers rely on specific tools to help with security. These tools could include encryption built into the database, SSL code built into the platform, or authentication methods such as Facebook authentication. The tools were perceived as being more secure than hand-rolling implementations themselves. For example, P4 discussed the use of Facebook for authentication, "The expectation is that all the crafty security stuff has been handled by them, because I assumed they'd be smart enough to have that locked down, given that they probably hired security people." However, participants noted that tool usage could be a double-edged sword. For example, participants who used Facebook for authentication had access to much of their users' Facebook profile. Developers discussed weighing the advantages of collecting this information in case it might be useful against the privacy concerns of the user.

Very few interviewees used or knew about existing tools specifically for privacy, such as privacy policy generators, or security audits. One interviewee described his experience with a privacy policy generator as being "good enough" for the time, but not able to handle complex cases. Security audits were only considered by one interviewee; he handled health information and was working with businesses that required audits.

Participants also relied on third-party tools for other uses, such as analytics or various other features. Participants seemed generally unaware of the privacy and security practices employed by third-party utilities used in the development of their apps. Many developers had not personally read the terms of service, were unsure if their lawyers or legal departments had done so, and may have even forgotten the names of the ad networks or web traffic analysis companies they had used. P3 described the need for more digestible information, saying, "if either Facebook or Flurry had a privacy policy that was short and concise and condensed into real English rather than legalese, we definitely would have read it."

3.3.3 Privacy Policies Are Not Considered Valuable

App developers find creating privacy policies to be a low priority or of low value, believing they only offer legal coverage and may turn off users.

Participants were particularly unconcerned about providing privacy policies. In one interview, P4 said, "I haven't even read [our privacy policy]. I mean, it's just legal stuff that's required, so I just put in there." Both P10 and P11 explicitly stated that they were not concerned, because they worked for small companies, with P10 saying: "I have not heard of any startups or small companies getting into trouble for privacy policies," said one, while P11 noted, "Big companies want to CYA [cover your ass], no one is going to go after a small guy like me. I don't generate enough revenue, so if you do sue me you won't get any money." Other developers stated that they did not collect personally identifiable information, and therefore were less concerned about transparency.

Most participants said that while their privacy policies can be accessed on the app website, they were not directly accessible from within the app. In addition, the type of information collected from users would be difficult to find: P8 admits, "We don't make it very obvious, exactly what data we're collecting. I guess it's kind of in the terms of use or privacy policy or something." Paired with the difficulty of quickly accessing an app's privacy policy, this suggests that users will find it tough to determine how their data is being collected and used by apps [76, 75, 35].

Furthermore, some developers were not convinced that users want privacy policies. P7 said users have "been groomed [into] thinking ... [data] is not private... Because it's all anonymous." They felt that as a result, data collected by their app would not surprise users or cause privacy concerns. P3 described the app developer and user relationship in stark terms: "we have consumers as customers. They either trust us or they don't." Some developers were aware of user concerns, noting, for example, the sensitivity of location data. As P8 put it: "it's definitely important to the user to know that their information is safe with [the app]."

When participants put an effort toward alerting users about information collection, they reported lower user retention. Two interviews reported this concern. "We've gone through pretty great lengths to try to make sure that people know exactly what we're collecting and why we're collecting it," describes P3, "So we end up losing out on some number of users because of warnings....they don't take the time to actually read...so they just sort of see this warning and they're like, oh, it must be something bad."

3.3.4 Trade-offs Between Privacy, Security, and Resources

Balancing the need for good security and privacy practices with the cost of actually implementing those practices was a struggle for participants in our interviews. Many discussed privacy and security as being part of the development process but not a top priority, and concerns like monetizing the app or limited resources often trump the desire to follow rigorous privacy and security standards. Some manage to support privacy and security, like P5, who states: "We are trying to balance where that line [between user concerns and the need to store information] gets drawn. I favor privacy."

P10 tellingly struggles with this trade-off when discussing his company's practice of borrowing from other privacy policies, saying, "I don't see the time it would take to implement that over cutting and pasting someone else's privacy policies.... I don't see the value being such that that's worth it."

When questioned about whether their personal feelings towards privacy affected their development decisions, participants gave mixed responses. Some supported privacy protection, such as P10 who said, "I personally have very strong feelings about user privacy," and P5 said that as a

supporter of privacy rights, he made an effort to collect as little user information as necessary for his app. Even self-described privacy advocates and security experts grappled with implementing privacy and security protection with limited time and resources.

Others, while voicing personal concern about privacy, discussed the need to work with clients' wishes. In reference to the privacy of user data in apps developed for his clients, P11 says, "What they want is what they want." Another developer was very invested in privacy protection, but expressed concern that with the threat of his app being copy-catted, advertising was a safer bet for earning revenue than pay-to-download.

This suggests that developers have to weigh their personal desire to respect privacy against the ability to monetize or sell their app, and in particular, developers who work as part of a larger company or who work on commission may be less free to implement good privacy practices than self-employed developers and those who work for small companies. Furthermore, developers consistently discussed the constraints such as time, effort, and money it would take to implement best privacy and security practices.

The cost of collecting and storing data is perceived as minimal. At the same time, interviewees indicated that the cost of developing the code or policies to delete old data or accounts is not prioritized. This is not a question of tools; many of the same tools that allow users to encrypt data also allow them to delete data. Instead, this is a pervasive belief that data may become useful in the future and is therefore worth the resources required to collect and store.

3.4 Survey Method

Based on the interview results, we formed two hypotheses about privacy and security behaviors in app development. We hypothesized that company size would be related to privacy and security behaviors and that revenue models would also be related to privacy and security. In order to test these hypothesis quantitatively, we performed an online survey of 228 United States app developers and product managers. The survey gathered relevant demographics about the developers and their companies, and examined how developers make decisions about privacy and security.

Our survey was designed to take less than 30 minutes, and participants were compensated with a $5 Amazon gift card. Participants were recruited though several online forums, such as reddit subgroups, technical Facebook pages, and through six United States cities on backpage.com. To avoid biasing participation, it was not advertised as a security or privacy survey.

We included four knowledge and attention check questions in our surveys to help us eliminate non-developers and invalid responses. Due to our stringent requirements, we discarded 232 results that either did not have valid responses or were outside the United States. We were left with 228 valid responses from within the United States.

The privacy and security behaviors we examined are those that were recommended by all five of the privacy and security guidelines for app developers that are discussed in Section 6.2. We describe the questions used to measure the privacy and security behaviors.

Security Behaviors

- SSL usage: By encrypting data going over the network, app developers can protect users from data snooping on insecure connections. We measured SSL usage with the question, "Do you use SSL when transmitting data?"

- Encrypting collected data: Encrypting data stored by the app, either in a database or on the phone, protects the user in the case of data breaches. We considered two variables: whether data was encrypted either in the database or when stored on the users' phones.

Privacy Behaviors

- Having a Chief Privacy Officer or equivalent: The existence of a CPO or equivalent indicates that the company is paying attention to privacy and has a specialist who is accountable for privacy. We measured this with the question, "Does your company have a Chief Privacy Officer (or equivalent)?"
- Providing a privacy policy: Privacy policies may indicate that the app company has considered their practices and is being transparent to the user. We measured this with the question, "How does your app inform users about what information it collects?" and the response "Privacy policy on website."

We recognize that there are concerns with self-reported data [102]. We present the results as app developers' own conceptions of their work, not as ground truth. Our findings may differ than those of previous research based on scans of the app stores. For example, our questions are on a per-developer basis, and developers may have created more than one app. Our results are for all platforms, and both free and paid apps. Furthermore, our survey was done in August 2013 and may be more recent than published papers' results.

Behavior	percent
Use SSL	83.8%
Encrypt data on phone	59.6%
Encrypt data in database	53.1%
Encrypt everything (all data collected)	57.0%
Revenue from advertising	48.2%
Have CPO or equivalent	78.1%
Privacy Policy on website	57.9%

Table 3.3: Percentage of respondents who reported various privacy and security-related behaviors. Participants could select multiple options.

3.5 Survey Results

We first present the demographics of our survey participants, including their training in privacy and security, and where they look for advice when making privacy and security decisions. We then discuss the app companies they work for, including size, revenue model, and use of third-party ad and analytics tools. We also present some exploratory work on data collected by app developers, including data types that have not been measured in previous work. We then describe our hypotheses about security and privacy behaviors; that they are correlated to each other, to company size, and to revenue. Finally, we report the results from testing the hypothesis.

Role	Participants
Programmer or Software Engineer	58%
Product or Project Manager	31%
Tester or Quality Assurance	20%
Manager	12%
CEO or President or Owner	8%
Marketing	4%
Student	4%
User Support	3%
Not currently working/Currently unemployed	1%

Table 3.4: Percentages of participants in different roles. Participants could select multiple options.

3.5.1 Participant Demographics

Most of our respondents were programmers, product managers, or quality assurance testers. The average age was 30 years old (range: 18-50 SD = 5.6). We did not collect additional personal demographics such as gender. Participants selected their professional role from a multi-select list. Our recruitment stated specifically that we were looking for app developers or product managers, so it was not surprising that 78% of participants were programmers or software engineers, product managers, or both. Other participants were testers, managers, and CEOs. The role breakdowns are shown in Table 3.4.

We asked participants to describe their formal privacy and security training. Our results directly contradicted our interviews, in which few people claimed to have formal privacy or security courses. However, most interview participants worked for small companies. In the survey, only 7.3% claimed to have no formal privacy or security training. 62.9% of respondents claimed to have taken a privacy or security training course. Many also stated that they had received corporate training on privacy or security (62.5%) or attended a professional development seminar or workshop (43.5%). App developers in companies of size 31-100 were the most likely to receive corporate training, and companies with only one employee were the least likely to receive training.

In order to determine how app developers were making privacy and security decisions, we asked participants from whom they sought advice about privacy and security. This is useful for two reasons: first, it provides some insight into the level of expertise available to developers, and second it may allow better framing of educational campaigns for app developers about privacy and security. Figure 3.1 shows from whom participants sought advice, based on their company size. The company size significantly affected whether participants sought advice from their social network, security or privacy exports in their company, or no one (Kruskall-Willis test, $p<.001$). Participants from companies with fewer than 9 employees were more likely to get advice from their social network, or to ask no one. Developers in larger companies (31-100 employees or 100+ employees) were more likely to ask a privacy or security specialist within their company.

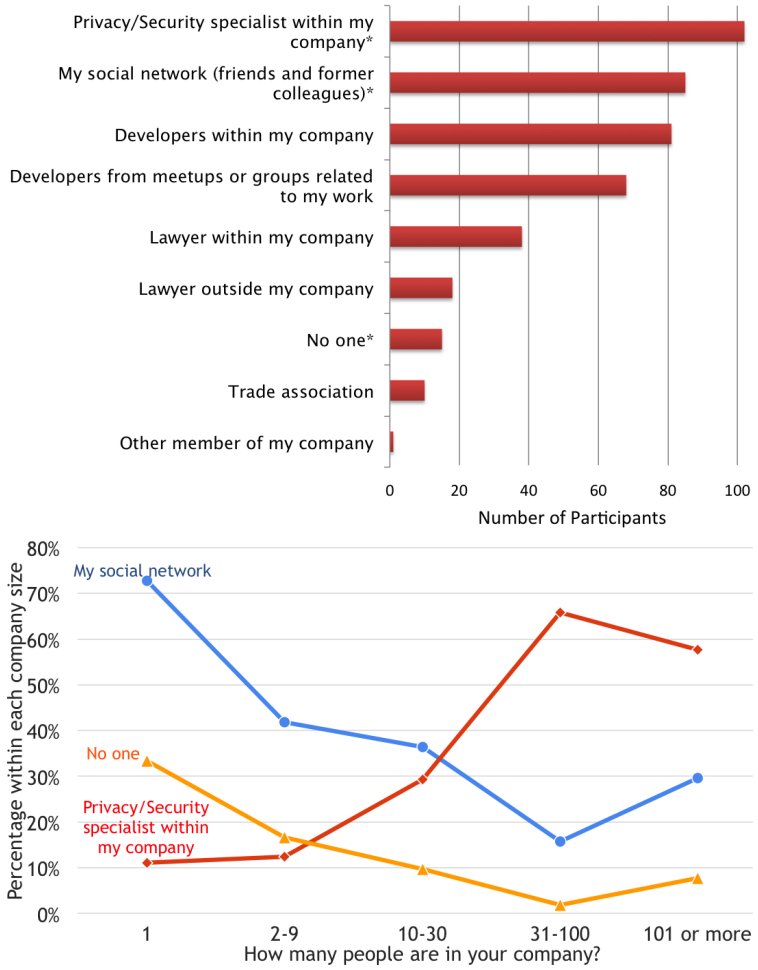

Figure 3.1: Survey participants' response to the question "Who, if anyone, do you turn to when you have questions about consumer privacy and security?" Responses significantly different based on size of company are marked with *. The bottom figure shows the 3 significant selections by company size. Developers at small companies rely on their social networks or no one, while developers at larger companies rely on specialists within the company.

3.5.2 App Company Characteristics

We discuss the categories of app companies represented by the survey participants. We do not claim that this is a proportionate sample of app development companies in the United States. Instead, we discuss the characteristics to put our other findings into context.

Equal numbers of participants were building or planning to build iOS (142) and Android (142) apps, with much smaller numbers for other platforms (38 for Windows, 10 for Blackberry and 6 for Palm, and 1 other). Over one quarter of participants (63) said they were developing for both Android and iOS. The survey participants represented different size companies and development groups. The percentages of developers in companies sized 1, 2-9, 10-30, 31-100, and 101 or more were 4.9%, 14.8%, 19.7%, 48.4%, and 12.1% respectively. However, the size of app development

Revenue Model	Total	Only Source
Advertising	48.0%	9.4%
Paid Download	44.8%	10.3%
Free trial, upgrade to premium	37.7%	8.1%
Subscription	25.1%	8.5%
Pay-per-use	21.5%	2.2%
In-app purchases	19.7%	4.0%
All of above except Subscription		7.6%
Advertising and Paid Download		8.5%
Hosting	3.6%	.9%
Other	3.6%	3.1%
None	2.2%	2.2%

Table 3.5: Revenue models of respondents. Respondents may have chosen multiple responses; the middle column represents all participants who selected that model, and the right column shows how many participants selected only that model. The app revenue models used are based on Leem et. al [114].

groups (employees working directly on the app) were typically between 2-30 people.

Participants were asked to categorize their app, using a list that was a combination of Apple iTunes store and Android Play store categories. All categories were represented, with Games (17.7%), Entertainment (12.5%), and Finance (10.8%) appearing most frequently.

Ad or Analytics company	Survey
Google analytics	82.1%
Google ads	64.1%
Amazon ads	43.9%
Flurry analytics	16.6%
AirPush	14.8%
AdMob	13.9%
No ads	13.5%
No analytics	12.6%
Medialets	11.2%
AdWhirl	8.1%

Table 3.6: Percentage of respondents who reported using various analytics companies. Participants could select multiple options. Only libraries with 10% or more of respondents are shown.

Apps may collect data for their own use, but interviewees also indicated data is collected for secondary uses such as advertising or analytics. Our interviews indicated that app developers were not always aware of the data collection of the third-party API's or toolkits they were using. Table 3.7 shows respondent's knowledge of third-party data collection practices. Just over one-third of app developers claimed that they knew exactly what data is collected by third-party tools.

These responses may represent more of the developers' self-perception than reality. For example, of the developers who claimed they did not use third-party tools, the majority answered separate questions about using third-party tools differently: 70% said they used at least one ad company, and 87% used an analytics company. This suggests confusion either about the question (different definitions of "third-party tools") or their own apps' behaviors.

Seventy percent of respondents stated that their apps were already earning revenue. As seen in Table 3.5, a variety of revenue models and types of apps were also represented, with advertising and paid downloads being the two most popular options. 52% of participants reported relying on more than one revenue model; the two most popular combinations were 1) advertising combined with paid download and 2) all revenue models except subscription and hosting. The more revenue models selected, the more likely they were already earning revenue (χ^2 tests $p<0.001$). The distribution of revenue models was similar for both iOS and Android. Due to the nature of our questions, these results could represent the revenue model used by an app developer across multiple apps that they have developed, as opposed to the models used within one app.

Response	percent
My app doesn't use third-party tools	41.7%
I know exactly what kinds of data the third-party tools are collecting	35.9%
I have some ideas about third-party data collection but don't know for sure	22.0%
I don't know	.04%

Table 3.7: Responses to "How familiar are you with the types of data collected by third-party tools?"

Overall, most app developers (87.4%) used at least one analytics company, with one in five using two or more analytics companies. Table 3.6 shows which companies were used by app developers. Most apps also used an advertising company: 86.5% selected one or more advertising companies in use, using on average 1.78 ad companies (SD=1.33). Interestingly, app developers were likely to use an ad company regardless of whether they relied on advertising for revenue. Of app developers who did not select advertising as a revenue source, 82% still resorting to using at least one advertising company. We speculate that app developers may be including advertising API's without earning money from ads; however this merits further exploration.

In our survey, 41.7% of developers self-report that they do not use a third-party tool. It is important to understand app developers' self-perception, as it will likely influence their need to consider third-party tools' data collection when creating privacy policies or handling data. If developers are not aware of or fail to consider some libraries, they will not report on their behavior when making privacy decisions. Our 2012 scan of free Android apps indicates that 50.2% of free Android apps did not use ads, analytics, social networks, or payment APIs, which is higher than our survey findings suggest [123]. We find that 36.3% of developers reported using exactly one ad library.

3.5.3 Collection of Sensitive Data

As we did not discuss the collection of sensitive data in our interviews, we did not formulate specific hypotheses to test. Therefore, we show the results of some exploratory analysis. Table 3.8

shows which data the app collected or stored. Due to an error with the survey, 5 participants did not answer this question. They were removed from the analysis of this question.

We asked about data that may be privacy or security sensitive. Several data items corresponded to Android or iOS permissions and warnings (such as location), but other data can be collected without warning the user. This includes which apps are installed, and sensor data from accelerometers. The user would only know about this data collection if it were included in a complete privacy policy. Other data that don't trigger permission notifications are credit card information or password; these are input by the user but require that the app developer handle them securely.

An average of 5.5 out of the 10 sensitive variables we asked about were collected. Based on our interviews, we were not surprised that most apps did not collect or store users' passwords or credit card information. Instead, apps that need this information may often rely on third-parties such as Facebook to do authentication or to handle credit card information. Unsurprisingly, apps collected information pertinent to their app, such as level attained in a game. It is startling that three quarters of app developers collected which other apps are installed on the user's device. Apps may do this to explicitly collaborate with other apps or services, such as a todo list app accessing a calendar app. However, information about installed apps can have privacy implications, such as family or health status if related apps are installed.

Data Type	Collect or Store (%)
Parameters specific to my app	83.9%
Which apps are installed	73.9%
Location	71.6%
Advertising ID	70.6%
Sensor information not location-related	63.0%
Phone ID	54.5%
Contacts	54.0%
Phone Number	44.1%
Password	35.5%
Credit card information	30.3%

Table 3.8: Percentages of respondents who collected or stored selected data.

We ran an independent samples t-test on revenue model and number of sensitive variables used. Only two types of revenue significantly impacted the mean number of data items used or collected: in-app purchases and advertising. The mean number items collected or stored was higher for participants with advertising revenue (μ=5.16 and std 2.73) than for those who did not use ads as a revenue source (μ=5.88 and std=2.49) and the difference was significant (t-tests p=.04). On the other hand, in-app purchases resulted in less data collected (t-test p=.04), with an average of 4.77 (std=2.65), while those that did not use in-app purchases averaged 5.68 (std=2.65).

The category of app also significantly affected the amount of data collected (ANOVA p=.007). Of the categories with 10 or more responses, finance used the most sensitive variables on average (μ=6.36) while entertainment collected the least (μ=4.73). Only 20% of respondents with a finance app had advertising revenue, while 57% of entertainment apps had advertising revenue.

Leontiadis et al. found that free apps required more permissions than pay-to-download apps [118]. Our findings support this. We find statistical differences in the amount of data collected by revenue (ANOVA, $p<0.001$), and find that the amount of data used by developers with paid-download revenue models only (μ=3.78, SD 3.03) is significantly different from the amount of data collected by advertising-only revenue models (μ=6.48, SD 2.40) (ANOVA multiple comparison, p=0.013 with Bonferroni correction).

3.5.4 Hypothesis Testing and Results

In this section, we describe our hypotheses about privacy and security behaviors and the results of testing each hypothesis. Table 3.3 summarizes the percentages of respondents who claimed to engage in the each privacy and security behavior. Table 3.8 summarizes the number of respondents who collected or stored the data types we examined.

Hypotheses 1: Behaviors are correlated

First, we hypothesized that security and privacy behaviors would be positively correlated, and that there would be developers who were generally concerned about privacy and security and demonstrated all or most behaviors, while others would not display any such behaviors. Our hypothesis is mostly supported; all behaviors are significantly and positively correlated at the p=.05 level except Privacy Policy and SSL, as shown in Table 3.9.

H1: *Security and privacy protective behaviors are correlated.*

	CPO		Encrypt Everything		Privacy Policy	
	ϕ	p	ϕ	p	ϕ	p
Encrypt Everything	·272*	<·001				
Privacy Policy	·159*	·018	·228*	·001		
SSL	·257*	·001	·217*	·005	·157	·063

Table 3.9: Correlations between the security and privacy behaviors. The Phi Coefficients (ϕ) indicate that the behaviors are generally positively but weakly correlated. * indicates significant correlation at the p=.05 level.

For hypothesis H2 and H3 we ran eight χ^2 tests separately. We conservatively correct the standard p-value of .05 with Bonferroni correction, and use a significance level of 0.006 (0.05 divided by the number of tests).

Hypotheses 2: Company size

We are aware that startups or app development companies with small teams and little investment may not have the resources, in terms of time or money, to invest in privacy and security. Therefore, we suspected that small companies may be less likely to engage in the privacy and security behaviors that require additional employees (a CPO), additional time (creating a privacy policy),

or additional resources. For example, encryption may require more equipment or software. Using SSL may require additional developer time or experience.

H2a: *Company size correlates to having a CPO.*
H2b: *Company size correlates to having a privacy policy.*
H2c: *Company size correlates with encrypting everything.*
H2d: *Company size correlates with using SSL.*

We found that the size of a company does help determine whether they have a CPO (χ^2 test $p<0.001$), whether they have a privacy policy (χ^2 tests $p=.002$), and whether they encrypt everything (χ^2 tests $p<0.001$). However, the company size was not correlated with SSL using the conservative corrected significance level (χ^2 tests $p=.009$). As one respondent wrote in an open-text field, "We are a small, two-person shop. Although we don't have CxO positions, we do understand the need to protect the privacy of our users. Our app embeds a privacy statement in an easily identifiable location."

The percentages of companies engaging in the above privacy and security behaviors grows as the company size grows, up to the 31-100 employee companies. For example, all of the respondents with company sizes of 1 said they did not have a CPO or equivalent, while only 58.8% of respondents in companies from 2-9 had someone responsible for privacy, compared to 89.6% and 92.6% of companies size 10-30 and 31-100 respectively. This is shown visually in Figure 3.2, and is similar for the other privacy and security behaviors. However, this trend of improved privacy and security practices does not hold for company sizes greater than 100. We speculate that app developers in larger companies may not be as aware of all their company's practices.

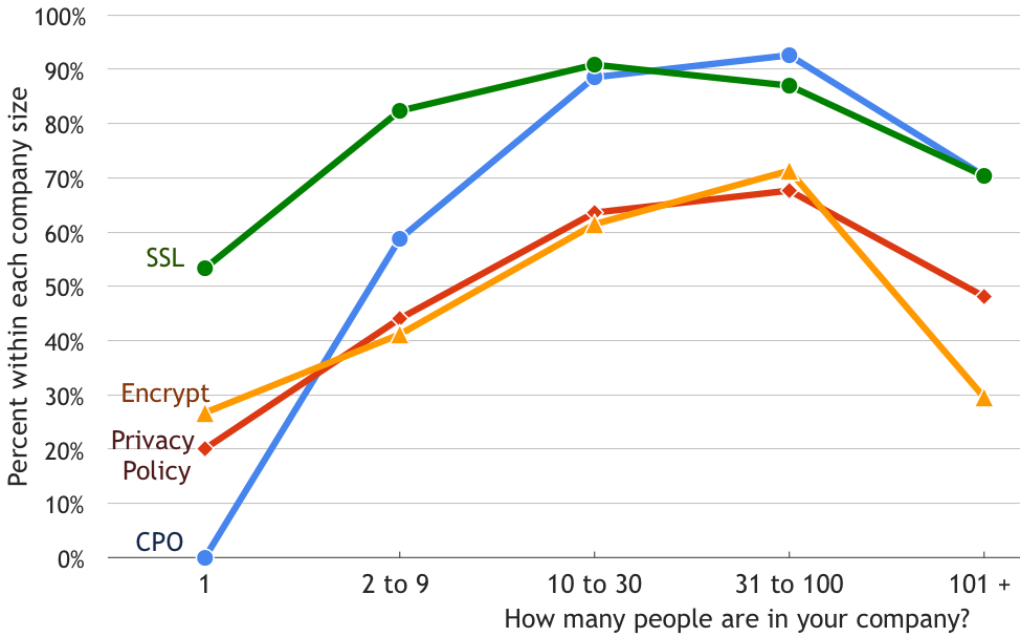

Figure 3.2: The size of the company is related to whether or not the company has privacy and security behaviors. Companies with 31-100 employees are the most likely to engage in these behaviors.

Hypotheses 3: Revenue model

We were curious about the impact of the revenue model on privacy and security behaviors, and hypothesized that certain revenue models, such as advertising, were less likely to show privacy and security behaviors

H3a: *Revenue model is correlated to having a CPO.*
H3b: *Revenue model is correlated to having a privacy policy.*
H3c: *Revenue model is correlated with encrypting everything.*
H3d: *Revenue model is correlated with using SSL.*

Since 47 unique combinations of revenue models were reported, we examine the most common models and combinations, which are shown in Table 3.5. All other combinations (with fewer than 10 responses) were combined into an "other" category. At our conservatively corrected p-value, none of the results were significant (CPO p=.035, encrypt p=.029, SSL p=.037, privacy policy p=.019). However, we note a few interesting cases. An advertising revenue model indicates low adoption of privacy policy, but is average on the other measures. However, we note that all 17 of the developers who used every model except subscription also claimed to implement all the privacy and security sensitive behaviors. The only common feature we found across all 17 of these developers is that they all received corporate privacy and security training as well as college classes.

3.6 Discussion

Our results indicate that many developers lack awareness of privacy measures, and make decisions in ad hoc manner. While most developers claimed to be using SSL and to have a CPO or equivalent, only slightly over half of our survey participants claimed to employ the other recommended privacy and security measures such as encrypting everything or having a privacy policy on their website. Our interview respondents discussed encrypting some, but not all of their data, and having little belief that privacy policies were useful. The survey respondents indicated a high level of data collection. Roughly three-quarters of developers collected information about the other apps installed on the users device. Some interviewees discussed collecting data that they didn't need, but thought might be useful in the future

While several government agencies, non-profit groups, and industry groups have developed guidelines for app developers on suggested privacy and security practices, the app developers we interviewed were not aware of and had not read these documents. This suggests that public policy around privacy and security is not reaching developers. In this section, we discuss hurdles to better privacy and security behaviors, and provide recommendations to encourage privacy-sensitive behaviors.

3.6.1 Third-Party Tools Should be More Transparent about Data Collection

Most app developers in our survey used third-party advertising or analytics services. Previous work shows that these libraries have permission to collect sensitive data [51, 123]. The developers we interviewed discussed their difficulties reading the policies and terms of use for the third-party

APIs or services that they integrated into their apps. Popular ad and analytics companies should provide information about their data collection to app developers in an easy-to-read format. They should explain both what they collect and the purpose of that collection. This information could be provided in two places: as part of a quick-start guide, so developers can review before integrating the code, and after the developers have configured the third-party settings, so they can review how their choices impact their users' privacy and write their privacy policies.

In addition, survey participants demonstrated some confusion about whether they were using third-party tools, providing contradictory responses in different questions. This indicates that tools that automatically detect and describe third-party data collection may be helpful for developers. To allow app developers to easily understand the data that third-parties collect from their users, we suggest that short, readable text about data collection be provided by ad and analytics tools. This short readable text should be available in strategic locations, not just in the longer privacy policies. For example, many ad or analytics tools provide getting started guides to app developers, which could include simple explanations of the data collected.

The developers we interviewed stated that if they read any terms of service or policy at all, it is likely to only be the terms of service of the app market or platform, such as the iTunes store or Android Play market. Therefore, we suggest that platforms also have a role in allowing app developers to understand what data is collected by third party apps. While more notice could be useful, platforms could also provide tools that integrate into the development process. These tools could help developers detect and understand what data is being collected by third-parties. For example, the emulators that developers use to test their code could provide logs of the permission requests.

Unfortunately, third-party tools may collect information about the smartphone user while having little or no relationship with the user, and thus have little incentive to protect user privacy. This may indicate a need for legislation to incentivize third-parties to provide clear information about their data collection to app developers and the end users.

3.6.2 With a Little Help From my Friends

App developers often mentioned searching for resources about security and privacy on the web. In addition, app developers in small companies rely on their friends and social networks for advice about privacy and security, while developers in larger companies may have experts within their company or legal counsel to turn to. Security and privacy advocates may find traction by intervening at a social level, such as by meeting with developers to discuss and improve their practices.

3.6.3 Legalese Hinders Reading and Writing of Privacy Policies

Less than half of small companies (fewer than 10 employees) informed their users about data collection through privacy policies on their websites. Several of our app developer interviewees had never read their own policy, and many others did not view it as a tool to communicate with users. Privacy policies were perceived as a tool that might protect them against lawsuits, but that small companies would not be targeted for lawsuits. This suggests that there is a need to emphasize that privacy polices need not be legalese, and can be an opportunity to communicate with their users. Furthermore, some interviewees expressed concern that full disclosure scares users away.

This suggests that required, standardized privacy notices might be a benefit for privacy-protective apps. Efforts of the government to develop such notices may provide guidance [20]. If all apps are required to provide notices, those who have good practices would not be punished for transparency.

Several interviewees believed that complying with the app stores' policies would provide sufficient legal protection, or that the app store would be monitoring them for compliance. This suggests that platform developers and market controllers are well-placed to encourage privacy and security behaviors. Platforms can highlight best practice notices and checklists, making them clear and accessible to app developers.

Useful privacy polices should be easy to generate for app coders and also be useful to app users. Platforms such as Android and iOS already help with this, but they can do more to address users' concerns. These platforms automatically detect requests for data permissions, and then generate notices for the app users. However, these notices fail to provide several crucial pieces of information that smartphone users want to know, such as why is the data being collected, and where is it going. For example, users may be interested in knowing that an app collected their location to provide them with targeted ads, but may be uncomfortable when that information is collected too frequently or shared with too many third-parties (see Chapter 5). Platforms could require that apps include additional fields explaining the purpose and destination of data along with the permission request. Platforms could then display this information about data collection to the user. This option would allow app developers to easily include privacy notices for their users.

3.6.4 Small Companies Need Privacy and Security Tools

The smaller companies were the least likely to engage in privacy and security behaviors. Companies with fewer resources are less able to devote time or money to privacy and security issues. Therefore, small companies may need additional help or resources to overcome the hurdles to developing privacy policies and encrypting data. Improved tools for developers that are integrated into the development of apps will allow for privacy to be built in, as opposed to tacked-on, or ignored. We offer two examples of integrating privacy into the development process so that it is not a separate task or afterthought.

Platforms can nudge developers to protect privacy by developing and clearly documenting privacy-preserving APIs. These APIs may discourage fine-grained data collection. For example, instead of API calls for location that return the specific latitude and longitude, platforms can provide region or zip code by default. By placing the documentation for the more privacy-protecting calls in prominent places, app developers might only request the more specific, privacy-invasive information when it is really needed [100].

Companies of all sizes could be nudged to minimize data collection with tools that help developers decide what data to collect and when to delete it. For example, many app developers use cloud storage to store their data. For example, one popular option is Amazon Cloud Storage. Currently storing data and keeping it forever is cheap and easy. However, storage solutions could make it easier for developers to delete old or expired date. For example, Amazon Cloud Service could require configuring the lifecycle of stored data, requiring an expiration date when the data is first specified. This would force developers to define retention periods for their data, and encourage removing old data.

We suggest that privacy and security tools should be specifically targeted small development companies with few resources. OS developers or open-source developers could focus on providing free tools to developers. These tools should be usable and not require legal expertise.

3.7 Conclusion

While there is general awareness of need for security measures, such as encrypting information or using SSL, there was a lack of understanding around privacy best-practices. Small companies rely on social networks and search engines for privacy and security advice. Privacy and security tools for developers must be quick, simple, and cheap, so that they can be used by time- and resource-constrained small companies. Platforms should make sure that it is easy to implement good security practices. App stores should provide privacy and security checklists, as they are uniquely positioned to reach developers. Third-party tools should make their data collection clear to developers and end users. More work is needed to make developing clear privacy policies a simple and routine part of app development.

One option to make privacy a higher priority is through public policy, regulation, and enforcement. We discussed the guidelines created by policymakers in Chapter 1.4.3. These guidelines currently are not enforced; app developers who know they will receive a fine for failing to follow such guidelines may be more likely to make privacy a priority. However, due to the fragmented nature of the app developer community, with thousands of small app development shops, enforcement is likely to be costly, time-consuming, and incomplete. Policy is more likely to have an impact if it can address a few bigger players, such as platform developers or ad companies.

A more light-handed alternative to policy is to integrate privacy tools into the development process in such a way that good privacy behaviors are less time consuming. App developers are more likely to protect data when it becomes part of their existing workflow. This could improve the privacy of the users without imposing burdensome time or money costs on the app developers. Such tools and nudges, however, require time and effort from the platforms, third-parties, and other players in the mobile eco-system. These players may lack any financial incentive to do so. However, increased focus from regulators, such as the White House on Big Data [95] and the NTIA multi-stakeholder efforts on privacy [20], may put additional pressures on everyone to protect the privacy ecosystem

We suggest that privacy in apps will improve only when there is more pressure to make privacy a priority. But this pressure should not just be on the myriad of independent developers alone, but also on other players in the app development ecosystem. These players include platform developers and app store management, as well as the analytics and ad companies, and data storage providers.

We have offered insight into app developers' privacy behaviors. Overall, helping app developers overcome the hurdles to good privacy practices is a crucial part of improving the privacy of their users. We suggest that all the players in the mobile eco-system play their part.

Acknowledgements

This research was funded in part by Google, NQ, John and Claire Bertucci Fellowship, and NSF grants DGE0903659, CNS1012763, and CNS1228813.

Chapter 4

Is Your Inseam a Biometric? A Case Study on the Role of Usability Studies in Developing Public Policy

In this paper, we present a case study of applying usable privacy methodologies to inform debate regarding a multi-stakeholder public policy decision. In particular, the National Telecommunications and Information Administration (NTIA) relied on a multi-stakeholder process to define a set of categories for short-form privacy notices on mobile devices. These notices are intended for use in a United States national code of conduct to assist mobile device users in making decisions regarding data collection. We describe, specifically, a 791-participant online study to determine whether users consistently understand these proposed categories and their definitions. We found that many users did not understand the terms in our usability study. The heart of our contribution, however, is a case study of our participation in this group as academic usable privacy and security experts, and a presentation of lessons learned regarding the application of usable privacy and security methodology to public policy discussion. We believe this work is valuable to usable privacy and security researchers wishing to affect public policy.

Public policy and regulation intersect with human computer interaction in many domains. In areas including voting machines, accessibility, and privacy, regulators may try to step in where market forces have failed. We argue that policy-making can and should be be informed by usability studies, and those policies in these areas that are not informed by the usable privacy and security community may be ineffective. We provide a case study of a multi-stakeholder process to standardize smartphone privacy notices in the United States. We present a user study, which we ran near the end of the process, that demonstrates the shortcomings of failing to take usability into account throughout the process. Furthermore, we discuss what lessons can be learned from our experience.

The U.S. National Telecommunications and Information Administration (NTIA) initiated a multi-stakeholder effort to develop a standardized short-form privacy disclosure on mobile devices. These standardized disclosures will show the user both what data is being shared and with which entities it is being shared. The year-long NTIA multi-stakeholder process (NTIA MSHP)

This chapter is based on a paper [38] and technical report [37] co-authored with Rich Shay and Lorrie Cranor.

lasted from June, 2012 to July, 2013. The NTIA stakeholders – representing app developers, consumer groups, and government – developed a Code of Conduct that provides guidance for app developers for a short-form privacy notice. This Code outlines seven categories of data and eight categories of third-party entities that apps should include in short-form privacy notices, but does not specify a format for these notices.

We present the results of a 791-participant online study in which we investigate whether participants are able to categorize realistic data-sharing scenarios using the NTIA MSHP categories. We also present the same categorization from four experts who participated in the NTIA MSHP process. Of the 52 examples given in our scenarios, participants showed low common agreement for how to classify the data or entity in 23 cases. Overall, we found that many of the proposed categories and definitions were not consistently understood by our participants, including our expert participants. We discuss categories that need clarification, and offer suggestions for improving the Code based on our findings. This study was undertaken by the authors independent from the NTIA MSHP, and is the first and only human-subject study conducted on those categories to date. Our results suggest that user studies should have been a larger part of the NTIA MSHP.

Based on our experience, we provide lessons learned for usable privacy and security experts who wish to participate in multi-stakeholder processes. Our contribution is primarily the insights and recommendations based on our experience participating in this process. We hope our discussion will enable and encourage usability experts to participate in public policy processes more readily and advocate for legislation and codes that are better informed by user studies.

We first discuss the background of the NTIA MSHP, along with related work. We then describe the methodology of our user study. We next present the high-level results from our user study. We explore the limitations of our study. Finally, we discuss recommendations and lessons learned for usability researchers who wish to improve policy making.

4.1 Usability and Public Policy

Usability and consumer testing have previously played a role in developing standards and policy for technology. We discuss several examples of usable privacy and security experts who have been involved in the public policy process. We first examine privacy issues in particular, and then briefly discuss two other examples of usability and public policy: voting machines and accessibility.

Independent academic research has evaluated privacy standard proposals created by the the The World Wide Web Consortium (W3C), such as the Platform for Privacy Preferences (P3P) and Do Not Track (DNT). Although P3P does not include a user interface standard, usability tests of prototype user agents conducted independently by members of the working group informed the standard's User Agent Guidelines [61]. While the DNT process refrained from defining a user interface, one independent academic user study examined the usability of an implementation of DNT [115], and another user study performed by the chair of the W3C DNT working group examined user understanding of DNT [129].

Academic researchers have proposed privacy label standards. Kelley et al. developed and tested a "privacy nutrition label" for websites. They also found that a tabular format was liked by users and facilitated policy comparison [106].

Consultants have also been engaged by policy makers to evaluate the usability of standards. For example, The U.S. Gramm-Leach-Bliley Act of 1999 (GLBA) required financial institutions

to provide a privacy notice to customers. To develop this notice, the Federal Trade Commission hired Kleimann Communication Group to conduct user studies and develop a privacy notice prototype. Their qualitative research involved iterating over prototypes with several studies — including focus groups and usability testing [4]. This prototype was then tested against several others with quantitative testing [119], and the results were used to develop the final 'model' form, which presents information in a tabular format [5].

Think tanks and user-advocacy groups have also been engaged in evaluating the usability of privacy notices and icons. For example, the Future of Privacy Forum researched consumer's responses to notices about online behavioral advertising (OBA). It found that transparency and choice increased people's comfort with OBA. That study also compared the effectiveness of different icons in communication about OBA [92]. Unfortunately, the icon revealed as the most effective was not selected by the ad industry.

Usability issues with voting came to national attention during the 2000 U.S. presidential campaign. In 2002, Congress passed the Help America Vote Act, which included helping states evaluate their voting systems. Norden et al. discuss voting machine and ballot usability, and provide four case studies in which usability experts evaluated voting systems and worked with public officials to improve usability [139].

Another issue at the crossroads of usability and public policy is the development of accessibility standards. Lewis and Treviranus explain that public policy impacts the accessibility of information technology content and services by influencing funding, setting standards, enforcing regulation, and promoting adoption. For example, Section 508 of the Rehabilitation Act helps adoption of standards by requiring all federal websites to meet accessibility standards. The authors encourage participation in standards development or related activities to influence public policy [120].

4.2 Multi-stakeholder processes in Privacy Policy

In this section we discuss the multi-stakeholder initiative to create standardized mobile privacy notices. More background on public policy is available in Chapter 1.4.3. Several studies have examined existing smartphone privacy notifications and found there is room for improvement. This was discussed in Chapter 1.4.5. The following paragraphs specifically address the National Telecommunications and Information Administration (NTIA) initiative on Mobile Application Transparency in 2012.

In 2012, the White House issued a report on consumer data privacy, which included a Consumer Privacy Bill of Rights [11]. The second principle in the bill of rights is transparency, which is summarized as: "Consumers have a right to easily understandable and accessible information about privacy and security practices." The White House report emphasizes the role of multi-stakeholder processes to develop and define privacy practices and technologies, and to develop "enforceable codes of conduct." It calls upon the Department of Commerce's National Telecommunications and Information Administration (NTIA) to lead multi-stakeholder processes. The NTIA launched one such initiative on Mobile Application Transparency in 2012. The result was a draft Code of Conduct for mobile short-form notices. That draft defines a standard short-form privacy notice for apps, which is not to be a substitute for a longer, complete privacy policy.

Multi-stakeholder processes are viewed by some as an improvement over industry self-regulation, in that more stakeholders have a voice. The development of a privacy Code of Conduct through a multi-stakeholder process is thought to facilitate the involvement of media, citizens, and academics, as well as lobbyists, non-profits, and industry. This was the first multi-stakeholder process conducted by the NTIA, and was considered a learning process. The NTIA MSHP included meetings every few weeks in Washington, DC that were open to the public and allowed for remote participation by calling in or viewing a webcast. Participants included lobbyists from companies involved in app development, representatives of consumer-advocate non-profits, and privacy lawyers representing interested companies.

The NTIA multi-stakeholder group struggled with the role of usability testing in drafting the policy. While a usability subgroup was initiated and met several times, no consensus was reached on what should be tested, or by whom. Some stakeholders argued that it had been so difficult to reach consensus on the wording of the code that they were unwilling to submit it to user testing. User testing risked dragging out the process longer than needed. Finally, the subgroup did not perform any user or usability studies. Some participants argued that usability was never a goal of the process.

The user study reported in this paper was initiated and run independently of the usability group, by our own research group at a university. As a participant in the user-study subgroup, we became aware of the practical issues in initiating a user study, and realized that if a user study were to be done, it would need to be done independently of the group, with our own design, initiative, and funds. Our goal was to examine one portion of the notice, in particular the understandability of the wording suggested for the short form notices. If our study found that there were problems with understandability, we hoped that this would influence the Code and the process, and allow the selection of improved terminology.

4.2.1 NTIA MSHP draft wording

The NTIA MSHP draft includes seven categories of information to include in in-app privacy disclosures. It also includes eight categories of entities with which data might be shared. The draft includes short definitions for all information types and entities – referred to throughout the paper as the "parenthetical" text – shown in parentheses below. We tested the wording used in the NTIA MSHP draft code published on April 29, 2013.[1] We deliberately did not change, add, or in any way modify the wording or punctuation.

The categories for data types are:

- Biometrics (information about your body, including fingerprints, facial recognition, signatures and/or voice print.)
- Browser History and Phone or Text Log (A list of websites visited, or the calls or texts made or received.)
- Contacts (including list of contacts, social networking connections or their phone numbers, postal, email and text addresses.)
- Financial Information (Includes credit, bank and consumer-specific financial information such as transaction data.)

[1] **http://www.ntia.doc.gov/files/ntia/publications/mobileappdraftapril29_2013_draft1b_fs.pdf**

- Health, Medical or Therapy Information (including health claims and information used to measure health or wellness.)
- Location (precise past or current location and history of where a user has gone.)
- User Files (files stored on the device that contain your content, such as calendar, photos, text, or video.)

The categories for entities with which data was shared are:

- Ad Networks (Companies that display ads to you through apps.)
- Carriers (Companies that provide mobile connections.)
- Consumer Data Resellers (Companies that sell consumer information to other companies for multiple purposes including offering products and services that may interest you.)
- Data Analytics Providers (Companies that collect and analyze your data.)
- Government Entities (Any sharing with the government except where required or expressly permitted by law.)
- Operating Systems and Platforms (Software companies that power your device, app stores, and companies that provide common tools and information for apps about app consumers.)
- Other Apps (Other apps of companies that the consumer may not have a relationship with)
- Social Networks (Companies that connect individuals around common interests and facilitate sharing.)

4.3 Methodology

We conducted an online survey using Amazon's Mechanical Turk crowdsourcing service (MTurk)[2] over a two-week period in May 2013. Participants were recruited with the text, "Give us your opinion about information about smartphone apps. This should take 15-25 minutes," and paid $1 for completing the survey.

Previous research has demonstrated that offline experimental results can be successfully replicated using MTurk [140]. Furthermore, while MTurk workers are younger and more technically savvy than the general US population, MTurk has been shown to provide a more diverse sample than a university lab survey [56, 97, 46]. Using MTurk has allowed us to conduct our study with a larger and more diverse sample than would otherwise have been possible.

We also invited NTIA MSHP members to participate in the same study. MSHP members answered two additional questions about their role in the process. MSHP participants were not compensated. The process for participating in the NTIA is open, but requires a time commitment and dedication to attend and participate in the meetings. These participants are considered experts, since they are familiar with objectives of the NTIA and have worked to shape the draft Code. We advertised the study to MSHP members through announcements by email and a brief presentation at one of their meetings. Response was limited, with only 4 experts (out of 25-50 participants) taking the survey. While we present their responses, we make no statistical claims about the results.

[2] `https://www.mturk.com/`

4.3.1 Survey Design

Our survey presented participants with a sequence of ten randomly-ordered smartphone-app scenarios. In each scenario, we described the app's purpose, the data it collects, and the entities with whom it shares that data. Some scenarios also included an explanation about why the data is collected. We then asked participants to categorize both the data being collected and the entities with which it is shared, according to the NTIA categories. An example scenario is below. All ten scenarios are provided in the Appendix B.

> The Fitness app integrates with your FitMonitor (FitMonitor is a special pedometer and activity monitor, purchased separately) to allow you to track and improve your fitness activities and level.
>
> Fitness app will collect information on how many steps you have taken, how long you've slept, and allow you to enter you weight and body fat.
>
> Fitness app will notify sports and health companies if you achieve certain goals, and these companies will send you valuable coupons as awards.

We attempted to represent every data category and every entity category from the NTIA draft in our scenarios. Our scenarios were designed to be realistic. Many scenarios were based on real apps or websites, though we changed the names and adjusted the wording in order to avoid confusion if the participant was already familiar with the real app. In some cases, we took descriptions of apps from the app stores or web sites. We guessed with whom data would be shared, as the apps typically did not reveal this. Our scenarios were more concise, explicit, and specific than typical privacy policies. In three cases, we used the names of real companies — Apple, Facebook, and Google — in order to investigate whether participants considered them to be social networks or operating systems. We included several scenarios that may be considered privacy sensitive. Two scenarios described collecting financial information and another described collecting the user's weight. The "FindMyKid" app allowed a user to set up tracking on someone's phone without that person being aware; such an app could be used by stalkers or abusive partners with physical access to a victim's phone.

4.3.2 Data and Entity Categories

After participants read the scenario, they were asked to categorize each type of data and third-party entity with whom the data would be shared, based on the NTIA MSHP short-form terms. We presented the categories using the exact same wording, in the same order, as used in the NTIA MSHP draft, published April 29, 2013.[3] We also added "None of the Above" and "Not Sure" options.

The NTIA provides both names and explanatory text for each category. In order to gain a better understanding of the utility of including this explanatory text, we conducted our study as a between-subjects survey. Participants in the *terms only* condition were shown only the category names in each scenario; participants in the *parentheticals* condition were also shown the NTIA's explanatory text for each category.

[3] **http://www.ntia.doc.gov/files/ntia/publications/mobileappdraftapril29_2013_draft1b_fs.pdf**

We designed our online survey after conducting eight in-person pilot tests, in which the survey-taker walked through the survey with the researcher and thought out loud. These pilots allowed us to refine our study design. For example, in these pilot surveys, we found that participants were skeptical about the scenarios giving them complete information about what data would be shared. As a result, they were apt to make inferences about additional types of data that might be shared. Therefore, we designed the survey so that participants would select a data or entity option only for elements mentioned explicitly in the scenario. Furthermore, we added a notice on every page stating, "The scenarios describe the data collection and sharing completely, so you do not need to guess anything outside of what is described." We also included two open-ended questions that were used as an attention check for quality results.

4.3.3 Data Analysis

Each of our participants was shown a sequence of ten scenarios; each scenario had at least one data item and at least one third-party entity with whom data is shared. Participants were asked to classify each data item and each entity according to the NTIA categories, or as "None of the Above" or "Not Sure." In all, participants were asked to make 52 categorizations. The data type items we asked participants to categorize are shown in the second column of Table 4.1; the third-party entities are shown in Table 4.2.

We cannot determine how many of our participants were "correct" in each scenario, because there is no ground-truth on which to base that assessment. This is the result of the stake-holder process, in which there were concessions but not necessarily agreement on the terms and their meanings. Thus, there is no way to determine whether a given response is inherently correct or incorrect. Given this lack of general correctness, instead our analysis focuses on how consistently our participants categorized the data items and entities. For each data item and entity, we considered the most-commonly selected category to be the *winner*. We then looked at the percentage of participants who selected the *winning* category for each data item and entity, and we call this percentage the *common understanding* for that data item or entity.

We classify each data item and each entity as being either *low common understanding* or *high common understanding*. A data item or entity in which more than 60% of our participants agreed on its categorization is considered to be *high common understanding* (that is, more than 60% of participants categorized it as its *winning* categorization). A data item or entity with 60% or lower categorization agreement is considered to be *low common understanding*.

4.4 Study Results

Our study found that participants and the NTIA experts had a low common understanding of many of the terms used in the NTIA MSHP notice. We begin with a description of our participants and then summarize our main findings. Detailed results of the study can be found in Tables 4.1 and 4.2. Breakdowns of how participants voted for each element can be found in Appendix B.

In the expert columns of Tables 4.1 and 4.2, we show all categories selected by two or more experts, with the number of experts who selected each category in parentheses. The terms in which the majority of experts and participants differed are in italics. If the conditions in the participant study had different winners, both are shown in the participant column. Some categories were

Scenario	Data	Expert Response	Winning Participant Response	Parenthetical[1]	Term only[1]	p-value
HipClothes	Inseam	Biometrics (2)	Biometrics	69.1	45.9	<.001*
	Waist Size	Biometrics (2)	Biometrics	69.6	46.4	<.001*
	Clothing Preference	None (3)	None	48	38	<.001*
	Location	Location (4)	Location	91.7	89.9	.494
Salsa	Call History	Browser History (4)	Browser History	88.5	87.5	.463
	Text History	Browser History (4)	Browser History	89.3	90.1	.184
	Video History	Browser History (4)	Browser History	51.5	70	<.001*
	Games Played	Browser History (3)	Browser History	45.9	50.5	.021*
	Photos	User files (3)	User Files	77.6	69.2	.005*
SuperTax	Photo of W2	Financial Info (3)	User Files	59.2	75.5	.001*
	Salary	Financial Info (4)	Financial Info	92.3	93.3	.502
	Interest Income	Financial Info (4)	Financial Info	92.5	91.8	.066
Fitness	Steps Taken	Health (2)	Biometrics	40.3	46.2	.225
	How Long Slept	*Health (4)*	Biometrics	39.7	44.2	.148
	Weight	*Health (4)*	Biometrics/Health	54.1	50.2	<.001*
	Body Fat	*Health (4)*	Biometrics/Health	53.3	49.5	.005*
EasyApply	Work History	*None (3)*	None /Financial Info	33.3	34.4	<.001*
	Medical Insurance	Health (3)	Health	85.9	81	.161
	Medical Payments	Health (4)	Health	59.7	52.2	.127
	Number of Children	None (3)	None	41.1	35.1	<.001*
	Marital Status	None (3)	None	43.5	35.1	<.001*
	Income	Financial Info (4)	Financial Info	88.5	91.6	.063
CallCalendar	Call Time	Browser History (4)	Browser History	91.2	86.8	.222
	Call Duration	Browser History (4)	Browser History	90.1	86.3	.189
	Name from Contact List	Contacts (3)	Contacts	71.2	82.5	<.001*
GoodDriver	GPS Location	Location (4)	Location	94.1	94.7	.788
	Gyroscope Bumps	None (3)	None	33.6	33.9	.252
FindMyKid	Location	Location (4)	Location	94.1	94.7	.176
iTunes	Credit Car Info	Financial Info(3)	Financial Info	96	92.3	.304
	Song and Artist Names	*None (3)*	User Files	57.1	53.1	.443
Bookstore	Book Title	None (4)	None	34.4	36.1	.502
	Home Address	None (2)	Location	49.1	58.7	.008*
	Credit Card	Financial Info(4)	Financial Info	94.1	91.1	.092

[1] Participant level of common understanding for winning term by condition (% who selected the winning participant response).

* Difference between conditions is significant at $p<.05$ with χ^2 test Benjamini and Hochberg FDR correction.

Table 4.1: Data Type categories selected for each term by NTIA experts and MTurk participants.

Scenario	Data	Expert Response	Winning Participant Response	Paren-thetical[1]	Term only[1]	p-value
HipClothes	OtherClothing Stores	*None (3)*	Consumer Data Reseller /None	31.5	33.3	<.001*
Salsa	Advertising Companies	Ad Networks (4)	Ad Networks	80.5	79.2	.520
	AdmeMetric[2]	*Consumer Data Reseller (3)*	Consumer Data Reseller	43.8	38	.086
SuperTax	State Agency	Govt. Entity (4)	Govt. Entity	93.9	96.2	.465
	Federal Agency	Govt. Entity (4)	Govt. Entity	94.7	95.4	.518
Fitness	Sports Companies[2]	*None (3)*	Consumer Data Reseller	38.4	26.8	.027
	Health Companies[2]	*None (3)*	Consumer Data Reseller	31.5	24.6	.022*
EasyApply	State Agency	Govt. Entity (4)	Govt. Entity	92	93.3	.208
CallCalendar	Carrier	Carrier (4)	Carrier	90	88.2	.173
	Google Calendar	Other Apps (3)	Other Apps	47.1	51	.066
GoodDriver	Traffic Data Company	None (2)	Data Analytics	59.7	58.4	.770
	Car Insurance	*None (4)*	Consumer Data Reseller	35.7	26	<.001*
	Car Rental	*None (4)*	Consumer Data Reseller	36.3	25.7	<.001*
FindMyKid	Parents Phone	None (3)	None	34.4	46.6	.034
	Local Police	Govt. Entity (4)	Govt. Entity	80	85.3	.333
iTunes	Facebook	Social Network (3)	Social Network	89.6	92.1	.714
	Apple iCloud	OS and Platforms (2), None (2)	OS and Platforms	37.9	34.9	.799
Bookstore	Facebook	Social Network (3)	Social Networks	88.8	90.6	.566
	GreatReading	Social Network (2), Other Apps (2)	Other Apps	37.6	40.1	.410

[1] Participant level of common understanding for winning term by condition (% who selected the winning participant response).
[2] 288 Responses Only
* Difference between conditions is significant at p<.05 with χ^2 test and Benjamini and Hochberg FDR correction.

Table 4.2: Third-Party Entities categories selected for each term by NTIA experts and MTurk participants.

abbreviated in the tables as follows: "Health, Medical or Therapy Information" has been abbreviated to "Health", "Browser History and Phone or Text Log" to "Browser History", "Financial Information" to "Financial Info", and "Government Entity" to "Govt Entity."

4.4.1 Participants

The four NTIA MSHP participants in our study, whom we call our *expert participants*, were a diverse group. They each held different professions and represented different stakeholders in the NTIA process; we do not report their demographics to preserve their anonymity. Expert participants were evenly split between our two conditions; because we had only two expert participants in both conditions, we do not report differences based on these conditions for expert participants.

For our MTurk participants, we analyzed data only for participants in the United States who had completed the survey, and we excluded participants who entered gibberish answers for open-text fields that were used as an attention check. This left us with 791 MTurk participants (375 parenthetical and 416 term-only). The data was collected in two batches, one of 503 responses and one of 288 responses. The second batch included three data entities accidentally omitted from the first. The data entities were: Sports and Health Companies in the Fitness scenario and AdMeMetric in the Salsa scenario; these are indicated in Table 4.2. We combine the results from these two batches, except when discussing the three questions that had only 288 responses.

51% of the MTurk respondents were female. Participants ranged in age from 18 to 73 years, with a mean of 33 and a standard deviation of 11 years. Participants took an average of 17 (median 15) minutes to complete the survey. Every US State was represented. Participants were generally educated: 38% have a Bachelors degree, and another 30% have some college. 82% own a smartphone.

4.4.2 Summary of User Study Findings

The NTIA MSHP has selected several categories of data sharing about which mobile users should be informed on short-form privacy notices. Our investigation looked at user and expert understanding of these categories. Our survey found that the categories were not well understood by our participants. Of the 52 examples of data sharing given in our scenarios, participants showed low (less than 60%) common agreement for 23 of them. Furthermore, our expert participants also disagreed among themselves on how to categorize some of the examples, and had different majority responses from the study participants for 13 examples. We find that the *Biometrics* and *Health, Medical or Therapy Information* categories were especially prone to disagreement. Further, participants struggled to categorize many of the third-party entities.

The main finding of this study is that the current set of NTIA categories does not appear to offer a high level of transparency for users. The lack of common understanding, even among experts, also suggests that app developers may have trouble generating accurate notices using these terms and definitions. Next, we will discuss our main findings and our recommendations.

Parentheticals Help (Sometimes). In most cases, the difference between the parenthetical condition and the term-only condition was not significant. When it was significant, the parenthetical usually resulted in greater agreement with the most-popular category. However, this was not always the case; some parentheticals appeared to confuse our participants. For example, the parenthetical text for *Browser History and Phone or Text Log*, *User File*, and *Location* appear to need some improvement to make them more useful to users.

Better Definitions Are Needed. Some categories were not well understood, either by participants or by NTIA experts. Therefore, we recommend that the Code provide further guidance on how to interpret the categories. This may include definitions and examples, including edge

cases. In particular, guidance is needed for all of the third-party entities except *Government Entities*, as well as the categories *Biometrics* and *Health, Medical or Therapy Information*. Further, experts should clarify whether location includes only information from sensors (such as GPS) or user-entered information (such as home address).

Ambiguous Data Items Need Clarification. Several types of data items were confusing to participants. Some data items could reasonably be classified in two categories (e.g., a photo of a W-2 is both a user file and financial information). This typically resulted in low common understanding. Furthermore, the Code of Conduct does not specify whether both categories should be listed or how one should be chosen. Some data items require an understanding of the platform architecture in order to classify them correctly (e.g., whether a contact name is stored in a call log or in a user file). As a result, app developers may correctly categorize a data type, but users may not understand the categorization.

In several cases, participants who saw the parenthetical text had less agreement than those who saw only the terms, indicating that the short phrases created confusion instead of clarification. In the case of home address, participants who saw the parenthetical were less likely to select *Location*, and were more likely to say *Not Sure* or *None*. In the case of video history, users who saw the parenthetical text may have been attracted to the word "video" in the *User File* description, and therefore choose that category over *Browser History and Phone or Text Log*.

For improved transparency on ambiguous or poorly understood data types, we recommend that implementors of the short-form specify the data being collected. For example, a short form notice with the text "Health, Medical or Therapy Info: how many steps you have taken, how long you've slept, weight, and body fat" may be more clear to users than "Health, Medical or Therapy Info." The specificity would alleviate the problems described above with ambiguous data types. Future research should investigate whether specific information is better understood, and whether implementors of a short-form notice should specifically say what is being collected instead of, or in addition to, the parenthetical text.

Third-Party Entities Are Poorly Understood. Many of the third-party entity categories were confusing to participants. Our results show that participants struggled with many of the third-party entities, except *Government* and *Carriers*. In particular, participants categorized six entities as *Consumer Data Resellers* while the experts only categorized one as such, typically choosing *None* instead. It may be that participants used this as a fallback choice for entities they didn't understand, while the experts had a much narrower definition in mind.

On the other hand, specificity about third-party entities will only be helpful if users recognize the name of the entity. Previous research suggests that users are not familiar with the names of advertisers, data resellers, or analytics companies [35, 115]. Further research is needed on describing third-party entities in a transparent way.

Uncategorized Data and Entities. There are some privacy-sensitive data that do not fit into any of the existing categories (and therefore need not be indicated in a short-form notice). These include identifying information such as user name, phone id, or SSN. Since not all data sharing falls into a category covered by the short-form notice requirements, the app may be sharing data without notifying the user through the short form. Our results show that participants did not often categorize data and entities as *None*, and preferred to place data in one of the categories. This suggests participants believe the categories encompass all possibilities. Therefore, information about the smartphone notices should emphasize that the short form does not notify users about all

types of data sharing.

Further User Testing is Needed. Our study is a concrete first step which indicates that more work is needed to develop a well-understood notice with categories and definitions that will be generally understood by American smartphone users. By providing realistic scenarios and asking survey participants to categorize data items and entities with whom data is shared, our work highlights that the categories are not well understood. However, this is not a typical task flow for users, and we did not test actual short-form notices. Nonetheless, if the NTIA MSHP had adopted a similar approach of using case studies to understand and categorize data sharing, it is likely they would have developed more understandable terms and definitions. In fact, when similar examples were raised in meetings the group moved on without reaching a consensus.

4.4.3 Demographic Factors

We looked at whether any of the demographic factors significantly affected participants' responses. We also examined at whether owning a cellphone, education, or knowing a programming language had any affect on choices for each item (χ^2 test with the conservative p $<$0.001 due to multiple tests). Only one term showed a significant difference among these factors: Education made a difference in how participants categorized *local police*. Participants with lower education (some high school or high school), more frequently selected *none* as the appropriate category for local police, while higher-education participants were more likely to recognize that they are *Government Entities*.

4.4.4 What Categories Are the Most Sensitive

We asked participants, "Which of the following types of data would you want to know about an app collecting?" and "Which of the following entities would you want to know if an app shared data with?" The response options were, "Want To Know," "Don't Care," and "It Depends." The response options were randomized between participants to avoid bias. In this analysis, we look at responses from participants who provided exactly one response to the question.

Among the types of data about which we asked participants, participants most wanted to know when "Financial Information" (89.5%) and "Health, Medical, or Therapy Information" (86.1%) was disclosed. The results for each datum are shown in Table 4.3.

Among the entities about which we asked, participants most wanted to know about data sharing with government entities (79.7%), followed by consumer data resellers (77.4%). For each of the entities about which we asked, over half of participants wanted to know when data would be shared with that entity. This is shown in Table 4.4.

It is worth noting that, while some entities and information appear more sensitive than others, over half of participants want to know about disclosure in each of the cases asked about. We also examined whether responses differed by condition or demographics. We examined the responses of our participants across three factors (omitting participants who did not indicate one): participant gender, the condition to which the participant was assigned, and whether the participant indicated using a mobile device. Separately for each of these three factors, we compared whether there was a significant difference in the proportion of participants who responded "Want To Know" to each entity and datum using a χ^2 test. All significance levels, separately for each factor, were corrected using Holm-Bonferroni correction.

Table 4.3: The percentage of participants who responded with "Want To Know" to the question "Which of the following types of data would you want to know about an app collecting?" for each entity.

Entity	% want to know
Financial Information	89.5%
Health, Medical or Therapy.Information	86.1%
Browser History and Phone or Text Log	82.5%
User Files	80.0 %
Contacts	79.8%
Location	71.1%
Biometrics	68.6%

Table 4.4: The percentage of participants who responded with "Want To Know" to the question "Which of the following entities would you want to know if an app shared data with?" for each entity.

Entity	% want to know
Government Entities	79.65%
Consumer Data Resellers	77.37%
Social Networks	74.97%
Ad Networks	72.31%
Data Analytics Providers	69.03%
Carriers	65.61%
Other Apps	63.34%
Operating Systems and Platforms	58.15%

We found no significant difference in response to entities or data when we look at participants by gender or by condition ($p > .05$). However, even with correction, if we compare participants who do and do not have a mobile device, we see a significant difference ($p < .05$) in wanting to know about disclosure to operating system and platform ($p = .002$), ad networks ($p = .008$), carriers ($p = .015$), consumer data resellers ($p = .012$), other apps ($p < .001$), and social networks ($p = .022$). In each case, participants who did not use mobile devices were significantly more likely to want to know about disclosure than those who do use mobile devices. We cannot determine, based on this data, whether users who are more privacy-sensitive are less likely to use a mobile device, or whether using a mobile device makes users less privacy-sensitive.

4.5 Limitations

This survey is designed to measure whether participants understand the NTIA categories by giving them an explanation of an app, and an explanation of the data shared, including such details as

Table 4.5: Third Party Entities with disagreement between participants and experts

Term	Expert	Participants
AdMeMetric	Consumer Data Reseller	Advertiser
Other Clothing Stores	None	Consumer Data Reseller
Sports Companies	None	Consumer Data Reseller
Health Companies	None	Consumer Data Reseller
Car Insurance	None	Consumer Data Reseller
Car Rental	None	Consumer Data Reseller

with whom the data is shared and the purpose of sharing the data. Participants may see more information than they would in practice. Our results for understanding, therefore, may be an overestimate of true understanding in practice. Further, as stated above, while we can measure the extent to which participants agree on how to categorize a given data item or entity, it is impossible to determine whether that categorization is "correct."

The task presented to survey participants more closely resembles a realistic task for an app developer than a user. A more realistic user task might be to provide a notice that uses the terms from the Code and to ask users what data they think an app is collecting and with what entities they believe it is shared. However, this is actually an even harder task because each data category could potentially cover many types of data, and it is not necessarily possible to infer what data is collected from a very brief description of an app.

This survey is limited to testing the particular terminology defined by the NTIA code. While the results indicate some categories are poorly understood, we do not test alternate wordings. Therefore, we are unable to offer better terminology; that may be an area for future work.

Furthermore, while we tried to present a broad swath of scenarios, we could not create a study that would present all possible scenarios to participants. There may be many more types of data that are ambiguous to users, or examples that are more clear than those in this survey.

Our pay rate of $1 for a 17-minute survey was well above the mturk rates studied by Buhrmeister et al., which showed that the lower rates did not effect the quality of results [56]. However, it is possible that the low pay could have impacted the quality of results.

Although we did extensive in-person piloting, we were not able to pilot extensively the survey with Amazon mturk participants. This was due to our deadline of completing the work with enough time to inform the NTIA MSHP before the final meeting. As mentioned in the results section, the data was collected in two batches; the second batch included three entities that were not in the first batch. Due to time constraints, we did not discard the first set. A χ^2 test between the two batches found no significant differences for the other questions. Therefore, we have reason to believe that combining the two batches did not impact the results of the questions.

Our recommendations for usable privacy and security practitioners are based on one case study. Although they are drawn from several informal discussions with other participants (such as personal conversations over the phone), we do not present them as results from a qualitative study.

4.6 Discussion

We released a technical report of our work on July 17, 2013, one week before the final NTIA MSHP meeting on July 25th. This technical report showed that the terminology in the short-form notice was not well understood, and further research was needed. However, by this point, participants were ready to reach consensus on the Code of Conduct and conclude the project. Although our study was discussed by the group, by then it was too late in the process to influence the Code very deeply. That said, we believe our study did have an impact, as future discussions indicated that user studies were planned [138]. Unfortunately, the process has already concluded, and the Code of Conduct has been announced, without plans to reconvene or address the usability issues. We fear that despite the best intentions of the participants, this will lead either to adoption of a short-form notice that does not meet its goals due to usability issues, or to app developers finding the notice flawed and therefore not adopting the voluntary Code.

Here we provide lessons learned, particularly aimed at academics or experts on usability who believe that regulation of technology should consider users and the human element. These are based on our own experience and personal off-the-record discussions with several stakeholders. We distill our lessons for academics and usability experts who also wish to avoid policies with requirements that are known to be unusable. This is organized into two subsections. The first subsection describes specific issues that hindered the integration of usability studies into the NTIA MSHP. We describe these issues with the goal of illuminating some gaps between academic HCI experts and the policymakers. We then offer some concrete suggestions for usability experts who wish to participate in multi-stakeholder public-policy-making.

4.6.1 Issues that hindered usability testing

In this section we describe specific issues that occurred during the NTIA process that led to adopting a Code of Conduct that was not well understood by the users in our study.

Disagreement about what 'usabilty' is. The main issue with conducting usability studies was that the stakeholders did not agree on what 'usability' meant. Although the stakeholders often recognized the need for user studies, they had different opinions about what should be studied, how it should be done, and what the results would mean. This is largely a result of the multi-stakeholder process, in which different stakeholders had different objectives and priorities, based on their experiences and whom they were representing.

For example, some stakeholders representing app developers felt that if usability tests showed some users were concerned by the notice and therefore did not download an app, this indicated a failure in usability. In contrast, some consumer privacy advocates argued that the notices should lead consumers to refuse the data collection practices of apps and download fewer apps. This difference in opinion may be familiar to those who have worked on notifications in other areas, such as authentication or P3P.

Other debates about usability included the role of icons in the notice, and whether icons could stand alone, or with text, or whether icons should be allowed at all [14]. Another debate was about which entities and elements the app needed to show on the notice: either only those the app collects or shares with, or the entire list with an indication that some things are not shared. Some stakeholders felt that usability tests should address these questions, while others felt the questions

were not relevant or were not the primary issues of interest. Our study did not address these issues, but we agree they should be examined.

Cost of usability studies. Part of the delay in starting usability studies was that it was difficult to resolve who should pay for the studies, and who should carry them out. The NTIA process itself did not have a budget for usability studies, so in order to pay usability consultants or private usability firms, some stakeholders would need to volunteer the funds. Although a stakeholder volunteered to search for usability consultants and request prices, it was difficult to get an estimate without knowing what would be tested. Furthermore, it was not clearly defined whether financial contributions from stakeholders would give them more control over the tests. The final cost of our study was under one thousand dollars, not including the value of graduate students' time to implement the study, and was paid by our lab's funds.

Process fatigue. After a year, many of the stakeholders were eager to complete the project. Although fatigue with the process may have contributed positively to the ability to come to a consensus, it also meant that the stakeholders were not willing to wait for usability results. This may be a different perspective than that of academics, who are often willing to dive into an interesting problem for several years.

Everyone is expected to have a bias. It has been said that policy makers at the federal level expect everyone to have a position. For example, several participants in the NTIA MSHP felt that future processes should request that all participants submit position papers. Academics may feel they don't take a 'position'; they strive to be neutral and let the results of the research stand as facts that support their claims. However, in controversial areas such as privacy, academics should be prepared to describe their position. In our case, our position was that the Code of Conduct should be usable; that both smartphone device users and app developers should understand the notice and the terms used.

4.6.2 Recommendations

In this experience, we found that although the drafters of the Code of Conduct generally recognized the value of user studies, they were unable to implement those studies. Our independent user study confirmed that categorizations described in the Code were confusing and suggested that further user studies could help create a more understandable privacy notification.

Engage early. Our independent study confirmed that usability tests were needed. We recommend that other researchers who have the resources can and should conduct user-tests to inform public policy, as it may not happen otherwise. We do not recommend waiting until stakeholders come to agreement about what to test. With the benefit of hindsight, we should have run our tests sooner to inform the process at an earlier stage.

Furthermore, we released a technical report before the final NTIA MSHP meeting, so that the results of the study would be available for the participants. We did not wait for publication in an academic journal or conference, as this may have delayed the results beyond the point of impact.

Our technical report could have been more useful if we had included a one-page executive summary. This may have been more relevant and useful to stakeholders and journalists than a full-length paper for understanding the issues within their time constraints.

Impact versus incentives. We recognize that academics may have little incentive to put their resources toward such studies, which may have more value to policymakers or a working group than to academics or reviewers. Indeed, our attempts to publish this paper in an academic conference

were initially thwarted by reviewers who felt the results of the study did not make a significant contribution to the field. Ultimately, we refocused the paper as a case study before publication. However, we believe that engaging with public policy can help prevent requirements with poor usability from being written into regulation. This may increase the impact of our research, as a community, in the long run.

Acknowledgements

This research was funded in part by NSF grants DGE0903659, CNS1116934, and CNS1012763 and John and Claire Bertucci Fellowship.

Chapter 5

"Little Brothers Watching You:" Raising Awareness of Data Leaks on Smartphones

Users are concerned about protecting their privacy on smartphones. In a telephone survey of 1,203 US adults, most were as concerned about the privacy of data on their smartphone as on their home computers. The majority of these participants oppose practices in which applications collect their contacts. Forty-six percent felt that wireless providers who collected location should not store it at all, and an additional 28% thought it should be deleted within a year [157]. A separate telephone survey of 2,254 US adults found that 57% of all smartphone app users "have either uninstalled an app over concerns about having to share their personal information, or declined to install an app in the first place for similar reasons" [53]. These surveys show users have expectations of how their privacy-sensitive data should be treated on the smartphones.

Existing interfaces typically fail to make users aware of relevant aspects of data sharing, e.g., destination, frequency, and purpose of the sharing. Without this awareness, it is difficult for users to make informed and optimal decisions about data sharing from smartphones. For example, Android requires that users make decisions about granting data access permissions before they install an application. The user is asked to agree to data sharing before she is able to evaluate the benefits of the application itself. In comparison, the iPhone interface provides a dialog asking for permission to send location data or address book the first two times an application requests that information. These interfaces do not notify users of the frequency, destination, or purpose of data sharing. None of these systems provide overviews about the information leaving the phone so that users can compare applications and types of information sent in a clear summary. In a recent field study of 20 Android users, we found that participants were often surprised by apps' data collection in the background and the level of data sharing [105].

In this chapter we present a smartphone app, *Privacy Leaks* that aims to improve users' awareness of privacy leakages as they occur on an Android phone. We use a term *privacy leakages* when referring to privacy-sensitive data being transmitted off the smartphone by applications in a way that is unexpected by the user. The prototype smartphone app was built on the TaintDroid

This chapter is largely based on a paper [35] co-authored with Jaeyeon Jung, Wei Lu, Lorrie Cranor, and Carolyn Nguyen.

platform [71], informing users about the frequency and destination of data being shared by an application in two different ways: (a) a visualization of the amount and types of information shared, after the data has been shared; (b) just-in-time (JIT) notifications at the moment the information is shared. Using the prototype, this work explores the following three research questions:

- What are participants' pre-existing understandings of data sharing with smartphone applications?
- Can runtime feedback via notifications and visualizations of data sharing on smartphones reduce the gap between users' understanding and actual privacy leakages without creating annoyance or confusion?
- What design guidelines can be drawn from participant feedback to improve smartphone privacy interfaces?

To create a concrete context for data sharing, we used a role-playing technique in our 19-participant lab study. Participants were asked to play two popular smartphone games and select one to recommend to a friend or family member. This simple task was performed twice: first on a regular Android phone and second on a phone running our prototype. Through a semi-structured interview, we first examine participants' misconceptions about data sharing. We then examine reactions to our interface and changes in understanding, and finally we look at desired control over data sharing.

This chapter makes two contributions. First, we find that some participants have a very limited understanding of data sharing by smartphone applications, yet have a strong desire to remain anonymous or to protect their children from potential harms. Without any consumer education or interfaces raising their awareness of privacy risks, these users would be left vulnerable. Second, we provide design guidelines to improve users' understanding of privacy leakages through just-in-time notifications and a summary visualization on the phone. However, improved awareness is only the first step toward helping smartphone users reduce privacy risks. We identify future research efforts to provide users with control over their data.

5.1 Related Work

We first discuss prior work that explored users' understanding—or lack thereof —of privacy and security risks of smartphone applications. We then describe work that designed tools to inform users about various security and privacy issues and to provide control over their data. We highlight how these previous studies influenced the design of our study method and the Privacy Leaks prototype.

5.1.1 User Understanding of Privacy & Security Risks of Smartphone Applications

Several studies demonstrate a lack of user understanding of privacy and security risks associated with installing smartphone applications. An Internet survey of 308 Android users and a laboratory study of 25 Android users found that only 17% paid attention to the permissions (including ones which grant an application access to privacy-sensitive data) when installing an application. They also found that only 3% of the Internet survey respondents demonstrated full comprehension of the

permissions screen [76]. Kelley et al. reported that Android users found it difficult to understand the terms and wording of the Android permissions [107]. Our study goes deeper into this lack of understanding and discusses users' misconceptions about data sharing with two popular game applications using a role-play technique.

To examine expectations and perceptions of smartphone security, Chin et al. interviewed and surveyed 60 users on how they would choose applications to install on their smartphones. They found that few participants considered the privacy policies when deciding which apps to install. Instead, referrals from friends or family, or on-line referrals were the predominant ways that users discovered new applications for their smartphones. Price, popularity, and recommendations from friends were important parts of the decision about whether to install [59]. Recognizing the value of recommendations, we designed our interview so that users were asked to make a recommendation to a friend or family member.

Lin et al. used crowd-sourcing to analyze users' "expectations of apps' access to phone resources." The crowd-sourced expectations were used to design a new privacy summary interface for installation, which was more easily understood and efficient than the existing interface [122]. Our study exposed users to the sharing of location data and phone identifier by two popular games, Toss It and Angry Birds, which are unexpected uses of data according to the crowd-sourced results.

5.1.2 Designing Control Over Data Leaks

Many of the pioneering studies on designing usable privacy notices are described in Chapter 1.4.5. In this section we additionally describe some additional tools that have been developed to help users control the data sharing from smartphones.

AppFence is an automated system for allowing privacy control, which can provide fake or filtered data to applications and prevent the sending of sensitive data [94]. The authors tested 50 popular Android Market applications using an automated testing methodology and found that AppFence reduced data leakage without side effects on two thirds of the applications. However, they found trade-offs between usability and privacy on the remaining applications.

Zhou et al. developed an application for Android with a privacy mode. This application, which requires modifications to the Android framework, allows users to set fine-grained privacy settings about whether an application will have access to real, spoofed, anonymous, or empty data. The authors did not address users' understandings of the settings or dialogs [165]. This work along with similar studies [94] helped us design some of interview questions on privacy control in order to determine whether the proposed interfaces match users' expectations and desires.

5.2 Designing Privacy Leaks

Our prototype is built over TaintDroid [71], which instruments the Android operating system to determine whether privacy-sensitive data is being transmitted off the phone by Android applications. Our prototype reads data transmission events generated by TaintDroid. These events both trigger a notification and are written to a database, where the event information can be accessed later for the visualizations. However, in general, not all data transmissions are deemed unexpected by users (e.g., location data being sent off to a map server when the user is using a navigation

app) and there needs to be a filter that can differentiate privacy leakages from legitimate data transmissions. We discuss how such a filter can be implemented in Chapter 5.6.

We used an iterative process to design the notifications and the layout for our visualization interface. This process included iterating over several designs by testing paper mockups and Android prototypes on colleagues who are not privacy experts.

We named our app *Privacy Leaks*, which may have led to a slight user bias about the information. Although it is consistent with our previous definition of leakage, this title may have negative implications about data sharing. We realized this after the study had been completed.

5.2.1 Notifications

The just-in-time notifications were intended to notify users at the moment data was being sent. In our prototype, the phone both vibrated and made a water-drop sound when privacy-sensitive data had been transmitted off the phone.

While we attempted to build a unique vibration, it is not clear whether the users would have been able to distinguish our vibrations from vibrations caused by an application. However, as shown in Figure 5.1, our prototype also included an icon and short text notice in the notification area, so users can check out the phone's status bar to see the source of vibration (e.g., Privacy Leaks notifications vs. text message arrivals).

5.2.2 Visualization

The visualization allowed users to compare, across apps, what information had been detected by TaintDroid as being shared recently. This type of visualization can be examined after an app has been used to see what was shared and would require the user to actively open Privacy Leaks to view the information.

We focused on a simple layout that could quickly give users a sense of shared information without using jargon or requiring technical knowledge of Android. Through our iterative design process, we selected which information to show and how to display it. We used a grid layout (similar to [106]) to show an overview of the data that had been leaked. The columns showed the type of data, and the cells in each grid showed the number of times the information was sent. The cells were shown in red that became progressively brighter as the number of times increased.

The main visualization (e.g., Figure 5.2) shows data leaked by all applications over a period of time; this period is configurable by the user.

Our prototype included a jargon-free one-sentence description of the information: "How many times did Apps leak information about you since [timestamp]?" The rows include the application icon to help the user easily identify the application. The columns are the permission-based fields that are sent. We created a second screen, seen in Figure 5.2, to show the destinations of the data for the individual applications, available by clicking on the application icon.

Due to limited screen-space, we were not able to display a column for every type of data that could be shared. Therefore, we made the following design decisions to choose which columns to show. Three columns are always shown. Other types of data are shown in additional columns only that appear if that data type has been sent. In particular, we always show Location, Phone ID, and Phone #. Location and Phone ID are the two most frequent types of leaks [71]. We also always included the field Phone # to clarify that Phone ID is not the phone number. Phone ID can be used

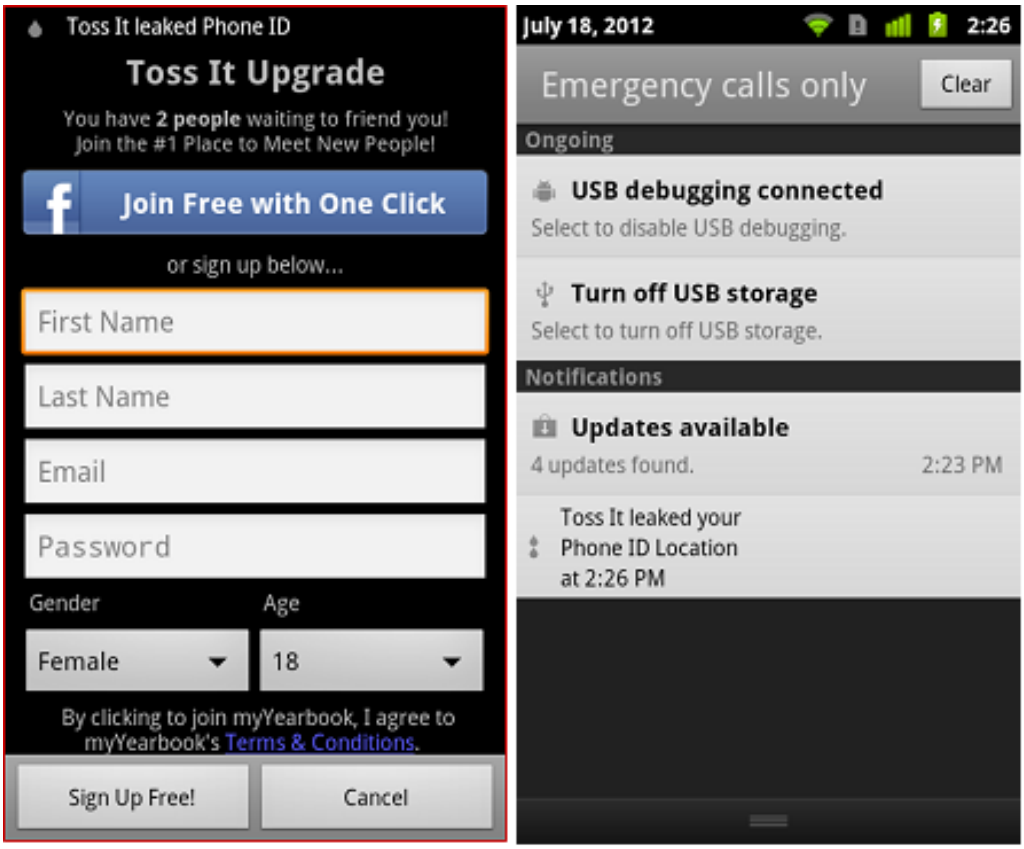

Figure 5.1: Notifications in the status bar (left) and in the notification drawer (right) by Privacy Leaks

to uniquely identify the phone. As stated in in the Privacy Rights Clearinghouse Fact Sheet, "The privacy concern here is that information could be shared with third parties and compiled with other data to create a detailed profile about you without your knowledge or consent" [12]. Location can be used to locate the phone, and Phone # can be used to identify and to call the phone. Other privacy-sensitive columns, such as Address Book, appear if and when an application sends off that information.

During the design process, we found that users were confused by the different types of phone identifiers such as: "IMEI," "IMSI," "SIM card identifier," "Device serial number," and "Android Id." We renamed and collapsed these to a single group, "Phone ID," to avoid overwhelming jargon. Similarly, we did not distinguish between types of location data: we collapsed "Last-known Location," "NET-based Location," and "GPS Location" into "Location."

Furthermore, we did not show which location was sent, such as the exact GPS coordinates. Nor did we show a timeline of when information was sent. Our paper mockups of such visualizations were not well received, but we believe they are both feasible visualizations and we are considering them for future work.

Applications may also send parameters along with the above privacy-sensitive fields. Understanding this data often requires technical knowledge of the application. Therefore, we did not show this information out of concern that it would overwhelm or confuse users.

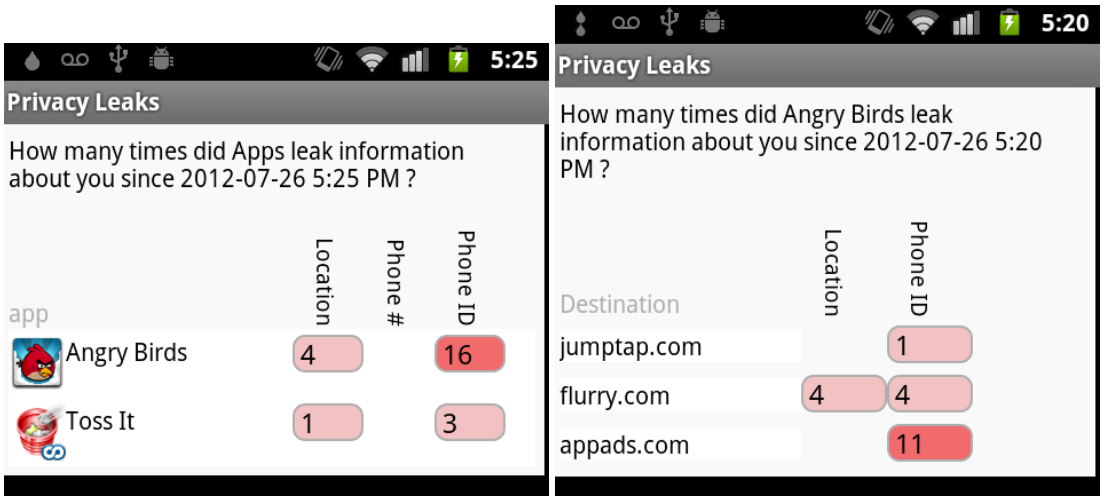

Figure 5.2: Main visualization screen of Privacy Leaks (left) and Application detail screen of Privacy Leaks (right)

Our prototype was also instrumented to allow configuration for research purposes, including configuring the time frame, refreshing data, turning off notifications, and exporting data. However, participants were not expected to use these options and the usability of the configuration settings was not a part of our user study.

After the design iteration, we noticed that the grid is somewhat similar to the Wall Street Journal's visualization of data sharing[1].

5.3 Study Methodology

We conducted a lab study of 19 participants in July and August 2013 to investigate their existing understanding of potential privacy leakages while using smartphone applications and to collect initial feedback on our Privacy Leaks prototype. We interviewed each participant for up to an hour in the lab. Interviews were structured in the following order: 1) the participant plays two games without Privacy Leaks and answers questions about the games and data being sent off the phone, 2) the participant plays the same games with Privacy Leaks and answers the same questions as before, 3) the participant is interviewed about data control, the usability of Privacy Leaks, and perceptions of desired data sharing. We explain each part of the interview in further detail in the next section.

5.3.1 Study Procedures

The first part of the interview served as a control to gauge the participants' impression of the games and examine their knowledge of data leakage. After arriving and being briefed about the study, participants were given an Android phone that had applications pre-installed. Then, they were asked to play and compare two games: Angry Birds and Toss It. Participants were provided with

[1] `http://blogs.wsj.com/wtk-mobile/`

copies of paper screenshots of the install process of both games, including the permissions screen. They had up to 7 minutes total to play and compare the two games. They were asked to evaluate the two games in order to recommend one to a friend or family member. This is somewhat longer than a typical session with an application of less than a minute [49]. However, participants were asked to think out loud and evaluate the applications, which typically lengthens the time to finish a task.

The participants selected a specific friend or family member, and then the researchers used that relationship (e.g., "wife" or "colleague") in all further questions to help the users create a specific and realistic scenario. Thirteen users selected a family member, such as their wife or nephew. Five participants selected a minor: for example, a child or younger sibling. The participants were asked to imagine that the friend or family member would be playing the game "on the bus, at the doctor's office, or waiting to meet you somewhere."

id	sex	age	condition	data sent
1	m	32	notifications	33
2	m	28	notifications	28
3	f	39	notifications	29
4	f	49	visualization only	21
5	f	43	visualization only	37
6	m	52	notifications	31
7	m	44	visualization only	23
8	m	23	visualization only	21
9	f	38	visualization only	29
10	f	21	notifications	18
11	f	43	notifications	14
12	m	38	visualization only	32
13	f	37	visualization only	17
14	m	38	notifications	25
15	f	37	visualization only	29
16	m	28	visualization only	38
17	m	26	notifications	36
18	f	44	notifications	63
19	f	20	notifications	17

Table 5.1: Participants' demographics, condition, and number of times data was sent off the phone while they used Privacy Leaks

After the first round of play, participants were asked to describe their recommendations of the games to their friends, and how they would describe the games on the app market. They were then asked about the information that was leaving the phones while they tested the games, and why the data was leaving and where it was going. This allowed us to evaluate their existing awareness and understanding of data sharing.

In the second part of the interview, participants were given a phone that was identical to that which used in the first part of the study, except it also had Privacy Leaks. The participants were told that an application was installed to notify them of data sharing, and that they had another 7 minutes to evaluate the same two games. After participants played the games, they were prompted to open Privacy Leaks to view the visualizations. Following this, participants were asked whether their recommendations changed, and were interviewed on their understanding and awareness of data leakages of the application. This allowed us to examine how users reacted to data sharing, and their revised understanding based on the information in Privacy Leaks.

The second part of the interview included the visualization-only and the JIT conditions. The interview questions about the games and the data sharing were the same across both the conditions, except for two differences. The participants who received JIT notifications were told, before the second part of the interview, "an application will inform you about information that is being shared

through notifications, such as vibration and the sound of water dropping." Also, participants in the JIT condition were asked additional Likert-scale questions in the third-part of the interview on whether they found the noise and vibration annoying or interruptive.

The third part of the interview consisted only of interview questions, and did not include game-playing or application use. Participants were asked to describe what they would do if they could "completely control the data leaving the phone." Participants were asked about specific elements of Privacy Leaks notifications and visualization, which allowed us to examine the usability and likability of the application. They were also asked about desired control over data leakage, and the risks and benefits of sharing data.

The interview was structured enough to allow comparisons between participants, but open enough to allow the researcher to probe about specific comments. The interview questions and instructions were the same across all participants. There were no written components of the interview—the questions were all asked and answered orally. The interview included a combination of open-ended, yes/no, and Likert-scale questions. Participants did have access to a printed copy of the Likert-scale values to refer to when answering the Likert-scale questions.

One researcher led all the interviews. One of two additional researchers took notes in the interviews. The interview was audio-recorded. The results were coded iteratively based on both the notes of the two researchers who were present at the interview and the audio transcripts. Two researchers sought themes within the responses, and then coded the results based on the theme list. They then iteratively re-coded based on re-evaluating the responses and discussion until agreement was reached between the coders.

Our results are entirely qualitative. We include the number of participants who responded with certain themes or ideas, but we do not intend to imply statistical significance, or that this represents a larger population.

5.3.2 Game Features

The two games used in the study were decoys; we wanted participants' attention on the primary task of selecting the game, as opposed to thinking about privacy. Both games involved a simple flick gesture to send an object on a trajectory, aiming at either pigs or a trash can. Therefore, the games were similar in their simplicity and style. While the games themselves were not of particular importance to our study, both games had features that were important to participants' conceptions of data sharing. Neither game uses location for functionality, but both send location information to the game developers and third-parties.

Angry Birds showed a banner ad, as shown in Figure 5.3a that several participants remarked upon. Angry Birds sometimes shows a full-screen ad as well. Since recruited participants were already familiar with Angry Birds, several commented on the possibility of viewing ads, even if none were displayed while they played the game during the lab study.

Some participants recognized that data would be shared if the game was social or allowed score sharing. As seen in Figure 5.3b, Toss It had four buttons for social networking or social games, such as challenging another player, at the bottom of the play screen. These buttons took users to a screen asking them to log in with their Facebook account. None did so.

(a) Angry Birds

(b) Toss It

Figure 5.3: Screenshots of the games used in this lab study

5.3.3 Participants

Ten male and nine female participants were recruited from the Seattle Metropolitan area by a service agency for our study. Our intention was to get variety and diversity, not to represent the USA population statistically. Participants were compensated with a choice of software gratuity. We screened to exclude people with a computer science degree. The average age was 35, in a range of ages from 20 to 52 years. Seven participants had a bachelor's degree, while 6 had completed high school and 6 had an advanced degree. Table 5.1 includes details on the participants' demographics. All participants were current Android users for at least 3 months, and had installed and played Angry Birds before participating in the study. To avoid priming the participants in advance about privacy or data leakage, the participants were told that the study was about Android games.

5.4 Initial Understanding

Participants played both games for a total of 3-7 minutes in the first part of the interview before making a recommendation and describing the game. We then asked them what data had left the phone while they played the games. They were therefore given a specific situation in which to evaluate data sharing, that allowed us to examine the understanding of data leakage before viewing Privacy Leaks. They were asked about why, when, and what data left the phone, and whether both games shared information.

5.4.1 Purpose of Sharing

We found that participants' level of awareness about the data that was shared could be roughly categorized into three groups:

- Group 1: Five participants stated explicitly that they had never before thought about information leaving the phone.
- Group 2: Eight participants believed that data was shared only with application developers for the purpose of improving the application.
- Group 3: Six participants understood that data was used for marketing but were surprised by the scope of data sharing, including the frequency of data sharing and the destination of data.

While the degree of awareness was different, none of the participants entered with a complete understanding of data sharing and the scope.

Participants belonging to Group 1 had never thought about data sharing before. P4 expressed her uncertainty about whether data left the phone, "Maybe it's not [leaving]. Maybe it's all in the phone. That's a tricky question. I don't know. Does it leave it?" P3's comments represent the idea that the game is self-contained, "It was my understanding once you downloaded it to your phone, it's on your phone. It didn't need to communicate with anything." In the first part of the interview, the participants were prompted with several open-ended questions about where, when, and why data was leaving, but they were often unable to answer. Several of these participants adopted a new understanding as they pondered our questions and thought out loud. These new ways of understandings fit into the next two categories.

Participants belonging to Group 2 believed the application is a self-contained environment. For example, P5 said, "If I'm within the Rovio game I'm thinking it [data] goes to Rovio. I didn't think if I'm within the application environment [data is leaving the phone]." Some participants commented on social networking. For example, P2 said, "Toss It [would share] with online communities if I had continued to start a challenge. Other than that it's not sending anything." P19 said, "Information is useful for analyzing the product. They can customize the game based on where and how long the game is played. I think it is about knowing the market." These participants were not aware that data was shared for the purpose of marketing, and thought their level or skill was sent in order to improve the game.

Participants belonging to Group 3 were aware of targeted advertising integrated with smartphone applications. However, even those who mentioned targeted ads were still confused about the mechanisms. P6, an older participant who had seen an ad for insurance in Angry Birds, stated that data was being shared for marketing. He said, "It didn't ask for age, education, doesn't know who is playing, but it might have email. A ten-year old wouldn't receive an ad for insurance," indicating an understanding that targeted ads could exist, but not sure how they would get enough information to target him.

5.4.2 Additional Perspectives on Data Sharing

Seven participants referred to the existence of terms and services but were not clear on what was included in these terms. For example, P18 said, "We give them all these permissions," when

referring to what data left the phone but she wasn't specific about what the terms were. P11 expressed uncertainty while correctly summarizing the situation "Does it need to ask permissions? I think it asks something when you download it. I guess you can't download it without allowing it." Only a few of the participants examined the printouts of the install screenshots that were on the table in front of them to find out what information was being shared. This suggests that even though users are aware of the permissions requests, they rarely see them as a resource for understanding data sharing.

P8 thought that data moves in a cycle with continuously coming and going. "Data can't always be stored in memory. It is in-going and out-going." In this (incorrect) perspective, data is shared because the limited memory space on the phone pushes the data out to remote servers to use them as temporary storage.

Overall, all nineteen participants had a limited understanding of whether and how often these smartphone games may collect the user's privacy-sensitive data. Next, we analyze participants' response after repeating the same task with Privacy Leaks.

5.5 Early Experiences with Privacy Leaks

This section discusses the participants' reactions to our Privacy Leaks prototype in the second part of the interview, after having viewed the visualization shown in Figure 5.2 and Figure 5.2. Ten participants in the JIT condition felt and heard JIT notifications in addition to seeing the visualizations. Participants were prompted with the same open-ended questions about where, when, and why data was leaving as in the first part of the interview, in order to gage the difference in understanding after using Privacy Leaks. We discuss overall reactions to our prototype implementation, including which new information violated the participants' initial understanding of data sharing.

5.5.1 Surprised by Actual Data Leakage

Across all groups, participants were most surprised by the frequency and destinations of the data. Usually, the information that was new, and did not fit into their previous understanding, was the most surprising.

Many participants were surprised by the frequency of data sharing, regardless of their initial perspective. However, participants belonging to Group 1 were very surprised by the frequency. They struggled to understand why the data was sent multiple times in the short time span they played the games. P3 said, "Why does it need to say what my Phone ID was more than once?"

Participants belonging to Group 2 were typically most surprised by the unrecognized destination URLs (e.g., flurry.com or admob.com). P1 expressed this concern about not knowing where the data was going: "Destination is surprising; that is a little concerning. It would be nice to have some sense of who is collecting the information." Participants were sometimes able to make assumptions about the destinations upon examining the URLs, as some have "ad" in their name. P7 said, "I've never heard of any of these companies. I assume they are using it for marketing." P10 had a similar comment while looking at the list of destination URLs, "Are those things supposed to mean anything to me? Oh. It's all advertisers."

Participants belonging to Group 3 were typically most surprised by the number of different destinations *and* the frequency of sharing. P19 expressed her anger that the game was sharing the

data with many companies, "I find Toss It slime as they let other companies collect information." She continued, "My eyes have been opened today. Every time you use the phone, every time you download an application [it] is not big brother watching you, but a lot of little brothers watching you. And they want to sell something to you."

5.5.2 Opinions of Privacy Leaks

Playing the games for 3-7 minutes was sufficient for participants to experience notifications and build up a history of data sharing to see on the visualization. Data was shared an average of 29 times per participant. As participants played each game for different amounts of times and accessed different parts of the games, the amount and types of information shared varied.

Overall, we found that participants liked our Privacy Leaks prototype and would want to install a similar app on their phone. Sixteen participants agreed or strongly agreed that "the information provided by Privacy Leaks is useful," and fifteen agreed or strongly agreed with the statement "I am likely to install an application like Privacy Leaks." For more information on responses, a histogram of responses to these particular questions is shown in Figure C.1 in Appendix C. Unfortunately, the Privacy Leaks app would only provide useful information on a special Android phone that is instrumented to run TaintDroid. Interested users can build and flash their Android phone with the TaintDroid system image following the instructions available in `http://appanalysis.org`. However, this additional step can be a substantial barrier for deployment.

A number of participants indicated their desire to install the Privacy Leaks tool immediately, and asked when it would be available on the application market. P2 said Privacy Leaks is "a great asset to have on your phone. It gives you information about where your data is going so you can choose [apps] more wisely." Most participants disagreed (3) or strongly disagreed (13) with the statement, "The information was irrelevant." P11 described the interface as "a good app for a person who is curious about data sharing", indicating that although she did not worry about data sharing, she thought it was useful for others. We are cautiously optimistic that this result indicates that there is a demand for privacy applications on Android smartphones that provide information about data sharing.

Typically, participants were able to read the text and numbers in the grid format and interpret them quickly. After using Privacy Leaks, participants were able to correctly answer questions about with whom information was shared, the type of information being shared, and which application sent the most data.

However, we did find that there were areas for improving the interface, as only 6 participants claimed they "understood what everything meant in Privacy Leaks." As discussed in Section 5.6, participants struggled to understand what Phone ID meant. Additionally, they did not know or recognize the different destination domains.

Three users initially failed to understand the purpose of Privacy Leaks, thinking that it was responsible for the data being shared. Similar results were found in a study on the impact of projecting excerpts from open network traffic on a wall in a public space [109]. More education would be needed about the purpose and goals of such tools to alleviate such confusion. This could be done through marketing, or providing an additional explanation of Privacy Leaks at install time. We also asked participants if Privacy Leaks was "accurate" and several stated they had no way of knowing. Two suggested it would take reviews in trusted media (e.g., "TechCruch" and

"BusinessWeek") to convince them that Privacy Leaks was trustworthy. Others said they would trust Privacy Leaks if it were from a well-known and trusted corporation or part of the phone's operating system.

5.5.3 Reactions to Just-In-Time Notifications

In the Just-In-Time (JIT) condition, 10 participants felt and heard the JIT notifications in addition to the visualizations. Due to the small sample size, we did not run statistical tests between the two conditions. However, there were some general differences between the groups that we describe.

Some participants were surprised by the frequency of the notifications. For example, P18 said, "I hear drops, this is going crazy! There are a lot of bleeps!" When playing the games with JIT notifications, participants often tried to figure out why data was being sent. P2 commented, "I'm trying to figure out when it actually sends data. I don't know what it just sent, [drop sound] not entirely sure what is being sent, I'm just loading it up. Probably checking for new content or updates." Participants questioned whether data was sent when they scored or reached a new level.

On average, participants found the sounds more annoying than the vibrations. However, this depended heavily on the individual. Participants suggested that Privacy Leaks should allow them to configure whether sounds and vibrations should be enabled. This functionality was already built into Privacy Leaks, but we did not include it in the study.

We anticipated that participants would overwhelmingly find the JIT notifications annoying or interruptive. Figure C.2 in Appendix C shows that participants had mixed reactions to the sounds and vibrations. For example, only one out of ten found the vibrations distracting. While five out of ten participants agreed that the sounds were distracting, 5 also said the sounds would allow them to keep working or playing without interruption. This may be due to the short amount of time using the notifications. Furthermore, by the time the participants received notifications, they had already been prompted with questions that they had a hard time answering, such as when data was sent. This probably increased their curiosity and therefore their appreciation of the notifications. Future work is needed on how users respond to JIT notifications over time.

Eight of the 10 participants who saw JIT notifications responded to the questions "The information provided by Privacy Leaks is accurate" affirmatively, while only 3 of the 9 who saw only visualizations were affirmative, typically saying they didn't know. This indicates that participants in the JIT condition were more likely to find Privacy Leaks accurate. It is possible that the audio, tactile, and visual feedback all combined to reinforce the information, making it seem trustworthier than the visualization alone.

5.5.4 Recommendations to Friends and Family

As smartphone owners rely on friends and family for app recommendations, we were curious about whether privacy leakage information would change participants' recommendations to friends and family [59]. Most (12) participants would not change their recommendation to their friend or family member about the game after using Privacy Leaks, saying the functionality was still the same. However, participants frequently said they would add that the data was being leaked. P14 said, "Angry Birds is still a fun game. I would probably inform her that they are tracking what you are doing." This indicates that game functionality was typically still more important to participants than data sharing.

However, some participants changed their recommendation. P16, who discussed recommending the game for a cousin in college said, "Yes, I would advise her not to play the Angry Birds after seeing the leaking."

All but 2 participants would add that information was being leaked if they were to write a description on the application market. For example, P11 said, "I might put a little note about Angry Birds talking to a couple companies I don't know." P18 described how he would recommend the game, "Probably say that both like to leak location and phone id. It is probably for marketing. It is important to let people know. Some people think it is helpful, some think it is invasive."

5.5.5 Privacy Preferences

Participants' existing privacy preferences impacted their reactions to the data sharing. Although we did not ask this directly, over the course of the interview six participants volunteered that they were not particularly privacy sensitive. They explained that data sharing was not overly concerning because they were not, for example, "paranoid" [P1] or "conspiracy theorists" [P17].

Two participants were even more sanguine about the data sharing. They were fully aware that data-sharing was a trade-off for free games, and were fine with this model. P11 said "It's not really a big deal to me. It can be a good thing. As long as they don't flash ads every second or something, I really don't mind." P7 said, "As long as that does not affect my life in negative way, I am ok to give the information away."

On the other hand, several participants had strong negative reactions to learning about data sharing after viewing Privacy Leaks. P5 said, "It really bothers me that this sort of thing happens, because I want to remain as anonymous as possible." P2 said, "This makes me an angry birdy."

5.5.6 Risks & Benefits of Data Sharing

In order to probe how users make decisions about data sharing, it is important to understand their concepts of risks and benefits of data sharing. We asked participants about the benefits and risks of sharing data with the questions, "Are there any benefits [risks] to you or [your friend] when the game shares information, and what are they?" We substituted the words "your friend" for the friend or family member they had selected at the beginning of the interview. We asked these questions at the end of the interview, to avoid biasing the interview, and therefore the participants had already viewed Privacy Leaks and knew how often data was shared and what the destination URLs were.

Fourteen participants thought there was no benefit overall to sharing the information with games such as Angry Birds. This is in contrast to Ur et al., whose participants often recognized that there may be economic benefits to themselves and the websites from on-line behavioral advertising (OBA). This may be due to the wording of the question or the context (data sharing from this particular game versus OBA in general). Alternatively, the participants in Ur et al. may have been better informed because they watched an informational video on OBA at the beginning of the interview, whereas our participants were asked to provide opinions without any education outside of our prototype [156].

Two of our participants mentioned that sharing location with certain applications was useful for functionality. Three participants mentioned targeted ads, but were unsure it was a benefit to them. For example, P19 said, "I guess customized ads are a benefit. It's a stretch. I don't click on

Perceived benefits	Perceived risks
none (14)	accidental purchases
free games (2)	porn
targeted ads (3)	virus
search and rescue (2)	annoying SMS
	price discrimination
	telemarketers
	find kids' backyard
	creepy
	social networks —(friends can access information)
	identity theft
	data breach
	worker with access to data 'goes postal'

Table 5.2: Responses to "Are there any benefits/risks to you or [your friend or family member] when the game shares information, and what are they?"

ads, so I'm not sure I can make that argument." P3 mentioned targeted ads as a possible benefit, but doubted the efficacy: "phone ID, location, that doesn't help them hone it on what I like." Some of these participants then concluded that there was no real benefit.

Only two participants stated that free games or improved functionality are benefits of sharing information. Two participants also mentioned search and rescue as a benefit of providing their location; their (mis)understanding was that first responders would be able to find them since their location had been shared.

Participants were also asked about the risks to themselves or their friend or family member when a game shares information. The risks mentioned by participants spanned an array of possibilities, as shown in Table 5.2 including accidental in-app purchases, getting a computer virus, and receiving annoying SMS messages. Some risks do not involve provable harm, such as someone knowing their kids' location, being creepy, or price discrimination. Some participants were particularly concerned about the risks to their children of data sharing, and were concerned that bad people could get access to the information.

5.6 Discussion & Future Work

We have built an interface that informed users about the frequency and destination URLs of data shared by smartphone applications. Most participants indicated that the application was useful and that they would like to install a similar application probably because it provided information participants could not get elsewhere in a simple layout with some explanatory text. However, future work is needed to improve the interface and validate through a field study how an improved interface can actually raise users' awareness of privacy leakages on smartphones. We first discuss near-term opportunities to improve the Privacy Leaks interface based on the suggestions from study participants. We then present a few suggestions to help users make informed decisions about data sharing with smartphone applications and control over their privacy-sensitive data.

5.6.1 Improving the Interface

There were three main areas in which participants frequently requested further information: phone ID, destination, and location. Participants felt that "Phone ID" was an unfamiliar term, and they were unsure about the implications of the Phone ID being shared. P18 expressed this confusion, "I don't know if it means type of phone or identifying the phone with the person." Participants were unclear if phone ID was just the model information, or if it also included their phone number and email address. Location was also confusing to some participants, who wanted a better understanding of how fine-grained "location" is. Participants suggested a roll-over for these columns that would explain these two fields more.

Participants did benefit from seeing the "Phone #" column, even if it was blank, as it helped them see that phone ID is not the phone number. Some participants asked about what other fields were possible. Despite the visual clutter of adding more columns, it may be helpful for users to see the possible privacy-sensitive types of data that could be sent, to help them distinguish what has been sent. For example, P14 became concerned about the capabilities of the phone and data sharing, wondering if the features of the phone could be combined to lead to inaccurate profiling, "Is that how they get those ads on top of the screen? They could take a picture of me and assume I'm into rap music." He wasn't clear that the application did not have control of the camera or other functionality.

Phone ID was often being sent along with additional information interpretable by the game developers. Our interface did not show the entire set of data sent with the Phone ID, as we thought it would be overwhelming or incomprehensible. However, additional information about the purpose of the data sharing might enable users to understand the frequency of data sharing.

As mentioned earlier, participants were confused about the destination domains, as the URLs were typically unfamiliar to them. Some users asked for the ability to click to open the domain in a browser. However, this may not be helpful, as many of the URLs do not have consumer-facing web sites, typically presenting toolkits aimed at application developers. While clicking on the domain may not be useful, other types of information could be provided for the user, such as brief descriptions of the company's purpose, and a link to the company's "opt-out" page, if one exists.

There are a number of ways to visualize data; we chose a simple grid format. The grid visualization highlighted the number of times different types of data were sent by applications. While this simple grid format allowed participants to quickly read the information displayed, it remains to be seen whether other visualization techniques (e.g., Wi-Fi Privacy Ticker [60]) would be more effective to improve users' understanding of privacy leakages associated with smartphone applications. Once an improved interface is developed, our plan is to run a field study to evaluate whether ongoing feedback on data leakages can be presented to users without causing too much annoyance and having users desensitized quickly over time.

5.6.2 Providing Usable Control

Informing users about data leakage is only a first step; users should also have control over data sharing. Today, even if users become aware of privacy issues with respect to certain smartphone applications, they have little choice but to uninstall offending applications on their phone. Although there are a few research prototypes that allow users to selectively hide privacy-sensitive information against data hoarding applications (e.g., AppFence [94]), it still remains open how

this prototype can be properly configured by users to protect their privacy. For instance, our participants had a hard time considering all the implications of blocking data sharing on all their applications. This is not surprising, considering how many participants were unaware of data sharing before the study, and had not had time to consider how they would control data. However, some participants suggested particular contexts in which they could imagine wanting particular control. These included mobile banking, being in a government building, and not wanting their car insurance company to find out they were texting while driving. This suggests that users have very specific desires for controlling information that might not fit into broad categories. Two participants made analogies to their computers when suggesting protective steps they could take. They discussed installing an anti-virus and deleting cookies as options they could take to control information sharing. Exploring usable privacy control mechanisms is another future direction to pursue by our team.

5.7 Limitations

Qualitative studies with small sample sizes have limitations, such as lack of statistical power. However, the in-depth interviews did allow us to get qualitative insights into participants' reactions to the notifications and visualizations, as well as their understanding of the information they were seeing. While we provided the numbers of participants that made similar statements, we do not claim that these could generalize to larger populations.

A lab study has limited ecological validity. The participants were not using their own phones, playing games of their own choice, and were not in privacy-sensitive locations (e.g., at home). With a greater range of locations, applications, and situations, participants may be more sensitive to the particular context of information sharing. In real settings, users may actually be more concerned with privacy leakage than when they are in a lab and using a lab phone. We did ask participants to imagine their friend or family member playing the game "on the bus, at the doctor's office, or waiting to meet you somewhere." However, we have no indication that participants considered the privacy-sensitive nature of any of these locations.

On the other hand, lab participants may exaggerate their concern and interest in the task at hand (understanding or caring about privacy leakages) to appear to be a good participant. Also, naming our app as Privacy Leaks may have biased some participants who would not have considered privacy risks otherwise. Furthermore, they may not actually do what they claim they will do when they are in the lab. Previous work has shown that survey participants may report that they engage in privacy-protecting behavior, but behavioral studies show that these self-reports are inaccurate [102].

While a field study would address some of the concerns about a lab study, we would only be able to measure quantitative actions that the participants take. We would not have access to the initial verbal reactions or questions that we have in a lab, nor would we be able to probe them for details as they are viewing the interface.

5.8 Conclusions

Our qualitative interviews provide insight into users' understanding of data sharing, both before and after being informed in real-time about data sharing from two smartphone games. Overall,

we have found that participants have misconceptions about data sharing occurring through smartphone applications, do care about applications that share privacy-sensitive information with third parties, and would want more information about data sharing. Thirteen out of 19 participants did not know that data would be shared for the purpose of advertising. Many had never considered it before; others believed data was only shared with the application's developers in order to improve the game or provide useful functionality; yet others understood that data was being shared for marketing purposes. Most participants were not aware and often not comfortable with the scope of data sharing done by these game applications, both in terms of amount of data shared and the destinations of the data. This is particularly troubling as the ad and app industry may be working on the assumption that users understand the trade-off between free apps and data sharing.

Moving forward, we continue to explore tools and interfaces that can improve users' awareness of privacy leakages while using smartphone applications and usable control mechanisms that can help users prevent unwanted data sharing with smartphone applications.

Acknowledgements

This research was funded largely by Microsoft, as part of an internship in the Technology and Policy Group.

Chapter 6

I Don't Remember, I Don't Recall: The Impact of Timing on Recall of Smartphone App Privacy Notices

6.1 Introduction

Currently, users of major smartphone platoforms are informed about data collection practices through permissions dialogs or longer privacy policies. Despite efforts to develop standardized privacy notices [20] and privacy metrics to inform users about privacy-intrusive apps [121, 83, 122], it is not clear *when* privacy notices should be shown to users. While it is well-known that people often ignore notices such as computer security dialogs or End User License Agreements, which may be shown at install-time [48], it is less understood whether there are optimal times to show privacy notices to maximize attention and recall.

This chapter assesses how timing can impact the recall – a measure of effectiveness – of an app privacy notice [31]. We conducted a web survey and a field experiment to compare whether participants better recall a privacy notice shown before installation, or before, during, or after app use. A follow-up web survey investigated app store notices in particular. We make the following contributions in this paper:

- We show that timing of smartphone app privacy notices has an impact on the recall of the notices. We show these effects both in a web survey and in a field experiment.
- We provide recommendations on how to integrate privacy notices into apps for improved recall. Since notices in the app store had low recall rates, our results indicate in-app notices are more salient, preferably shown at the beginning or during app use.
- Since there are other benefits to providing notices in the app store, we offer design guidelines for improving privacy notices shown in the app store based on a follow-up web survey.

This chapter is largely based on a paper co-authored by Rebecca Balebako, Idris Adjerid, Florian Schaub, Alessandro Acquisti and Lorrie Faith Cranor (under submission). The title is based on a song by Peter Gabriel. Other lyrics in this song are relevant to privacy notices: "Strange is your language and I have no decoder. Why don't you make your intentions clear?"

6.2 Related Work

A brief overview of the current state of privacy notices on smartphone platforms is provided in Chapter 1. We provide some additional information relevant to this study specifically. The role of human decision making on privacy and security notices is also discussed in Chapter 1. We discuss user studies examining privacy notices on smartphones that are relevant to this study on timing of warnings.

6.2.1 Standardized Smartphone App Privacy Notices

U.S policy-makers recognized that different notices may be confusing to users, and National Telecommunication and Information Administration (NTIA) attempted to develop a code of conduct for standardized short-form privacy notices for smartphone apps [20]. This notice includes a list of data elements and third-party entities about which users should be informed. Industry associations have backed the code and have developed several examples of how it could be implemented [88]. The security company Lookout released an open-source implementation of the standard, called Private Parts [160]. While study participants did not understand all the terms used in the NTIA code (see Chapter 4), we chose to use it in this work because it is realistic, standardized, and not biased toward any platform.

A meta-analysis of consumer warnings across a variety of products found that recall of the notice was one of five dimensions – including also attention, judgment of risks, comprehension, and compliance with the warning – that define the warnings' effectiveness [31].

6.2.2 Privacy Notice Timing on Smartphones

In the context of smartphones, some work has studied the impact of notifications on user activity, such as interruptions caused by phone calls [50] or messaging [143]. However, we specifically examine privacy notifications related to the app that is the primary task rather than external interruptions.

Privacy information may influence users choice of apps at install time, if the information is provided in a clear manner. When shown a 'privacy checklist' in the Google Play Store, users would select the app requesting fewer permissions [108]. Additionally, when asked to compare similar apps with different permission requests, users demonstrated that they were willing to pay more for apps requiring fewer permissions [68]. A 'risk' score generated by examining an app's permissions was shown to be effective

in helping users choose an app that requested fewer permissions [84]. However, users may follow different paths when selecting an app. They may not always be comparing two similar apps with different privacy impacts, but instead may select an app based on peer recommendations, and may forgo comparison shopping. Therefore, some work, including ours, examines reactions to notices within the context of specific apps. Participants in a lab study were shown just-in-time notifications after a few minutes of playing Angry Birds and Toss It in Chapter 5. Participants in a field study were shown location notifications while using the same apps they normally used on their phones [80]. In both studies, and in contrast to our study, the notices appeared multiple times if the permission was requested frequently. These studies found that users appreciated the notices, and the latter found that users did take privacy-protective actions, such as uninstalling apps, when

they saw notices during app use. These results align with our findings that notices during app use are effective. However, we further provide a clear comparison of timing effects without additional confounds by showing the same notice at different points in time.

Our field study is the first to investigate the timing of privacy notices by asking users to download and install an app on their own phone. Furthermore, we are the first to investigate the impact of timing on the memorability of privacy notices.

6.3 Methods

We investigated whether the time at which a user sees a privacy notice impacts her recall of the notice. Participants installed and used an app specifically designed for this investigation, and saw a privacy notice. We measured whether the timing of the notice significantly impacted correct recall of the notice's content. We also verified that participants found the privacy notice to be relevant and worth remembering by asking them to evaluate the notice.

This chapter describes a web survey, a field experiment, and a follow-up web survey. The first web survey and the field experiment were used in combination to examine app timing, as they each have specific benefits. Web surveys allow quick access to a large, relatively diverse participant pool. Our field experiment aimed to be ecologically valid, in that users were installing our app on their own phones on their own time, with the vagaries and distractions that may occur naturally in such a situation. The follow-up web survey was based on the findings of the first web survey and the field experiment, and allowed us to examine some variations on the app store notice condition that are not available in the real app store.

This section describes the methods that are similar across the web surveys and field experiments.

We measured participants' recall of the privacy notice, and their self-reported desire to have and remember the notice within the context of the employed app. All participants completed five steps: 1) consent form and demographic questions, 2) install and play the app, 3) experience a distractor or delay, 4) answer recall questions, and 5) evaluate notice. To ensure that participants responded to recall questions of the notice based on memory, they were not allowed to return to previous steps to revisit the notice.

All participants were recruited to use a smartphone app. There were not initially told that the research was about privacy or notifications, but were informed of the purpose at the end of their participation.

Quiz app

We designed and deployed a simple quiz app. Our objectives in selecting the app content were to create an app both entertaining and distracting, and that could be completed in a few minutes. Therefore, we developed a history quiz that asked eleven questions about the inventions of less famous inventors (see Figure 6.2). Before beginning the quiz, the app showed two screens: first a paragraph of instructions, and second a page to enter an email address (field experiment) or code

Figure 6.1: The privacy notice.

Figure 6.2: A quiz question from employed app.

(web survey) to link participants to their consent form. After answering the eleven quiz questions, participants saw their score. The app was developed in HTML and JavaScript using PhoneGap.[1]

We created two similar Android app store entries; one entry showed the privacy notice (shown in Figure 6.1 and the other did not.[2] In the first web survey, the privacy notice was the right-most image (as seen in Figure 6.3). For the field experiment, we hypothesized that the notice might be more salient as the first (leftmost) image. We did not find that this improved recall.

At the time of the field experiment, our apps were not rated, and had neither ratings nor comments from users. The description of the apps in the store stated that it was part of a research project and included a link to the consent form approved by our IRB. In the "Developer's Website" section of the app store, we included a link to our website detailing the steps of the experiment and including a link to the consent form. The privacy policy link in the store pointed to the image of the privacy notice shown in Figure 6.1.

Privacy Notice Design and Verification

To create a realistic notice, our design was based on the aforementioned code of conduct for standardized privacy notices [20]. More specifically, we employed Private Parts,[3] an open source

[1] `http://build.phonegap.com`

[2] `https://play.google.com/store/apps/details?id=com.rebeccahunt.historyquiz0`

[3] `https://github.com/lookout/private-parts`

implementation of this standard, which we modified to match our app's color scheme (see Figure 6.1). Our notice informed users that the app collected "Browser History" and that this data would be shared with "Ad Networks." We selected only one data and one entity to avoid confounds or information overload for the participants. To insure that the notice was the same across all conditions, including in the app store where we could not show an interactive notice, we moved the explanatory text – which appears only after clicking a Private Parts element – next to the privacy icons.

The data and entity shown were selected based on an preliminary web survey we conducted to test possible notices. We randomly assigned participants to view one of three notices, and measured whether they claimed to like, remember, or be concerned about different notices if shown in a "History Quiz" app. We recruited 238 participants on MTurk in December 2013. Participants took an average of 4.5 minutes to complete the survey and were compensated $0.25. The notice conditions were: 1) Share browsing history with social networks, 2) Share location with advertising networks, and 3) Share contact information with data analytics. We chose these options based on previous findings that users would want to know about these data and entities [37].

There were no significant differences in whether respondents claimed to care about, understand, or remember the notices (χ^2-test). Regardless of notice, participants said they wanted the notification and would remember it a day later. We did find some differences in comprehension by including multiple-choice and open-ended questions to measure understanding. These responses showed lower rates of understanding for "social networks" and "data analytics." Therefore, our final notice used "advertising networks." Because our results did not show differences between the data types, we selected one based on previous research. We speculated that users may be inured to collection of location, as location data is collected by many apps [36, 123]. On the other hand, previous research [122] has found that users are very uncomfortable with apps taking contact information without a clear purpose. We selected the middle ground of user concern: "Browser History." We further ascertained the appropriateness of these choices for the notice in our web survey.

6.3.1 Timing Conditions

The conditions varied based on the moment in time at which the privacy notice was shown. The first web survey and the field experiment had five conditions, which represent privacy notice timings that occur in existing apps and platforms. The app store condition varied slightly from the other conditions: the notice in the app store did not occupy the full screen, whereas all the other notices were shown full screen. The timing conditions were:

Not Shown. The privacy notice was not shown to the app user. This is the control condition.

App Store. The notice is displayed as a screenshot in the app store (see Figure 6.3), similar to previous work on showing privacy indicators in the app store [108, 84]. This is the only way that Android allows privacy notices to be displayed at install time.

Before Use. The notice is displayed after the app splash screen, before the first page of the app with instructions. This resembles the timing of EULAs or notices that are shown before app usage.

During Use. The notice is displayed several steps into the app. This was meant to mimic a just-in-time notification as used in iOS.

Figure 6.3: App store with the privacy notice.

After Use. The notice is shown after completing the last question of the app quiz. This would be the timing of a privacy notice shown to summarize data sharing and collection, after the app has been played, similar to the timing of a summary notice [35].

In the follow-up web survey, we investigated variations on the app store timing condition. The two additional conditions used in the second web survey are introduced in the section discussing the follow-up web survey.

Our two web surveys and field experiment were between-subjects experiments, and participants were randomly assigned to one condition. The app, privacy notice, recruitment, and all associated materials were identical across conditions.

6.3.2 Questions Following App Usage

In both web surveys and the field experiment, participants completed the same exit survey at the end. The questions allowed us to evaluate their recall of the notice and app, and to evaluate the notice. The questions used to measure recall of the privacy notice were, "With whom does the app share data?" and "What information was collected by the app?" The questions were multiple-choice, with six possible answers, including "I don't remember." Participants were also

asked two multiple-choice questions about the contents of the quiz questions, the color of the app background, and whether they remembered seeing the privacy notice.

At the end, participants were shown the privacy notice again and were asked to evaluate it. These questions were used to measure whether participants perceived the privacy notice's content as important and whether they wanted to remember it. Participants were told, "This is the privacy notification for the app. Please note, we did *not* collect this information, but please imagine your reaction if this really occurred on your phone." Six 5-point Likert-scale questions about the notice included positively-biased questions such as, "The privacy notification gave me information I care about" and negatively-biased questions such as "This notification could be improved so I understand it better." Four 5-point Likert-scale questions were used to evaluate participants' opinions of the timing of the notice. The questions included whether the timing was disruptive, unexpected, allowed them to make decisions, and whether they could pay attention to the notice. These questions also included a "Not Applicable/Don't Remember" option, as participants in the control group did not see the notice and therefore were not positioned to evaluate the timing.

6.4 Web Survey Results

In this section we describe the results of our first web survey, which examined the impact of privacy notice timing on recall.

6.4.1 Web Survey Participants

Web survey participants were paid $1.01 and were recruited via MTurk. Two hundred and seventy-seven U.S. participants completed the survey. Participants completed the survey in a median of 9.08 minutes (range 2.82-29.6). The participant group was diverse. Nearly half of participants (49%) were female (1 participant declined to state gender). Almost half (48%) had a bachelors degree or graduate degree. While the ages ranged from 18 to 69 years, the median age was 29 year. Forty-five out of 50 U.S. states were represented. Most of our participants owned and used a smartphone (95%), although we did not recruit for smartphone owners, and specifically stated that owning a smartphone was not a prerequisite for the web survey. There were no significant differences in the following demographics across timing conditions: age (ANOVA F=1.67 p=16), gender (χ^2_8=12.4, p=.135), and smartphone type owned (χ^2_{20}=19, p=.524) respectively).

The timing conditions were randomly assigned, and there were between 39 and 67 participants in each condition, as seen in Table 6.1.

6.4.2 Web Survey Analysis

The web survey had two main results. First, the timing condition did impact the ability to recall the notice. Second, participants, overall, claimed to find the notice useful, and indicated that they would want to still remember it a day later.

Recall of the Privacy Notice

Most participants did not feel confident in their recall of the privacy notice when asked, "Do you remember seeing the privacy notice?" after the distraction. Only 36.5% responded either, "I

condition	participants	recall rates	
not shown	67	2	(3%)
app store	57	10	(17%)
before use	67	25	(37%)*
during use	20	18	(43%)*
after use	39	11	(28%)*

Table 6.1: Number of participants in web survey, and correct recall of both the data and entity described in privacy notice, by condition. Values significantly different from "not shown" are marked with * (Mann-Whitney U with Bonferroni correction).

remember most of it," or "Yes, I remember it well," while the remainder responded that they did not remember it at all, or only remembered it vaguely.

Only 24.5% of participants correctly remembered both the data (*browser history*) and entity (*ad networks*) shown in the privacy notice. More participants remembered the data (40.8%) than the entity (31.4%). Both were recalled better than simple chance of selecting one of the six options (16.6%). Self-reports of remembering the notice did positively correlate with the ability to correctly identify the elements on the notice (r_Φ=.546 p=.001 for data and r_Φ=.602 p=.001 for entity). There was a positive correlation between correct recall of the data and the entity (r_Φ=.515 p=.001). We used an ordinal variable "RecallCorrect" with three levels: 1) did not remember any part of notice, 2) remembered at least one part of notice, 3) remembered both data and entity from the notice correctly.

When the notice was shown before, during, or after app use, participants remembered it more accurately than when shown in the app store, see Table 6.1. A Kruskal Wallis test revealed a significant effect of timing condition on RecallCorrect (KW χ^2_4=70.2, p=0.001).[4] Post-hoc tests (Mann-Whitney U with Bonferroni correction) showed significant differences between 'not shown' and each of the three app use conditions (before/during/after), as seen in Table 6.2. The three app use conditions were not significantly different from each other. Participants who saw the notice in the app store were less likely to remember the notice than those who saw it during app use. This indicates that notices shown at the time of app use are most beneficial for retention and later recall of the notice.

People's self-reported frequency of reading privacy policies was a good indicator of their memory of the notice – of those who stated that they read policies 'rarely' or 'never,' 15% remembered the notice, while 30% of those who read 'Sometimes' or 'Always' correctly remembered the notice (KW χ^2_2=11, p=.004).

In the web survey, the distractor was a set of IUIPC web privacy concern questions [125]. Responses to the IUIPC scale ([125]) in the categories "Control" and "Collection" were not correlated with RecallCorrect, although "Awareness" was a weak predictor (one-way ANOVA, F=5.97, p=.015). Participants recall was not affected by the following demographics: age (ANOVA, F=.38, p=.54), education (KW χ^2_2=.267, p=.875), gender (KW χ^2_2=1.06, p=.590), and owning

[4]Throughout this paper, for the Kruskal-Wallis (KW) tests on RecallCorrect, we examined significance after Benjamini-Hochberg corrections.

condition	not shown		app store		before use		during use	
	r	p	r	p	r	p	r	p
app store	−0.45	.27						
before use	−2.3	.001*	−1.2	.005*				
during use	−2.9	.001*	−1.8	.001*	0	1.0		
after use	−2.2	.001*	−1.1	.006*	0	1.0	−0.04	0.93

Table 6.2: Web: r (effect size) and p-values of pairwise comparisons on recallCorrect using Mann-Whitney U with Bonferroni correction. Significant results marked with *.

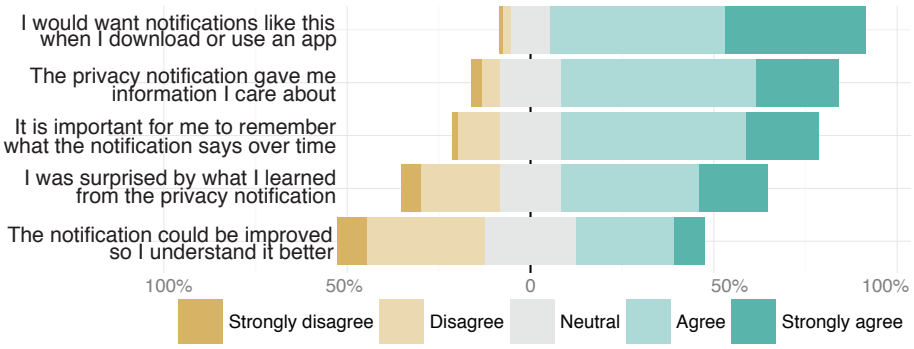

Figure 6.4: Web survey participants want the notification and want to remember it.

a smartphone (KW χ^2_2=2.31, p=.315). Overall, previous preferences and timing are the main predictors of whether the participants remember the privacy notice.

Despite not remembering the notice well, participants remembered other aspects of the app. They were able to identify the inventors asked about in two separate questions (88.1% and 67.9%), as well as the app's background color (80%). These aspects of recall were not correlated to the timing condition, indicating that we did not, by chance, have an uneven distribution of recall skills between conditions. Better recall for the app content was to be expected because participants focused on answering the quiz questions (primary task), while the interaction with the notice was a secondary task.

Evaluation of Privacy Notice

We would not expect participants to remember a notice unless they care about it and would want to remember it. Figure 6.4 shows the results of the Likert-scale questions used to evaluate the privacy notice. Our results validate the notice's relevance, showing that, overall, participants wanted to remember it and felt it had information they cared about. We note that liking the notice does not imply that they liked the data collection described in the notice. The responses to these Likert-scale questions did not significantly depend on the timing condition (KW test with Bonferroni correction, α=.01).

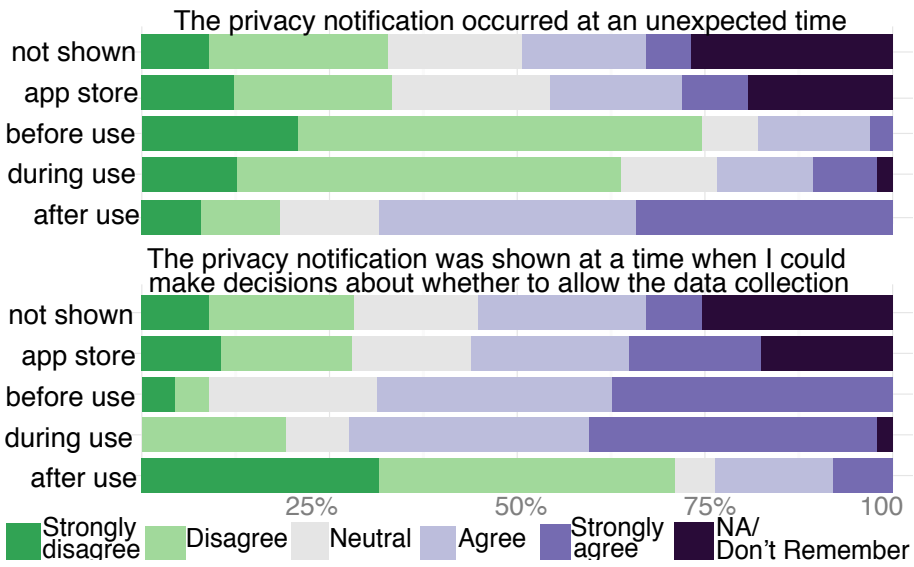

Figure 6.5: Web survey responses about timing of privacy notice. Participants in after app use condition were more negative about timing.

We also evaluated participants' reactions to the timing of the notice. We did not remind participants what timing condition they were in.

The timing condition significantly impacted participants' responses to two questions about the timing of the notice: "The privacy notification was shown at a time when I could make decisions about whether to allow the data collection" (KW χ^2_4=32.4, p=.001) and "The privacy notification occurred at an unexpected time" (KW χ^2_4=44.2, p=.001) (see Figure 6.5). Participants in the "after use" condition reacted negatively to the timing of the notice, and said more frequently that the timing was unexpected and that they could not make decisions about the data collection.

Differences between iPhone and Android owners

Our web survey resembled the Android store across all conditions and participants. The web participants themselves owned different types of smartphones. About half owned Android smartphones (54.2%) and 38.6% of participants owned an iPhone. Since iPhone and Android show privacy notices at different times, users of different platforms may be habituated to different timings. However, we did not find significant differences between Android and iOS owners in terms of recall of the notice or in participants' rating of the timing of the notice (KW χ^2_2=.13, p=.94).

6.5 Field Experiment Results

The web survey indicated that the timing of a notice impacts users' ability to recall the privacy notice. However, the ecological validity is limited by the browser-based setting. Thus, the goal of the field experiment was to measure whether timing of the privacy notice also had an impact on participants' memory of the notice when the app was installed and used on participants' own

phones in their own environments. By running a field experiment, participants are subjected to distractions and variable conditions similar to what they would encounter when installing the app outside of a study. Also, by installing on their own phones, as opposed to phones provided by the experimenter, participants may exhibit realistic privacy concerns.

The field experiment consisted of the five steps described in the Methods Section, similar to the web survey. Field experiment participants installed the app and completed the quiz. Twenty-four hours after completing the app quiz, participants received an email with a link to the exit survey, in which they answered recall questions and evaluated the notice. If participants completed all these steps, they were e-mailed a $5 Amazon gift code.

6.5.1 Field Experiment Participants

We recruited 126 participants from three university participant pools: Phone-Lab at SUNY Buffalo[5] (n=29), Notre Dame University[6] (n=37), and CBDR at Carnegie Mellon University [7] (n=42). We also posted ads on craigslist and reddit, which yielded 18 additional participants. Our participant pool skewed young. While the range of ages was 18–55, 80% of our participants were 30 or younger (median=23.5). Our participants were well educated, as 57% had a bachelors or graduate degree; 46.8% were female, and the rest were male. Participants were based in 24 different U.S. states. There were no significant differences between conditions in age (ANOVA, F=1.67, p=.16), gender (χ^2_4=.716, p=.949), U.S. state (χ^2_{100}=113, p=.171), or education level (χ^2_{12}=14.1, p=.297). The field experiment was conducted only on Android. Compared to the online survey, participants were slightly younger, all used Androids, and resided in fewer US states, but otherwise the participant groups were similar. Table 6.3 shows the number of participants in each condition.

To get an idea of the participants' familiarity with installing apps, participants were asked to self-report how often they installed apps ("rarely," "sometimes," "often," "daily"). While this is a subjective measure, most participants (61.1%) stated they sometimes install apps, with only a small group stating, "often" or "daily" (19.8% combined). We asked participants what they reviewed when deciding to install an app. Most participants stated that in general they consider the description of the app (83.3%) and app ratings (76.2%). Slightly more than half described looking at the permissions (57.9%).

condition	participants	recall rates	
not shown	35	3	(9%)
app store	21	3	(14%)
before use	30	10	(33%)*
during use	24	5	(20%)*
after use	16	6	(37%)*

Table 6.3: Number of participants in field experiment, and correct recall of notice by condition. Values significantly different from "not shown" are marked with * (Mann-Whitney U with Bonferroni correction).

[5] **www.phone-lab.org**
[6] **www.nd.edu**
[7] **cbdr.cmu.edu**

While 126 participants completed the field experiment, additional participants started the experiment but dropped out at various steps. Of the 204 participants who filled out the consent web form, 61 failed to download the app, and an additional 6 started but did not complete the app quiz. Of those who finished the app, 9 did not complete the exit survey they received 24 hours after completing the app. There were no significant differences between conditions in terms of completing the app quiz. To determine whether people dropped out due to the privacy notice, we contacted everyone by email that filled out the consent form but did not complete the app, asking for a short explanation. Of the 15 responses we received, only one cited concerns related to the privacy notice. Other responses indicated that people forgot or had technical issues downloading the app.

Participants were asked to rate and review the app before answering recall questions. Participants were rather neutral about the app when asked to rate it from 1 to 5 stars, with the median score being 3 stars. While some participants found the app "simple," stating that it resembled a quiz they could take online, others enjoyed the educational aspect of learning about history and called the app "interesting."

Of those participants who completed the exit survey, 90% did so within 48 hours of finishing the app, the median time being 26.3 hours after completing the app. However, six participants took 3 - 7 days to complete the exit survey. In the app store condition, participants saw the notice slightly earlier than in the other conditions. The median time participants took to download and finish the app was 6 minutes, which is negligible compared to the minimum 24 hour delay before participants were asked to recall the notice. Therefore, we do not think that seeing the app more recently in the app use conditions impacted the recall rates.

6.5.2 Field Experiment Analysis

The field experiment had two main results, which were in agreement with the web survey results. First, the timing condition did impact the ability to recall the notice. Second, participants, overall, claimed to find the notice useful, and indicated that they would want to still remember it a day later.

Recall of the Privacy Notice

Participants did not feel confident that they remembered seeing the privacy notice. When asked, "Do you remember seeing the privacy notice?" 54% of all participants said they only remembered it "vaguely," while 21% said they did not remember it at all. Only 5% said they remembered the notice well. Unlike the web survey, self-reported response of remembering the notice did not correlate with the ability to correctly identify the elements on the notice (r_Φ=.207 p=.021 for data and r_Φ=.121 p=.184 for entity).

While the recall rate is lower than that of the web survey – likely due to the longer delay – the trends are similar. More participants remembered the data than the entity. Overall, just over one-third (37.3%) of field experiment participants correctly identified that the privacy notice said data was shared with the entity "Ad Networks." A smaller percentage (26.2%) correctly identified that the privacy notice said that it would collect "Browser History" data. About one-fifth of participants correctly remembered both aspects of the privacy notice (21.4%). Both of these percentages are better than if multiple choice answers had been selected randomly (16.6%). Correct recall of the

condition	not shown		app store		before use		during use	
	r	p	r	p	r	p	r	p
app store	0	1.0						
before use	−1.3	.001*	−0.78	0.05*				
during use	−1.2	.004*	−0.61	0.13*	0	1.0		
after use	−1.1	.006*	−0.61	0.12*	0	1.0	0	1.0

Table 6.4: Field experiment: r (effect size) and p-values of pairwise comparisons on recallCorrect using Mann-Whitney U with Bonferroni correction. Significant results marked with *.

two aspects of the privacy notice was positively correlated (r_Φ=.548, p=.001); for example, 81% of participants who remembered the entity type also remembered the data.

The timing condition was a significant predictor of recall, with all conditions during app use yielding better recall rates than the app store or control conditions. The percentage of participants who correctly recalled both aspects of the notice is shown in Table 6.3. Of the participants who saw the notice, those in the app store condition were the least likely to remember it. Overall, timing had a significant impact on "RecallCorrect" (KW χ^2_4=24.1, p=.001). Post-hoc tests (Mann-Whitney U with Bonferroni correction) showed significant differences between 'not shown' and all three conditions during app use (before, during, and after play), but the difference between "app store" and "not shown" was not significant. Differences between the three within-app conditions were also not statistically significant. These pairwise comparisons are shown in Table 6.4.

A statistically significant difference in RecallCorrect exists between participants who self-reported to read privacy policies frequently versus those that did not (KW χ^2_3=16.1, p=.001). Of participants who indicated that they read privacy policies 'Always' or 'Sometimes', 30% correctly recalled the notice, while of those who selected 'Never' or 'Rarely,' only 11% correctly recalled the notice. This indicates that pre-existing preferences and behaviors impact the user's ability to remember the notice.

Two additional variables impacted RecallCorrect. First, the self-reported frequency of installing apps on their phone impacted RecallCorrect (KW χ^2_3=11.2, p=.010). The more frequently they installed apps (e.g. 'Often' or 'Daily'), the more likely they were to correctly remember the notice. Second, we expected that participants would be less likely to remember the privacy notice when there was more time between using the app and answering the recall questions in the exit survey. We found an impact of this delay on RecallCorrect (KW χ^2_4=12.8, p=.012).

There were no statistically significant differences between demographic groups in RecallCorrect. Region (KW χ^2_{25}=28.6, p=.282), gender (KW χ^2_1=2.09, p=.148), age (ANOVA, F=.01, p=.91), or education level (KW χ^2_8=5.67, p=.683) did not affect RecallCorrect. As in the web survey, previous preferences and timing are the main indicators of whether the participants remember the privacy notice.

Participants had a much better recall of other aspects of the app than the privacy notice. The majority correctly identified the background color (74%), and were able to identify two inventors described in the quiz in two questions (86.6% and 57.1%). Memory of these aspects of the app did not correlate to the timing condition, indicating that the ability to recall the app in general was evenly distributed between conditions. In general, participants' memory of the privacy notice was

not correlated with their memory of the other aspects of the app. That is, correctly identifying the people in the app quiz or the background color did not correlate to correctly remembering the data or entity in the privacy notice (χ^2-test, corrected α=.017). This further suggests that any ability to remember the notice, or not, was not simply a matter of remembering the app overall, but isolates the effect of timing as an impact. This also indicates that privacy is treated as a secondary task, as the primary task (history quiz) was better retained.

Evaluation of Privacy Notice

As with the web experiment, we verified that our privacy notice was perceived as relevant by participants and that they wanted to remember it. Our findings support that the notice was appropriate for this experiment.

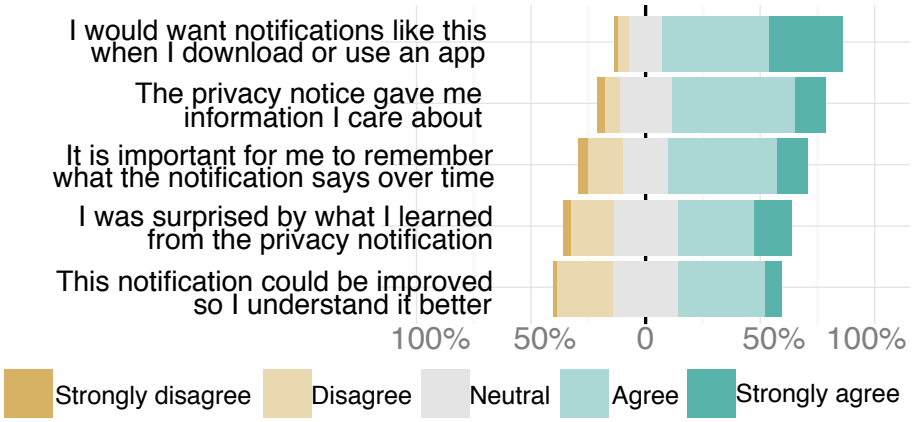

Figure 6.6: Field experiment participants want the notification and want to remember it.

Overall, participants stated that they wanted to see the notice when downloading or playing an app (78%), wanted to remember the information in the notice a day later (60%),and cared about the information shown in the notice (66%). Half of the participants found the content of the notice surprising (50%). The results of the questions are shown in Figure 6.6.

To evaluate participants' opinions of the timing of the notice, we asked participants the same Likert-scale questions as in the web survey (disruptive, unexpected, could make decisions, and could pay attention). Unlike the larger web survey, there was no significant impact of the timing condition on the responses to these questions (KW, p=.076, .444, .057, .022 respectively α=.0125 with Bonferroni correction).

6.6 Follow-up Web Survey on App Store Notices

The first web survey and field experiment indicated that when the notice was shown in the app store participants had low rates of recall. We used an app store design that matched what had been proposed by the multi-stakeholder group that developed the notice code of conduct [20]. This design did not require any changes to the app store itself as the privacy notice could be inserted as a screenshot. However, the notice was notice displayed prominently. In the follow-up web survey

described in this section, we evaluated whether the notice in the app store was less effective due to the small size and distractions (such as other elements describing the app).

6.6.1 Follow-up Web Survey Participants

Our follow-up web survey used the same method as the first web survey, described in the Methods Section. Three of the conditions (Not Shown, App Store, and During Use) were the same as the previous web survey and field experiment. We added two new conditions designed to show the notice more prominently. The two new conditions were:

App Store Popup. The privacy notice was shown to the user as a pop-up after the permission dialog popup. The app store was greyed-out, and the privacy notice dominated the screen (see Figure 6.8).

App Store Big. The notice is in the same location as the screenshots in the app store (see Figure 6.7), but the image is as wide as the store, and replaces other screenshots.

Figure 6.7: App store with the big privacy notice shown in place of screenshots.

Figure 6.8: Privacy notice as a popup displayed after the Android permission screen.

Web survey participants were paid $1.01 and were recruited via Amazon MTurk. The median age of the 326 participants was 31 years (range 19–69). Forty-six percent of participants were

condition	participants	recall rates	
not shown	63	1	(2%)
app store	52	3	(6%)
app store big	84	12	(14%)*
app store popup	69	18	(26%)*
during use	58	26	(45%)*

Table 6.5: Number of participants per condition in follow-up web survey, and correct recall of notice by condition. Values significantly different from "not shown" are marked with * (Mann-Whitney U with Bonferroni correction).

female; 3 participants opted not to select gender. Almost half (49%) had a bachelors degree or graduate degree. Forty-four out of 50 U.S. states were represented. Most of our participants owned and used a smartphone (94%). There were no significant differences between timing conditions and the following demographics: age (ANOVA F=1.09 p=.36), gender (χ^2_8=8.11, p=.423), and smartphone type owned (χ^2_{12}=18.7, p=.096) respectively). Participants completed the survey in a median of 8.71 minutes (range 2.68-27.8).

The timing conditions were randomly assigned, and there were between 52 and 84 participants in each condition. Table 6.5 shows the number of participants in each condition for the web survey and the field experiment.

6.6.2 Follow-up Web Survey: Recall of the Privacy Notice

The follow-up web survey found that the app store notice was recalled at better rates when it was displayed more prominently in the app store than when it was just one of many screenshots in the app store. However, when the notice was displayed during app use, participants remembered it more accurately than any of the app store conditions, as shown in Table 6.5.

A Kruskal Wallis test revealed a significant effect of timing condition on RecallCorrect (KW χ^2_4=81.2, p=0.001). Post-hoc tests (Mann-Whitney U with Bonferroni correction) showed significant differences between 'not shown' and and all of the conditions except 'app store', as seen in Table 6.2. The two new app store conditions were not significantly different from each other, but were significantly better than the previous app store condition, indicating that size and prominence improves recall. However, despite the improvements with the new app store conditions, participants who saw the notice in any of the app store conditions were still less likely to remember the notice than those who saw it during app use, and this difference was significant. This indicates that despite our efforts to improve the app store notice, during app use notices still had better rates of recall.

6.7 Limitations

We did not study the impacts of habituation on users' ability to recall the notice. Although we used a privacy notice modeled after a standardized notice, there is little indication that many app

developers have adopted this notice yet. It is possible that if this notice is widely adopted across apps, smartphone users may begin to ignore them, no matter when they are shown.

In the flow of our experiment, participants were asked to install a specific app and were directed to that app's Play Store page. Therefore, our results may apply to situations in which a smartphone user knows the name or link of the app they want, and will not be comparing between apps. If participants had been asked to select between comparable apps, they may have paid closer attention to app store privacy notices.

Although participants were using their own phones in the field experiment, they were aware that they were enrolled in a study, and may have implicitly trusted the researchers to protect their privacy. This may have impacted the level of attention paid to the notice. We tried to mitigate this by making the app as realistic as possible without unnecessary explicit references to the study in the install or app use process.

In the exit questionnaire, participants were asked to evaluate the notice after they had been asked to recall the content of the notice. This may have caused a bias in that participants who felt they did poorly in the recall questions may have been more critical of the notice. Therefore, we did not use these evaluation variables to examine or predict recall.

We only used one app and one notice to be able to isolate the timing effects. The app itself was rather innocuous and would not have raised concerns by the nature of its content; other apps may yield different results.

We do not assume that recall will necessarily change behavior, as we recognize that many elements go into a smartphone users' decisions. Furthermore, we specifically studied notice here, and did not examine users' ability to or desire to control data sharing.

6.8 Discussion

In two web surveys and in a field experiment, we investigated participants' recall of a privacy notice after installing and playing a history quiz app. We specifically examined how varying the time at which the notice was shown impacted participants' ability to recall the message. We find that participants' who viewed the notice before installing – in the app store – were the least likely to remember the notice. In fact, seeing the notice in the app store as a screenshot – the only option currently available to app developers who wish to show a privacy notice in the app store – was not significantly better than not seeing the notice at all. Seeing the app notice during app usage resulted in better recall. Although participants remembered the notice shown after app use as well as in other points of app use, they found that it was not a good point for them to make decisions about the app because they had already used it, and participants preferred when the notice was shown during or before app usage.

A notice shown in app use may be more salient to users, leading to the better recall we found. The fact that the notice interrupted the app usage may have helped the user pay attention to it. Further work is needed to examine habituation to notices shown during app use and determine how frequently the notices should be displayed, e.g.: on first run-only, periodically, depending on context, or depending on type of information.

When the notice is shown in the app store as one screenshot of many, it competes with other information on the screen (such as app title, developer, the install button), while the notice shown during app usage was a modal dialog that occupied the entire screen. Our second web survey

attempted to understand if this is why the app store was ineffective, by testing options to display the notice more prominently in the app store. The more prominent conditions in the app-store had better rates of recall than not showing the notice, but were still recalled significantly less than when the app was shown during app use. Since the app store options we tested in the follow-up survey are not currently available to app developers, we propose that app store designers consider offering new options for app-store privacy notices that allow the notices to be shown in larger sizes with fewer distractions.

While we found that participants did not remember the notices in the app store well, we believe there are nevertheless benefits to showing privacy notices in the app store. When shown in the app store, users can make informed decisions before they purchase or install an app. This may be particularly valuable for privacy-concerned users.

In this work, participants were directed to look at a specific app, which is similar to the real-life installation flow if a consumer has decided to install a specific app without comparing it to other apps. This may occur when an app was recommended by a friend, it was the top search result, or the app was linked to in a web article or app. Our results show that in these circumstances, users may ignore privacy information in the app store. However, as they use and evaluate an app, smartphone users may make more decisions about whether to continue to use the app, uninstall it, change the privacy preferences (when available), or even upgrade the app. If users have forgotten or never paid attention to the privacy notice information, they will not be able to make informed decisions about privacy. In these cases, a notice shown during app usage would be useful and memorable.

Chapter 7

Conclusion

Large amounts of data can be collected about users through smartphones. Data might be collected passively through sensors, such as GPS or gyroscopes, implying that users may not be aware of the collection. Alternatively, users may enter data themselves, but be unaware of with whom and how it is shared, as data is often shared, sold, or resold to a network of companies. The companies include not only app developers, but advertising companies, analytics companies, platforms, telecommunications companies, and the government.

The proliferation of data sharing in the smartphone data ecosystem may have negative ramifications for the smartphone user, creating privacy and security risks. We examine the risks from smartphone data sharing, and look for opportunities to mitigate the risks. Through these studies, my co-authors and I identified the risks and explored mitigations available to some of the stakeholders, including users, app developers, platform developers, and policy-makers. Users who wish to mitigate risks may rely on privacy notices for information. I examined smartphone privacy notices and identified opportunities to improve them.

To motivate the need to mitigate risks of harm, I revisit some of the harms from smartphone data sharing as discussed in Chapter 2. Based on expert interviews, the three most harmful results from smartphone data sharing were financial, physical and social harm. The harms varied in scope and could depend on the context. For example, the scope of social harms could vary from embarrassment that lasted a few minutes, to harms that result in divorce, jobs lost, or suicide. The harms faced by a person may vary by individual; not everyone has the same risk of being stalked or of facing financial discrimination from redlining or job discrimination.

7.1 Contributions of this thesis

This thesis contributed to identifying risks and mitigations to smartphone sharing in the following ways:

- I developed a taxonomy of privacy and security risks to smartphone users from data sharing based on expert interviews. This includes the most harmful and the most likely risks.
- I compiled and categorized mitigations that may reduce the above risks. Through this categorization, I find that notice is not sufficient to mitigate all the harms; data minimization and better security are needed.

- I examined the privacy and security behaviors of app developers. Through the qualitative and quantitative findings, I identified that small development companies, who are a large portion of app development companies, need usable tools implement good privacy and security behaviors.

This thesis contributed to improving smartphone privacy notices in the following ways:

- I uncovered participants' mental models of smartphone data sharing through a lab study.
- I developed guidelines for improving privacy notices, including specific suggestions on terms.
- I find that smartphone users have better recall of privacy notices shown during app use than in the app store, and recommend that notices be shown during app use (before the data is shared), in addition to in the app store, when the data sharing is particularly risky.
- Through participation in a process to develop public policy for a smartphone privacy notice, I offer insights to usability experts who would like to engage in public policy.

In the following sections, I offer guidance for public policymakers on how to mitigate the risks of smartphone data sharing. I will then discuss opportunities for future work.

7.2 Implications for Public Policy

Recognizing that privacy notice is an important part of mitigating the risks of smartphone privacy, I provide guidance to improve and develop privacy notices. I point out some concerns with smartphone privacy notices, and reasons why they are not sufficient in mitigating all risks.

I describe the role public policy should have in engaging all the stakeholders in privacy protection. I discuss additional opportunities to mitigate risks of smartphone data sharing that do not put the burden solely on the smartphone user. Instead of expecting users to read and understand notices, regulation could require other players in the smartphone data sharing ecosystem to improve user privacy.

7.2.1 Improving Notice

Improved smartphone notices must provide users with more relevant information. Users appreciate knowing three aspects about data sharing that are not included in most existing notices: *how often data is shared, with whom data is shared*, and *purpose of data sharing*. I suggest these three aspects be included in privacy notices or warnings in addition to *what* information is being shared. The first two aspects – how often and with whom – are specifically supported by our research. Participants appreciated knowing *how often* data was shared in the study described in Chapter 5. The high frequency of data sharing in simple games was surprising and alarming to our participants. In the same study, I found evidence that participants were interested in the second aspect of data sharing: *with whom* data was shared. Further supporting this, the majority of our study participants in Chapter 4 expressed a desire to be informed if data was shared with the third-party entities (such as government entities or consumer data resellers) listed in the NTIA Code of Conduct. Related work has found that users are interested in understanding the third aspect of data

sharing: *purpose*, or why data was shared [122]. I examined participants' perceptions of *purpose* in Chapter 5.

I found that users are more likely to remember the notice when it was shown during app usage than when it was shown in the app store in Chapter 6. To allow users to consider privacy as they make decisions about uninstalling, upgrading, or continuing to use an app, I recommend that crucial notices should be shown during app use for the most important types of data sharing. This should be in addition to providing full information before the app is downloaded.

7.2.2 Defining Usability for Smartphone Privacy Notices

Our research demonstrates that effective, understandable privacy notices are difficult to design, and should not be considered "low-hanging fruit." In Chapter 4, I demonstrated that usability tests are important for informing the design of consumer-facing notices. I also noted that policymakers struggled to define "usability" for privacy notices. Despite interest in running usability tests, policymakers may not know what metrics should be used to evaluate whether a privacy notice is usable. Different stakeholders have different concerns; app developers do not want to frighten users, while privacy advocates may want users to be fully informed of the risks. To meet everyone's goals, notices should neither make benign data collection seem inappropriately scary, nor should they hide important information relevant to the privacy or security of an individual.

I provide some guidelines on how to do usability tests to develop new smartphone privacy notices. Specifically, I describe what usability tests should measure. These guidelines should help develop privacy notices that are acceptable to all stakeholders and accurately convey the risks.

I recommend that smartphone privacy notices should be tested for the following qualities:

- Concrete user understanding of terms or icons – Are users able to match the information shown in the privacy notice with real-life app sharing situations? That is, are consumers able to understand any general categories described in the notice in concrete and specific terms? An example of testing for this was demonstrated in Chapter 4, in which users were asked to map real-life scenarios – such as an app sharing inseam information – to the categories enumerated in the NTIA Code of Conduct.

- User understanding of risks – Are users able to understand the risks that may occur from the data sharing? Are they able to put the sharing in context of their own lives and determine what impact it may have on them? In Chapter 5, I examined users' perceptions of risks after using a prototype that notified them of data leaks. Further attempts to design privacy notices should go further in exploring user understanding of risk, and may wish to specifically explore the vulnerable populations described in Chapter 2.

- User attention – Are users able to switch their attention to the notice? Does the timing, placement, or format of the notice attract users' attention sufficiently? There are several ways to measure attention, including through eye-tracking studies. I examined attention by measuring recall of the notice in Chapter 6.

- Habituation – Do users continue to switch attention to the notice over time or with repetition? I did not study the impact of habituation in this work. While the just-in-time notifications discussed in Chapter 5 were repeated, that study was not designed to examine habituation effects. However, I recognize that habituation is an important aspect of

real-world notice and choice [55]. Therefore, I recommend that this be included in future studies.

These four components of usability are neccessary for developing privacy notice designs. As designs are developed or iterated upon, these aspects are crucial in measuring whether the notices will be effective in helping users mitigate risks. While this not a complete list of the usability aspects of a privacy notice that can be examined, it is a starting point that can guide the design of usability studies.

Considering the App Developer

If app developers are expected to implement a privacy notice, they should be included in user studies. App developers can only implement privacy notices accurately if they understand the terms themselves or have access to specialists who do. As I discussed in Chapter 3, app developers may not have these resources or specialtize. In particular, further work could examine whether app developers have a concrete understanding of terms or icons in the privacy notice.

Furthermore, there is a risk that app developers don't know that they are sharing data, or with whom they are sharing. As I discussed in Chapter 3, app developers may not be aware of what is being shared with third-parties. Therefore, regulators should put pressure on third-party libraries, such as ad and analytics companies, to make their data collection policies obvious to the developers. We discuss some specific ideas in the next section.

7.2.3 Notice and Choice is Not Sufficient

Our work is not the first to propose that notices may not effectively mitigate risks. These concerns have been discussed in many domains from health to security (for example [42, 81]), and therefore I do not revisit the full discussion here. Instead, I focus on issues specific to privacy notice for smartphone data sharing.

I enumerated the risks to users from smartphone data sharing in Chapter 2. I found that a number of the risks cannot be mitigated by notice and choice alone. Notice and choice put the burden of risk mitigation on the smartphone users, but users are not always empowered to protect or verify their data. To reduce risk, the experts recommended improving security or reducing the amounts of data collected and stored.

Privacy notices should help users understand the risks of data sharing, but this may be difficult due to the contextual nature of the risks. The experts interviewed in Chapter 2 highlighted that the actual risks may depend heavily on the individual or the context. For example, an app that shares location with other users may be useful to parents who want to know where their children are, but may be dangerous if installed on the phone of a woman who is being stalked by an abuser. In these cases, it is unclear what risk a privacy notice should describe.

Even for the risks that can be mitigated through notice or choice, there are significant barriers to effective notice and choice. These go beyond the usability of specific notices, and cannot be designed away with better icons or timing. Users will not be able to effectively reduce risk without understanding the data sharing ecosystem. At this point, however, most of the data sharing is hidden from their view, and they are not aware of it. As I found in Chapter 5, many smartphone users were not aware of the amount of data shared with advertisers and analytics companies. In

Chapter 4, I found that survey participants showed low understanding of terms, such as "Consumer Data Reseller," that required an understanding of the data sharing ecosystem.

If we rely on notice, then we are expecting users to have a more complete understanding of the data sharing ecosystem and improved digital literacy. For consumers to protect their privacy interests, they must understand the implications of data sharing; they must be aware of and understand the data sharing ecosystem. Our studies have shown that most users do not have this level of understanding and awareness. Furthermore, to fully understand the privacy implications, users must understand data aggregation, how databases are used, and de-anonymization. I have not yet seen efforts to educate policymakers or app developers, much less the general US population, about these issues. Due to the small size of smartphone privacy notices and the limited attention given to them, I am skeptical that smartphone privacy notices are an appropriate place to provide education about data aggregation, and de-anonymization.

7.2.4 Selecting Whom to Regulate

As there are a number of stakeholders in the data sharing ecosystem, there are also various options when selecting whom to regulate. Regulation aimed at specific stakeholders may be more effective than that aimed at users. When considering whether it is effective to regulate a stakeholder, policymakers may want to consider several factors, including: the ability of government to enforce regulation, the ability of the stakeholder to create meaningful change, and existing market incentives that would encourage compliance to regulation. In the next paragraphs, I discuss each of these issues in the context of the smartphone data sharing ecosystem.

Enforcement: Regulation should be enforceable. Enforcement may be difficult when there are a large number (e.g. hundreds of thousands) of small, independent companies to audit, verify, and enforce compliance. As described in Chapter 3, app developers fall into that category. In comparison, there are currently very few major smartphone platforms in the United States (less than 5). There are many ad companies, perhaps in the hundreds, but this is still significantly fewer than the number of app developers. Therefore, I recommend that regulation focus on platforms, app stores, and data analytics and advertising companies, as opposed to focusing on app developers or the smartphone users. This would allow regulators to enforce policy efficiently.

Meaningful Change: Platforms are remarkably well placed to create meaningful change in protecting user privacy and security. The major smartphone platforms in the US control many important aspects of the smartphone experience. They manage the app stores, which is the major point at which users and app developers intersect and users make decisions. App stores can improve the location of privacy notices and information, and improve the security screening. Platforms also offer software development kits (SDKs), which are used by app developers to build apps. SDKs should include nudges to help developers improve privacy, either through privacy warnings in the build, increasing information about data collected, or design of documentation or function calls to nudge app developers to collect less information.

Incentives: Stakeholders that have market incentives to improve privacy and security may be likely to comply without lawsuits or costly enforcement efforts. Large platforms that encourage app development may have incentives to comply with good privacy regulation. For example, many large companies recognize that any apps used in the platform may have privacy and security flaws. In personal conversations, employees of Google and Facebook stated that that privacy flaws in apps reflect badly on their own companies' reputations. Users who have a bad privacy or security

experience with an app blame it on the platform, not on the app. Therefore, platforms may have an incentive to provide good privacy and security experiences for their users, and to improve privacy and security tools for app developers.

On the other hand, the lack of incentives may indicate that there are externalities that can only be addressed through legislation. Stakeholders, such as cloud storage companies, consumer data resellers, and advertising companies may have few market incentives to protect the privacy and security of smartphone users. They are not consumer-facing, and likely face little pressure from their direct users to protect consumer privacy. Unless the smartphone users begin to pay for apps without data collection over free apps that rely on advertising, ad and analytics companies are not likely to be influenced by user preferences. Public policy could encourage these companies to improve security, minimize data collection, and provide notice and choice.

7.2.5 Going Beyond User Privacy Notices

Currently, privacy notices put the burden on the smartphone user to read and understand the notices, and then make reasonable and rational decisions about the risks. I advocate that app developers and platforms have a greater role to play in protecting smartphone users. They should consider privacy and security when developing software. In Chapter 3, I uncovered several hurdles that prevent app developers from integrating privacy and security into their apps. I also pointed out that tools can be developed to encourage developers to make good privacy and security decisions.

Regulators can mandate that platforms or third-party data collectors develop the privacy tools for app developers and make them available at no additional cost to app developers who use their products. The tools can take the form of better documentation, improved programing interfaces, or privacy notices integrated into app development. In the next few paragraphs, I give some provide specific ideas of tools that could be mandatory.

Privacy policies of third party data collectors are not readable. However, even if privacy policies were readable, app developers have no way to monitor the actual data collected about their users by third-parties. Even conscientious app developers have little knowledge or control over what data is collected by third-parties and how frequently. Therefore, regulation should require that platforms include a tools that to allows app developers to monitor the data collection by third-party tools. I envision a tool for app developers that can be run during development stage, before launching the app. This tool would allow app developers to run their app and see the frequency of data collection from third parties. This tool might be similar to the prototype described in Chapter 5, in that it would provide both just-in-time notifications and summaries of data collection to app developers. Platforms could require that app developers run their app in this tool before submitting it to the app store. App developers would be better informed about data collection. If they are or their users would be uncomfortable with the data collection, app developers could then select different data collectors. The regulation would require that platform developers provide the tool, and that app stores require the tool be used before the app is published.

Regulation should also require all data collectors to provide control to users, allowing users to opt-out of data collection. As users are unfamiliar with third-party data collectors, control that is only available on the collectors' websites is unlikely to be useful. Therefore, smartphone platforms would include an interface to the opt-out. In this case, regulation would require that platform developers provide an interface, and that third-party data collectors allow meaningful choice through the interface.

Overall, I recommend that future public policy should focus less on asking app developers to write privacy notices and expecting users to read them, and more on opportunities for platforms and data handlers to improve privacy and security. These are just some, but not all, options for different stakeholders to help protect consumer privacy. There are many opportunities to reduce risk that don't rely on user notice and choice. Policymakers should examine the other privacy rights described by the White House (see Chapter 1 and [11]), which include focused collection and security. An in-depth study of the other stakeholders, perhaps similar to that done on app developers in Chapter 3, may help determine what regulation on purpose and data minimization would work and how it could be enforced.

7.3 Future Work

In this thesis, I tried to go beyond criticizing existing notices. Instead I found opportunities for improving notice. I developed a privacy notice and tested the timing of notices. However, I did not develop the perfect privacy notice that allows for user understanding and attention. Nor have I identified and tested all the possible mitigations available. In this section, I provide some areas of future work that may help improve smartphone privacy. I discuss both future work possible in the domain of human-computer interaction and user experience, which can be considered a direct extension of this work. I also provide some opportunities for future work in other domains that may complement and clarify our findings.

Notices that inform about risks: The smartphone privacy notices I designed and tested listed data types being shared or third-party entities with whom data is shared. While it may be difficult to include an explanation of the risks in the privacy notices due to the individual nature of risks, as I noted above, there may be opportunities to get closer to informing users about risks. For example, Android permissions were redesigned in the summer of 2014 to notify users about the most important permissions, and further work can be done to determine if this helps users understand their own risks.

Usability for App Developers: There are more opportunities to look at app-developer decision making around privacy and security. Most user studies have focused on consumers, but software development tools should be usable as well. As the myriad of independent developers are making privacy decisions, more work can and should be done on what nudges work to encourage privacy.

Stakeholder incentives: This thesis identified several stakeholders in the data sharing ecosystem. However, I did not thoroughly examine the incentives and objectives of each stakeholder. There may be room for economists to improve the understanding of the incentives for different stakeholders. This may include identifying and clarifying privacy externalities that result from the data sharing ecosystem. By describing the privacy concerns in economic framework including incentives and externalities, the case for public policy intervention may become more clear.

User Education: Our work has identified a gap in user understanding about data sharing. Knowledge of the data sharing ecosystem could be improved and user education about data management may be appropriate. In order for smartphone users to make informed decisions to protect their privacy, they would need to have some understanding of databases, data aggregation, machine learning, and anonymization techniques. Further study on how best to inform smartphone users (or the general population) about these issues is needed. Perhaps this work would include educators, teachers, or game designers.

Appendix A

Assessing and Mitigating the Risks of Smartphone Data Sharing

Interview Script

Thank you for agreeing to participate in this interview. As you have read in the consent form, your participation is voluntary. All your responses will be kept anonymous. You may stop the interview at any time. We will be recording the audio of the interview for transcription. Do you have any questions about the process before we begin?

I have turned on the audio recording. Please confirm you are ok with being recorded.

- What is your professional title and industry?
- Tell me a bit about your background and expertise.

I'm doing research on the risks of data flowing from smartphones, with the end goal of designing better user interfaces so users can make informed decisions. In particular, I'm interested in the harms, risks, and privacy concerns that could occur from smartphone data sharing. I'd like your thoughts on how harms can occur and what the smartphone user can do to prevent them. I am also interested in privacy concerns, which may not involve physical or financial harm, but that a smartphone user would find uncomfortable or undesirable. I will be asking about what users can do or need to know to prevent harms and risks. If there is information they need but don't currently have access to, please include that in your response.

- How would you define data sharing from smartphones, in terms of what the data is and where it goes?
- What harms could come to smartphone users from data sharing? [Expert should brainstorm list of harms]
- To recap, you've mentioned the following harms [interviewer repeats harms mentioned]:

I have some other harms that have been mentioned in research and by smartphone users. I'm mentioning these to help with brainstorming. Please feel free to add to this list or object to any items on the list.

- Malicious apps stealing financial information
- Apps sharing location information leading to stalking
- Business sharing their data sets that then becomes de-anonymized
- Data breach leading to financial harm or identity theft
- Apps sharing behavioral information with social circles, leading to embarrassment or problems with friends and family
- Un-encrypted data sent over a public network, leading to stealing of financial or sensitive information
- Premium texting or downloading unwanted software

Do you have any more to add?

- I'm interested in categorizing the harms into two lists: the most harmful and the most likely. Of all the harms we discussed, which are the most likely.
- Which ones could cause the most damage or harm?
- My mother just got a smartphone. What should she do and what does she need to know to prevent this harm?
- How should the smartphone interface or OS change to protect her?
- What should an app developer do to prevent the harms or concerns?
- What should regulators or public policy be doing to mitigate the risks? [If government stakeholder]
- Are any of these harms or concerns different if children are involved?
- Are there any other vulnerable populations?
- Looking forward 5-10 years, what will change when it comes to privacy and security?
- That concludes my interview questions. What questions do you think I should have asked, or should ask future experts?
- Is there anything else you would like to add?

Appendix B

Is Your Inseam a Biometric? A Case Study on the Role of Usability Studies in Developing Public Policy

Scenarios

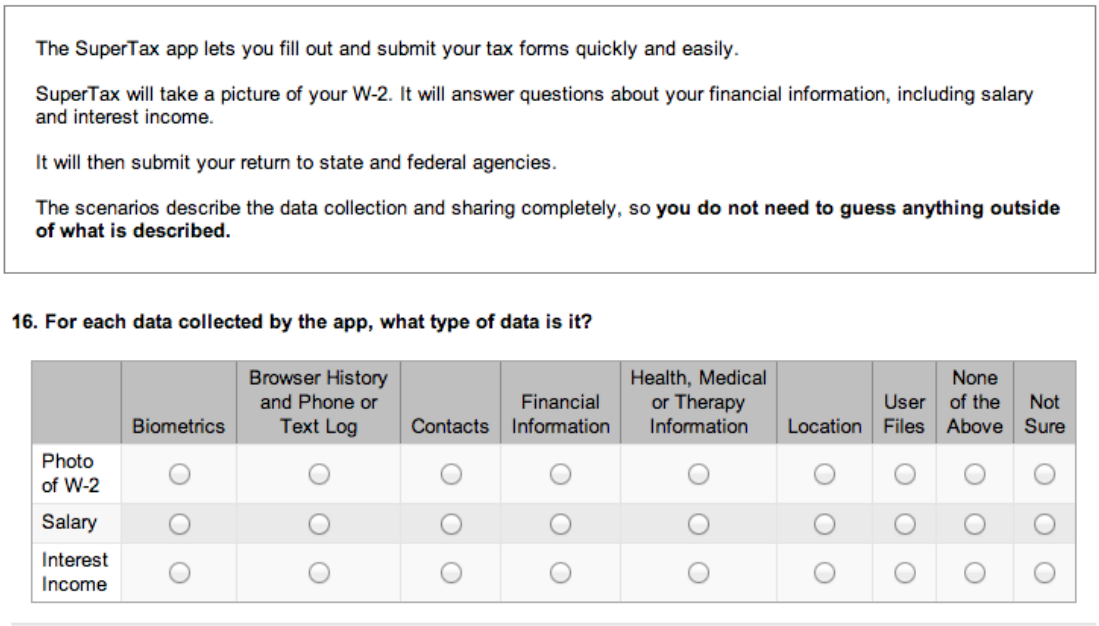

The SuperTax app lets you fill out and submit your tax forms quickly and easily.

SuperTax will take a picture of your W-2. It will answer questions about your financial information, including salary and interest income.

It will then submit your return to state and federal agencies.

The scenarios describe the data collection and sharing completely, so **you do not need to guess anything outside of what is described.**

16. For each data collected by the app, what type of data is it?

	Biometrics	Browser History and Phone or Text Log	Contacts	Financial Information	Health, Medical or Therapy Information	Location	User Files	None of the Above	Not Sure
Photo of W-2	○	○	○	○	○	○	○	○	○
Salary	○	○	○	○	○	○	○	○	○
Interest Income	○	○	○	○	○	○	○	○	○

Figure B.1: Screenshot of one scenario in the terms-only condition, showing how participants were asked to categorize the data types.

The text used to describe each scenario is presented here. A complete copy of the survey is available at **`http://alturl.com/vbmki`**. **HipClothes** The HipClothes app recommends cloth-

ing to you, and also shows you the stores closest to your location where you can find the clothes in your size.

The HipClothes app requests your inseam, waist size, and clothing preferences.

It also will share your information with two other clothing store chains that are owned by the same company.

Salsa The Salsa app allows you to make video calls, phone calls, text messages and include games, and picture sharing. Salsa stores your history in your online Salsa account.

Salsa collects your call, video, and text history, including copies of which pictures were shared, and information about which games were played.

Salsa will shows ads, and does so by sharing your information with advertising companies. Salsa will also share your information with AdMeMetric, which will resell information to companies that will provide you with coupons.

SuperTax The SuperTax app lets you fill out and submit your tax forms quickly and easily.

SuperTax will take a picture of your W-2. It will answer questions about your financial information, including salary and interest income.

It will then submit your return to state and federal agencies.

Fitness app The Fitness app integrates with your FitMonitor (FitMonitor is a special pedometer and activity monitor, purchased separately) to allow you to track and improve your fitness activities and level.

Fitness app will collect information on how many steps you have taken, how long you've slept, and allow you to enter you weight and body fat.

Fitness app will notify sports and health companies if you achieve certain goals, and these companies will send you valuable coupons as awards.

EasyApply This EasyApply app can be used to apply for government benefits such as Child Health Plus, Family Health Plus, Medicaid, and the Family Planning Benefit Program.

You will enter your income, work history, and whether you have any existing medical insurance and medical payments. You will also supply information about how many children you have, and your marital status.

EasyApply will save this information, and will submit your application to the state agency who will determine what benefits you and/or your children are eligible for.

CallCalendar The CallCalendar is an app that logs your phone activity and adds it to your Google Calendar.

You can select the type of calls to log (incoming, outgoing, and missed) and the calendar to log them in. CallCalendar will save your call log, including time, duration, and name of the person from the contact list.

CallCalendar will share your phone call information with your cellphone carrier so your cellphone carrier can improve its services. It will also share this information with Google Calendar.

GoodDriver The GoodDriver app is an application for your smartphone that will keep you and others safe on the roads.

It will use your GPS to detect your speed and location. It will use your gyroscope to detect road conditions (such as bumps). It will use your speed to tell you about traffic congestion and problems.

It shares information with that a company that specializes in traffic data so that congestion and problems can be predicted and analyzed. Your driving information will be sold to car insurance companies and car rental companies, who will offer you better rates for good driving.

FindMyKid The FindMyKid app can be installed on your childs phone to track its location and show you his or her whereabouts.

Without interrupting your child, you can see where he or she is at any time from your phone or on-line. FindMyKid app collects your child's location from his or her phone.

This app shares your child's location information with you (your phone). It will also share with local police, in case of emergency, with a simple button interface.

iTunes The popular iTunes app for playing music, and developed by Apple, is now available on Google Android phones.

You can enter song and artists names, which is stored by Apple. You can make purchases by entering your credit card information, which is saved by Apple for further purchases.

iTunes will share information about what you are playing with Facebook. Your songs are stored on the Apple iCloud service.

Bookstore The Bookstore app allows you to purchase books from your cell phone.

You will pay using a credit card and enter your home address where the book will be shipped. Bookstore app will save this information in your online account so that you can use Bookstore online or from any device.

It also shares information about your purchase with Facebook and GreatReading (an app that organizes local book clubs).

User Response Graphs

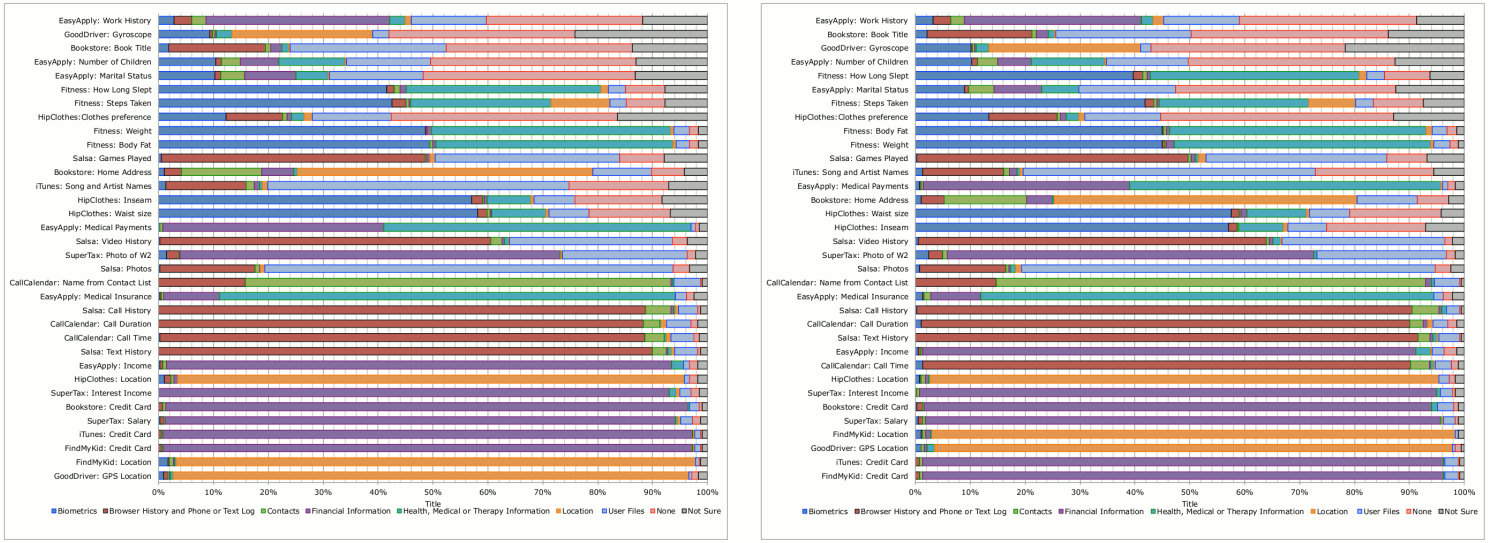

Figure 3: Participants' categorization of data types in parenthetical condition (left) and term-only condition (right).

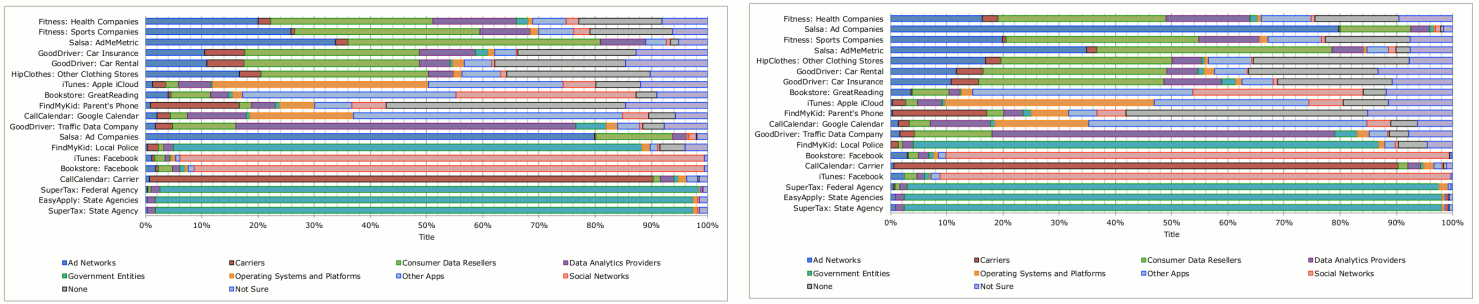

Figure 4: Participants' categorization of third-party entities in parenthetical condition (left) and term-only condition (right).

Appendix C

"Little Brothers Watching You:" Raising Awareness of Data Leaks on Smartphones

Interview Script

Welcome to our study. My name is ... and this is ... who will be taking notes.

Thank you for coming. Before we begin, let me tell you some important information about the study. We will be recording what is said in this interview, but everything will be anonymous. Your name and identifying information will be stored separately from your comments.

Please think out loud as you go through the tasks. That is, tell us what you are thinking as you go. Our goal is to evaluate our tools; not you. Everything you say, including confusion and questions, is very valuable to us.

Imagine that a family member or friend has just acquired an Android. They would like your advice on which game they should install. Imagine they will be playing these games during on the bus, waiting in the doctor's office, or maybe while they wait to meet you somewhere. Please take a minute to choose someone and tell us their relationship to you.

C.0.1 First Part of Interview

We will be giving you an Android phone with two free games, which we just installed before this interview. We are asking you to try these two games and decide which one you recommend to your friend. One game you are already familiar with is Angry Birds. The second game is called Toss It. Have you already played Toss It?

Screenshots from the install for each game are provided. You are welcome to refer to these in addition to actually playing the games. You will have up to 7 minutes to decide which game you prefer. Remember to think aloud.

[Participants played the games for 7 minutes or less.]

- Which game would you recommend and why?
- How would you describe each game to your friend?

- What would you write about each game in the app market?
- What information do you think was leaving the phone in the past 7 minutes while you played the games?
- Who was the information being shared with?
- Why was the information leaving the phone?
- Which application was sharing the data?
- What were you doing when the data was shared?

C.0.2 Second Part of Interview

This second phone has the same two games freshly installed. We have also installed an application that will inform you about information that is being shared [through notifications, such as vibration and the sound of water dropping]. You will have 7 minutes to play these games again. Imagine that you are evaluating these two games for your friend to use as he is waiting, for example, at the bus stop, at the doctors, or meeting you somewhere. Remember to think aloud. The app we installed is called Privacy Leaks, and you may look at it after playing the games.

[Participants played the games for 7 minutes or less. After playing games, participants were prompted to view Privacy Leaks.]

- Have your recommendations to your friend changed and why or why not?
- Would you describe these games the same way?
- What would you write about each game in the app market?
- Was there a relationship between when data was shared and what you were doing?
- Who was the information being shared with?
- Why was information leaving the phone?
- What type of information was being sent the most?
- Which application sent the most data and what data was being sent?

C.0.3 Third Part of Interview

Now imagine that these two games and Privacy Leak were on your own phone.

- Imagine that you had complete control over how your data was shared. What would you do?
- What if you could... would that be ok?

  - Stop information being sent when I'm in a particular location (e.g. at work, at home).
  - Stop information from being sent to particular companies or websites.
  - Stop information from being sent when I'm doing certain things (e.g. driving, sleeping).
  - Stop information from being sent by particular apps.
  - Stop certain information being sent, such as phone Id, or location, regardless of app or anything else.

The next few questions are specific to Privacy Leaks.

- Do you think the information given by Privacy Leaks was accurate?
- Would would you tell your friend about Privacy Leaks.
- What would you write in the app market about Privacy Leaks?
- What would you change about Privacy Leaks?

I'm going to ask a series of questions about the apps that you can respond to on a scale of 1-5, with one being "Strongly agree" and five being "Strongly disagree." [Interviewer places paper with likart-scale on table for reference.] Feel free to elaborate.

- The information provided by Privacy Leaks is useful.
- I understood what everything meant in Privacy Leaks.
- The sounds are distracting. [JIT condition only]
- The vibrations are distracting. [JIT condition only]
- I am likely to install an application like Privacy Leaks.
- The sounds would allow me to keep working or playing without interruption. [JIT condition only]
- The vibration would allow me to keep working or playing without interruption. [JIT condition only]
- The information was irrelevant.
- The information provided by this tool is confusing.

Final few questions

- Would you pay extra for a game that didn't send this information?
- Are there any benefits to you or your friend when the game shares information, and what are they?
- Are there any risks to you or your friend when the game shares information, and what are they?

Histograms of responses to Likert-scale questions

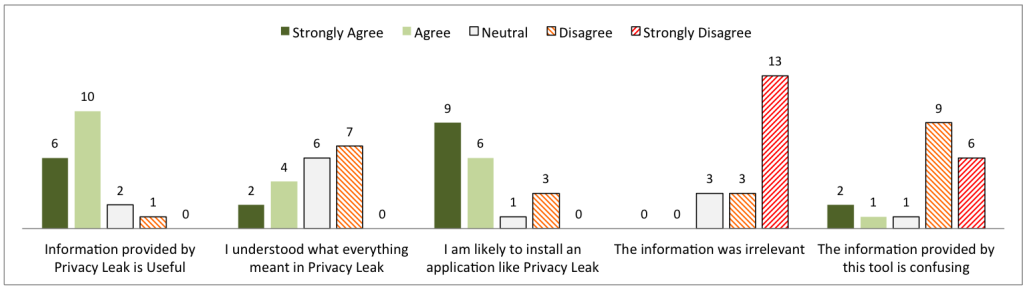

Figure C.1: Responses to Likert-scale questions about Privacy Leaks

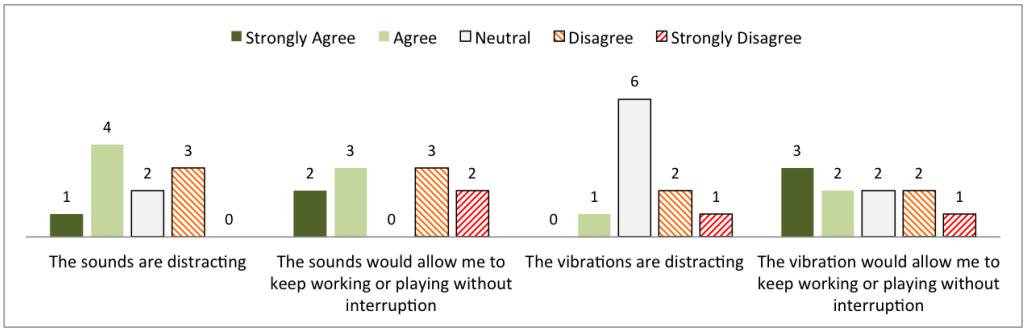

Figure C.2: Responses to Likert-scale questions about Just-In-Time Notifications

Appendix D

I Don't Remember, I Don't Recall: The Impact of Timing on Recall of Smartphone App Privacy Notices

Survey Questions

Page 1: App Review: Please tell us what you thought of this app.

1) How would you rate this app? (5 stars means a great app)
2) Please write a review of the app. (Imagine this is for the app store)*
3) I read the privacy policies of smartphone apps and websites.*
( ) Always
( ) Sometimes
( ) Rarely
( ) Never
4) How did you hear about this survey?*
5) Did someone who did the study before you did tell you anything about the study before you did it? If so, what did they tell you?*How did you hear about this survey?*

Page 2: Game Review Questions

[Questions 6-12 were shown in random order. We have marked the correct answers here.]

6) What was the title of the app?*
( ) US History Questions
( ) Inventions in US History Quiz
( ) History of US Inventions Quiz
(*) US Inventors History Quiz

7) Do you remember seeing the privacy notice?*
( ) No, not at all
( ) I remember most of it
( ) Vaguely
( ) Yes, I remember it well

8) Which of the following people were you asked about in this app?*
( ) Louis Armstrong
(*) Elijah McCoy
( ) Willie Brown
( ) Benjamin Banneker

9) What information was collected by the app?*
( ) Financial Information
( ) Which other apps are installed on my phone
(*) Browser History
( ) User Files
( ) I don't remember
( ) Nothing

10) With whom does the app share data?*
( ) Government entities
( ) Social Networks
(*) Ad networks
( ) Consumer Data Reseller
( ) I don't remember
( ) No one

11) Which of the following people were you asked about in this app?*
( ) Barack Obama
(*) Valerie L. Thomas
( ) Ralph Ellison
( ) Frederick McKinley Jones

12) What color was the background of the app?*
( ) Green
( ) Red
( ) Blue
( ) White
(*) Black

Page 3: Purpose of the Study:

[Debrief on purpose of the study, and image of privacy notice shown here].

13) Please select whether you agree or disagree with the following statements about when you saw the notification: [questions shown in random order]

The privacy notification was shown at a time when I could pay attention to it.
The privacy notification was shown at a time when I could make decisions about whether to allow the data collection.
The privacy notification disrupted my use of the app.
The privacy notification occurred at an unexpected time.

14)Please select whether you disagree or agree with the following statements about the information in the privacy notice: [questions shown in random order.]

I would want notifications like this when I download or use an app.
It is important for me to remember what the notification says while Im using the app over time.
This notification could be improved so I understand it better.
The privacy notification gave me information I care about.
I was surprised by what I learned from the privacy notification
I expected the app to collect my browser history and share it with ad networks.

15) Is there anything you would like to know that wasn't clear from the notification?
16) Is there anything else you would like to tell us about the privacy notification?

Bibliography

[1] Mobile app developers: Start with security. Bureau of Consumer Protection, **http://business.ftc.gov/documents/bus83-mobile-app-developers-start-security**.

[2] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. Organisation for Economic Co-operation and Development, **http://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm**, 1980.

[3] Privacy Online: Fair Information Practices in the Electronic Marketplace, A Report to Congress. Federal Trade Commission, **http://www.ftc.gov/reports/privacy2000/privacy2000.pdf**, 2002.

[4] *Evolution of a Prototype Financial Privacy Notice: A Report on the Form Development Project*. FTC and Kleimann Communication Group, Inc, 2006.

[5] Final Model Privacy Form Under the Gramm-Leach-Bliley Act; Final Rule. **http://www.gpo.gov/fdsys/pkg/FR-2009-12-01/html/E9-27882.htm**, December 2009.

[6] FTC staff issues privacy report, offers framework for consumers, businesses, and policymakers. Federal Trade Commission, **http://www.ftc.gov/opa/2010/12/privacyreport.shtm**, 12 2010.

[7] Mobile Web Application Best Practices, W3C Recommendation. W3c, **http://www.w3.org/TR/mwabp/**, 2010.

[8] Expert Elicitation Task Force White Paper. U.S. Environmental Protection Agency, 2011.

[9] App store stats summary. 148apps.biz, **http://148apps.biz/app-store-metrics/**, Sept 2012.

[10] Best practices for mobile application developers: App privacy guidelines. Future of Privacy Forum and the Center for Democracy And Technology, **https://www.cdt.org/files/pdfs/Best-Practices-Mobile-App-Developers.pdf**, July 12 2012.

[11] Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy. The White House, **http://www.whitehouse.gov/sites/default/files/privacy-final.pdf**, 2012.

[12] Fact sheet 2b: Privacy in the age of the smartphone. Privacy Rights Clearinghouse, **https://www.privacyrights.org/fs/fs2b-cellprivacy.htm**, Aug. 2012.

[13] How To Manage Privacy Settings On Your iPhone, iPad, iPod Touch In iOS 6. iPhoneHacks.com, **http://www.iphonehacks.com/2012/10/ios-6-manage-privacy-settings-iphone-ipad-ipod-touch.html**, October 2012.

[14] NTIA app group inching toward usability testing. Politico.com, **http://www.politico.com/morningtech/0113/morningtech9850.html**, Jan. 18 2012.

[15] Privacy design guidelines for mobile application development. GSMA Mobile Privacy, **http://www.gsma.com/publicpolicy/privacy-design-guidelines-for-mobile-application-development**, February 2012.

[16] Seizing opportunity: Good privacy practices for developing mobile apps. Office of the Privacy Commissioner of Canada, Information and Privacy Commissioner of (Alberta, British Columbia), **https://www.priv.gc.ca/information/pub/gd_app_201210_e.asp**, Oct. 2012.

[17] Fact Sheet 39: Mobile Health and Fitness Apps:What Are the Privacy Risks? Privacy Rights Clearinghouse, **https://www.privacyrights.org/fs/fs39/mobile-apps**, 2013.

[18] Mobile privacy: A better practice guide for mobile app developers. Office of Australian Information Commissioner, **http://www.oaic.gov.au/privacy/privacy-resources/privacy-guides/guide-for-mobile-app-developers**, Sept. 2013.

[19] Mobile privacy disclosures, building trust through transparency. Federal Trade Commission, **http://www.ftc.gov/os/2013/02/130201mobileprivacyreport.pdf**, 2013.

[20] Privacy multistakeholder process: Mobile application transparency. National Telecommunication and Information Administration, **http://www.ntia.doc.gov/category/privacy/u**, Jul. 2013.

[21] Android statistics, google play stats. appbrain.com, **http://www.appbrain.com/stats/stats-index**, May 2014.

[22] Emc privacy index. EMC.com, **http://www.emc.com/collateral/brochure/privacy-index-executive-summary.pdf**, May 2014.

[23] Smartphone app developers in the us: Market research report. **http://www.ibisworld.com/industry/smartphone-app-developers.html**, 2014.

[24] Alessandro Acquisti. Nudging privacy: The behavioral economics of personal information. *Security & Privacy, IEEE*, 7(6):82–85, 2009.

[25] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *Security & Privacy, IEEE*, 3(1):26–33, 2005.

[26] A. Adams and M.A. Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):40–46, 1999.

[27] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proc. of SOUPS*, page 9. ACM, 2013.

[28] Chaitrali Amrutkar, Patrick Traynor, and Paul Oorschot. Measuring SSL Indicators on Mobile Browsers: Extended Life, or End of the Road? In Dieter Gollmann and Felix Freiling, editors, *Information Security*, volume 7483 of *Lecture Notes in Computer Science*, pages 86–103. Springer Berlin Heidelberg, 2012.

[29] Carmela Aquino Andrew Lipsman. 2013 mobile future in focus 2013 mobile future in focus. *comScore.com*, February 2013.

[30] George J. Annas. HIPAA regulations, a new era of medical-record privacy? *New England Journal of Medicine*, 348(15), April 2003.

[31] Jennifer J Argo and Kelley J Main. Meta-analyses of the effectiveness of warning labels. *Journal of public policy and marketing*, 23(2):193–208, 2004.

[32] Charles Arthur. Is your private phone number on facebook? probably. and so are your friends'. The Guardian, **http://www.guardian.co.uk/technology/blog/2010/oct/06/facebook-privacy-phone-numbers-upload**, October 6, 2010.

[33] Rebecca Balebako, Cristian Bravo-Lillo, and Lorrie Faith Cranor. Is notice enough: Mitigating the risks of smartphone data sharing. *I/S: A Journal of Law and Policy for the Information Society*, Forthcoming, 2014.

[34] Rebecca Balebako and Lorrie Faith Cranor. Improving app privacy: Nudging app developers to protect user privacy. *IEEE Security & Privacy*, pages 60–63, July/August 2014.

[35] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Cranor, and Carolyn Nguyen. "A Lot of Little Brothers:" Measuring User Confidence in Smartphone Security and Privacy. In *Proc. of SOUPS*, 2013.

[36] Rebecca Balebako, Abigail Marsh, Jialiu Lin, Jason Hong, and Lorrie Faith Cranor. The Privacy and Security Behaviors of Smartphone App Developers. *Workshop on Usable Security (USEC)*, 2014.

[37] Rebecca Balebako, Richard Shay, and Lorrie Faith Cranor. Is Your Inseam a Biometric? Evaluating the Understandability of Mobile Privacy Notice Categories. *Cylab Technical Report*, 2013.

[38] Rebecca Balebako, Richard Shay, and Lorrie Faith Cranor. Is Your Inseam a Biometric? A Case Study on the Role of Usability Studies in Developing Public Policy. *Workshop on Usable Security*, 2014.

[39] M. Becher, F.C. Freiling, J. Hoffmann, T. Holz, S. Uellenbeck, and C. Wolf. Mobile security catching up? revealing the nuts and bolts of the security of mobile devices. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 96–111, 2011.

[40] Noam Ben-Asher, Niklas Kirschnick, Hanul Sieger, Joachim Meyer, Asaf Ben-Oved, and Sebastian Möller. On the need for different security methods on mobile phones. In *MobileHCI '11*, pages 465–473. ACM, 2011.

[41] Omri Ben-Shahar and Carl E. Schneider. The failure of mandated disclosure. *U. OF PENN. L. REV.*, 159, 2011.

[42] Omri Ben-Shahar and Carl E Schneider. *More Than You Wanted to Know: The Failure of Mandated Disclosure*. Princeton University Press, 2014.

[43] Z. Benenson, O. Kroll-Peters, and M. Krupp. Attitudes to IT Security when Using a Smartphone. *Proc. of the FedCSIS*, pages 1179–1183, 2012.

[44] Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal Ubiquitous Comput.*, 15(7):679–694, October 2011.

[45] Kevin Benton, L Jean Camp, and Vaibhav Garg. Studying the effectiveness of android application permissions requests. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2013 IEEE International Conference on*, pages 291–296. IEEE, 2013.

[46] Adam J Berinsky, Gregory A Huber, and Gabriel S Lenz. Using mechanical turk as a subject recruitment tool for experimental research. *Submitted for Review*, 2011.

[47] 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.

[48] Rainer Böhme and Stefan Köpsell. Trained to accept?: a field experiment on consent dialogs. In *CHI 2010*, pages 2403–2406. ACM, 2010.

[49] M. Böhmer, B. Hecht, J. Schöning, A. Krüger, and G. Bauer. Falling asleep with angry birds, facebook and kindle: a large scale study on mobile application usage. In *Proc. of MobileHCI*, pages 47–56. ACM, 2011.

[50] Matthias Böhmer, Christian Lander, Sven Gehring, Duncan P Brumby, and Antonio Krüger. Interrupted by a phone call: exploring designs for lowering the impact of call notifications for smartphone users. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 3045–3054. ACM, 2014.

[51] Theodore Book, Adam Pridgen, and Dan S Wallach. Longitudinal analysis of android ad library permissions. In *Proc. of MoST 2013*, 2013.

[52] Reinhardt A Botha, Steven M Furnell, and Nathan L Clarke. From desktop to mobile: Examining the security experience. *Computers & Security*, 28(3):130–137, 2009.

[53] Jan Lauren Boyles, Aaron Smith, and Mary Madden. Privacy and data management on mobile devices. *Pew Internet and American Life Project*, August 2012.

[54] Cristian Bravo-Lillo, Lorrie Cranor, Julie Downs, and Saranga Komanduri. Bridging the gap in computer security warnings: A mental model approach. *Security & Privacy, IEEE*, 9(2):18–26, 2011.

[55] Cristian Bravo-Lillo, Lorrie Cranor, Saranga Komanduri, Stuart Schechter, and Manya Sleeper. Harder to ignore? revisiting pop-up fatigue and approaches to prevent it. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014.

[56] Michael Buhrmester, Tracy Kwang, and Samuel D Gosling. Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011.

[57] M. Ryan Calo. Against notice skepticism in privacy (and elsewhere). *NOTRE DAME LAW REVIEW*, 647, 2012.

[58] Ann Cavoukian. Privacy and biometrics. Information and Privacy Commissioner, Ontario, September 1999.

[59] E. Chin, A.P. Felt, V. Sekar, and D. Wagner. Measuring user confidence in smartphone security and privacy. In *Proc. of SOUPS 2012*. ACM, 2012.

[60] S. Consolvo, J. Jung, B. Greenstein, P. Powledge, G. Maganis, and D. Avrahami. The Wi-Fi privacy ticker: improving awareness & control of personal information exposure on Wi-Fi. In *Proc. of Ubicomp 2010*, pages 321–330. ACM, 2010.

[61] Lorrie Cranor. *Web privacy with P3P*. O'Reilly Media, Inc., 2002.

[62] Lorrie Faith Cranor. A Framework for Reasoning About the Human in the Loop. *UPSEC*, 8:1–15, 2008.

[63] Lorrie Faith Cranor. The economics of privacy: Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.*, 10:273–445, 2012.

[64] Lorrie Faith. Cranor, P. Guduru, and M. Arjula. User interfaces for privacy agents. *ACM TOCHI*, 13(2):135–178, 2006.

[65] Amy Cravens. A demographic and business model analysis of today's app developer. Technical report, GigacomPro, funded by App Developers Alliance, September 2012.

[66] M. Dietz, S. Shekhar, Y. Pisetsky, A. Shu, and D.S. Wallach. Quire: Lightweight provenance for smart phone operating systems. In *Proc. of 20th USENIX Security Symposium*. USENIX Association, 2011.

[67] Manuel Egele, David Brumley, Yanick Fratantonio, and Christopher Kruegel. An Empirical Study of Cryptographic Misuse in Android Applications. In *Proc. of CCS '13*.

[68] S. Egelman, A.P. Felt, and D. Wagner. Choice architecture and smartphone privacy: There's a price for that. In *Workshop on the Economics of Information Security (WEIS)*, 2012.

[69] Serge Egelman, Adrienne Porter Felt, and David Wagner. In *The Economics of Information Security and Privacy*, pages 211–236. Springer, 2013.

[70] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proc. of CHI*, pages 319–328. ACM, 2009.

[71] W. Enck, P. Gilbert, B.G. Chun, L.P. Cox, Jaeyeon Jung, P. McDaniel, and A.N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proc. OSDI 2010*, pages 1–6. USENIX Association, 2010.

[72] William Enck, Damien Octeau, Patrick McDaniel, and Swarat Chaudhuri. A study of android application security. In *USENIX security symposium*, 2011.

[73] Sascha Fahl, Marian Harbach, Thomas Muders, Lars Baumgärtner, Bernd Freisleben, and Matthew Smith. Why Eve and Mallory Love Android: An Analysis of Android SSL (in)Security. In *Proc. of CCS '12*.

[74] Peter Farago. Rise of the New Middle Class: Indie iPhone App Developers, Part I. flurry.com, **http://www.flurry.com/bid/24163/Rise-of-the-New-Middle-Class-Indie-iPhone-App-Developers-Part-I**.

[75] A.P. Felt, S. Egelman, and D. Wagner. I've got 99 problems, but vibration ain't one: A survey of smartphone users' concerns. In *Proc. SPSM*, 2012.

[76] A.P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: User attention, comprehension, and behavior. *Proc. of SOUPS*, 2012.

[77] Elizabeth Fife and Juan Orjuela. The privacy calculus: Mobile apps and user perceptions of privacy and security. *International Journal of Engineering Business Management*, 5(6):7, 2012.

[78] Kelsey Finch. The All-New IAPP Mobile App Privacy Tool. IAPP, **https://privacyassociation.org/news/a/the-all-new-iapp-mobile-app-privacy-tool/**.

[79] Giulia Fiorese, Michela Catenacci, Elena Verdolini, and Valentina Bosetti. Advanced biofuels: Future perspectives from an expert elicitationsurvey. *Energy Policy*, 2013.

[80] Huiqing Fu, Yulong Yang, Nileema Shingte, Janne Lindqvist, and Marco Gruteser. A field study of run-time location access disclosures on android smartphones. In *Usable Security (USEC)*, 2014.

[81] Archon Fung, Mary Graham, and David Weil. *Full disclosure: The perils and promise of transparency*. Cambridge University Press, 2007.

[82] Susanne Furman and Mary Theofanos. Preserving privacy–more than reading a message. In *Universal Access in Human-Computer Interaction. Design for All and Accessibility Practice*, pages 14–25. Springer, 2014.

[83] C. Gates, N. Li, H. Peng, B. Sarma, Y. Qi, R. Potharaju, C. Nita-Rotaru, and I. Molloy. Generating summary risk scores for mobile applications. *Dependable and Secure Computing, IEEE Transactions on*, 11(3):238–251, May 2014.

[84] C.S. Gates, J. Chen, N. Li, and R.W. Proctor. Effective risk communication for android apps. *Dependable and Secure Computing, IEEE Transactions on*, 11(3):252–265, May 2014.

[85] Nathaniel S. Good, Jens Grossklags, Deirdre K. Mulligan, and Joseph A. Konstan. Noticing Notice: A Large-scale Experiment on the Timing of Software License Agreements. In *Proc. of CHI*. ACM, 2007.

[86] M.C. Grace, W. Zhou, X. Jiang, and A.R. Sadeghi. Unsafe exposure analysis of mobile in-app advertisements. In *Proc. of WiSec*, 2012.

[87] Ralph Gross and Alessandro Acquisti. Information revelation and privacy in online social networks. In *Proc. of WPES*, pages 71–80. ACM, 2005.

[88] Joseph Lorenzo Hall. NTIA Multistakeholder Process Delivers Increased App Transparency. **https://cdt.org/blog/ntia-multistakeholder-process-delivers-increased-app-transparency/**.

[89] Jun Han, Emmanuel Owusu, Le T Nguyen, Adrian Perrig, and Joy Zhang. Accomplice: Location inference using accelerometers on smartphones. In *COMSNETS*, pages 1–9. IEEE, 2012.

[90] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 2647–2656, New York, NY, USA, 2014. ACM.

[91] Kamala D Harris. Privacy on the go, recommendations for the mobile ecosystem. Attorney General California Department of Justice, **http://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/privacy_on_the_go.pdf**, 2013.

[92] M. Hastak and M. Culnan. Online behavioral advertising "icon" study. *Future Of Privacy Forum*, 2010.

[93] Giles Hogben and Marnix Dekker. Smartphones: Information security risks, opportunities and recommendations for users. *European Network and Information Security Agency*, 710(01), 2010.

[94] P. Hornyack, S. Han, Jaeyeon Jung, S. Schechter, and D. Wetherall. These aren't the droids you're looking for: retrofitting android to protect data from imperious applications. In *Proc. CCS*, pages 639–652. ACM, 2011.

[95] White House. Big data: Seizing opportunities, preserving values. **http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf**, 2014.

[96] Sami Hyrynsalmi, Arho Suominen, Tuomas Mäkilä, Antero Järvi, and Timo Knuutila. Revenue models of application developers in android market ecosystem. In *Software Business*, pages 209–222. Springer, 2012.

[97] Panagiotis G. Ipeirotis. Demographics of Mechanical Turk. Technical Report CeDER-10-01, New York University, 2010.

[98] Betsy Isaacson. Immersion, An MIT Media Lab Creation, Uses Email Metadata To Map Your Connections. Huffington Post, **http://www.huffingtonpost.com/2013/07/10/immersion-email-metadata_n_3567984.html**, July 10, 2013.

[99] Lee Jae-Won. Retail stores track consumers' smartphones through wi-fi.

[100] Shubham Jain and Janne Lindqvist. Should I Protect You? Understanding Developers Behavior to Privacy-Preserving APIs. In *Workshop on Usable Security (USEC) 2014*, 2014.

[101] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proc. of CHI*, pages 471–478. ACM, 2004.

[102] Carlos Jensen, Colin Potts, and Christian Jensen. Privacy practices of internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63:203 – 227, 2005.

[103] Adam N. Joinson. Looking at, looking up or keeping up with people?: motives and use of facebook. In *Proc. of CHI '08*, 2008.

[104] Chuck Jones. Apple's app store about to hit 1 million apps. Forbes.com, **http://www.forbes.com/sites/chuckjones/2013/12/11/apples-app-store-about-to-hit-1-million-apps/**.

[105] Jaeyeon Jung, Seungyeop Han, and David Wetherall. Short paper: Enhancing mobile application permissions with runtime feedback and constraints. In *Proc. of the workshop on Security and Privacy in Smartphones and Mobile devices*, 2012.

[106] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices: an online study of the nutrition label approach. Proc. of CHI 2010, pages 1573–1582. ACM, 2010.

[107] 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*, 2012.

[108] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. Privacy as part of the app decision-making process. In *CHI 2013*, 2013.

[109] Braden Kowitz and Lorrie Cranor. Peripheral privacy notifications for wireless networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 90–96. ACM, 2005.

[110] Balachander Krishnamurthy and Craig E. Wills. Characterizing privacy in online social networks. In *Proc. of WOSN '08*, 2008.

[111] Mary Kynn. The 'heuristics and biases' bias in expert elicitation. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171(1):239–264, 2008.

[112] Mariantonietta La Polla, Fabio Martinelli, and Daniele Sgandurra. A survey on security for mobile devices. *Communications Surveys Tutorials, IEEE*, 15(1):446–471, 2013.

[113] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research methods in human-computer interaction*. Wiley, 2010.

[114] Choon Seong Leem, Hyung Sik Suh, and Dae Seong Kim. A classification of mobile business models and its applications. *Industrial Management & Data Systems*, 104(1):78–87, 2004.

[115] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. Why johnny can't opt out: A usability evaluation of tools to limit online behavioral advertising. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, pages 589–598. ACM, 2012.

[116] Pedro G. Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorescu, and Lorrie Faith Cranor. What matters to users? factors that affect users' willingness to share information with online advertisers. In *Proc. SOUPS*, 2013.

[117] Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, and Guzi Xu. What do online behavioral advertising privacy disclosures communicate to users? In *Proc. of WPES*, WPES '12, pages 19–30, New York, NY, USA, 2012. ACM.

[118] Ilias Leontiadis, Christos Efstratiou, Marco Picone, and Cecilia Mascolo. Don't kill my ads!: balancing privacy in an ad-supported mobile application market. In *Proc. of HotMobile '12*, pages 2:1–2:6. ACM, 2012.

[119] Alan Levy and Manoj Hastak. Consumer comprehension of financial privacy notices: A report on the results of the quantitative testing. *Federal Trade Commission*, pages 62890–62994, December 2008.

[120] Clayton Lewis and Jutta Treviranus. Public policy and the global public inclusive infrastructure project. *interactions*, 20(5):62–66, 2013.

[121] Ilaria Liccardi, Joseph Pato, and Daniel J Weitzner. Improving user choice through better mobile apps transparency and permissions analysis. *Journal of Privacy and Confidentiality*, 5(2):1, 2014.

[122] J. Lin, S. Amini, J.I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and purpose: Understanding users' mental models of mobile app privacy through crowdsourcing. *UbiComp 2012*, 2012.

[123] Jialiu Lin. Understanding and capturing people's mobile app privacy preferences. Technical Report Ph.D Thesis CMU-CS-13-127.

[124] David M Lodge, Susan Williams, Hugh J MacIsaac, Keith R Hayes, Brian Leung, Sarah Reichard, Richard N Mack, Peter B Moyle, Maggie Smith, David A Andow, et al. Biological invasions: recommendations for US policy and management. *Ecological Applications*, 16(6):2035–2054, 2006.

[125] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 15(4):336–355, 2004.

[126] Moxie Marlinspike. Why 'I have nothing to hide' is the wrong way to think about surveillance. Wired, June 13, 2013.

[127] Caroline Mayer. Don't be dumb about smartphone privacy. *forbes.com*, 3 2013.

[128] Harry McCracken. Who's Winning, iOS or Android? All the Numbers, All in One Place. **http://techland.time.com/2013/04/16/ios-vs-android/**, April 16 2013.

[129] Aleecia McDonald and Jon Peha. Track gap: Policy implications of user expectations for the 'Do Not Track' internet privacy feature. In *TPRC*, 2011.

[130] Aleecia M. McDonald and Lorrie F. Cranor. The Cost of Reading Privacy Policies. *I/S: A Journal of Law and Policy for the Information Society*, 4(3):540–565, 2008.

[131] Aleecia M. McDonald and Lorrie Faith Cranor. Americans' attitudes about internet behavioral advertising practices. In *ACM Workshop on Privacy in the Electronic Society (WPES)*, October 2010.

[132] Mary A. Meyer and Jane M. Booker. *Eliciting and analyzing expert judgment: a practical guide*, volume 7. Society for Industrial Mathematics, 1987.

[133] M.G. Morgan, B. Fischhoff, A. Bostrom, and C.J. Atman. *Risk communication: A mental models approach*. Cambridge University Press, 2001.

[134] Joe Mullin. CA to app devs: get privacy policies or risk $2500-per-download fines. arstechnica.com, **http://arstechnica.com/tech-policy/2012/12/ca-to-app-devs-get-privacy-policies-or-risk-2500-per-download-fines/**, Dec 2012.

[135] Ildar Muslukhov, Yazan Boshmaf, Cynthia Kuo, Jonathan Lester, and Konstantin Beznosov. Understanding Users' Requirements for Data Protection in Smartphones. *ICDE 2012*, pages 228–235, April 2012.

[136] Alexios Mylonas, Anastasia Kastania, and Dimitris Gritzalis. Delegate the smartphone user? security awareness in smartphone platforms. *Computers & Security*, 2012.

[137] Alexios Mylonas, Marianthi Theoharidou, and Dimitris Gritzalis. Assessing privacy risks in android: A user-centric approach. In *Risk Assessment and Risk-Driven Testing*, Lecture Notes in Computer Science, pages 21–37. Springer International Publishing, 2014.

[138] Saira Nayak. What's next for the NTIA mobile app transparency code? TrustE blog, **http://www.truste.com/blog/2013/08/01/ntia-mobile-app-transparency-code**, Aug. 1 2013.

[139] Lawrence Norden, Whitney Quesenbery, and David C. Kimball. Better design, better elections. Brennan Center for Justice at New York University School of Law **http://www.brennancenter.org/publication/better-design-better-elections**, 2012.

[140] Gabriele Paolacci, Jesse Chandler, and Panagiotis Ipeirotis. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5):411–419, 2010.

[141] Irene Pollach. What's wrong with online privacy policies? *Commun. ACM*, 50(9):103–108, September 2007.

[142] S. Prabhakar, S. Pankanti, and A.K. Jain. Biometric recognition: security and privacy concerns. *IEEE Security & Privacy*, 1(2):33–42, 2003.

[143] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. Large-scale assessment of mobile notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 3055–3064, New York, NY, USA, 2014. ACM.

[144] Roman Schlegel, Kehuan Zhang, Xiaoyong Zhou, Mehool Intwala, Apu Kapadia, and X Wang. Soundcomber: A stealthy and context-aware sound trojan for smartphones. In *NDSS 2011*, pages 17–33, 2011.

[145] Steve Sheng, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Cranor, and Jason Hong. Improving phishing countermeasures: An analysis of expert interviews. In *eCrime Researchers Summit, 2009. eCRIME'09.*, pages 1–15. IEEE, 2009.

[146] Manya Sleeper, Rebecca Balebakok, Sauvik Das, Amber McConahy, Jason Wiese, and Lorrie Faith Cranor. The post that wasn't: Exploring self-censorship on facebook. In *16th ACM Conference on Computer Supported Cooperative Work and Social Computing*, Feb 2013.

[147] Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquist, Lorrie Faith Cranor, and Norman Sadeh. i read my twitter the next morning and was astonished: a conversational perspective on twitter regrets. *Proc. of CHI*, 2013.

[148] Daniel Smilkov, Deepak Jagdish, and César Hidalgo. Immersion, a people-centric view of yout email life. **https://immersion.media.mit.edu/**, 2013.

[149] Daniel Solove and Woodrow Hartzog. The FTC and the New Common Law of Privacy. *COLUMBIA L. REV*, 2014.

[150] Daniel J Solove. "I've Got Nothing to Hide" and Other Misunderstandings of Privacy. *San Diego law review*, 44, 2007.

[151] Daniel J. Solove. Five myths about privacy. Washington Post, June 13, 2013.

[152] Daniel J Solove. Privacy self-management and the consent dilemma. *Harv. L. Rev.*, 126:1879–2139, 2013.

[153] A Pew Research Center/USA TODAY Survey. Public Split over Impact of NSA Leak, But Most Want Snowden Prosecuted. **http://www.people-press.org/files/legacy-pdf/6-17-13%20NSA%20release.pdf**, June 17, 2013.

[154] Eran Toch, Justin Cranshaw, Paul Hankes Drielsma, Janice Y. Tsai, Patrick Gage Kelley, James Springfield, Lorrie Cranor, Jason Hong, and Norman Sadeh. Empirical models of privacy in location sharing. In *Proc. Ubicomp*, pages 129–138, 2010.

[155] Daniel W Turner. Qualitative interview design: A practical guide for novice investigators. *The Qualitative Report*, 15(3):754–760, 2010.

[156] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proc. SOUPS*, 2012.

[157] J. Urban, C. Hoofnagle, and S. Li. Mobile phones and privacy. *UC Berkeley Public Law Research Paper*, 2012.

[158] Jennifer Valentino-Devries, Jeremy Singer-Vine, and Ashkan Soltani. Websites vary prices, deals based on users' information. Wall Street Journal, **http://online.wsj.com/article/SB10001424127887323777204578189391813881534.html**, December 24, 2012.

[159] Narseo Vallina-Rodriguez, Jay Shah, Alessandro Finamore, Yan Grunenberger, Konstantina Papagiannaki, Hamed Haddadi, and Jon Crowcroft. Breaking for commercials: characterizing mobile advertising. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, pages 343–356. ACM, 2012.

[160] Ric Velez. Lookout Open Sourced Its 'Private Parts,' You Should, Too. **https://blog.lookout.com/blog/2014/03/12/open-source-privacy-policy/**, March 12 2014.

[161] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. "I regretted the minute I pressed share:" a qualitative study of regrets on Facebook. In *Proc. of SOUPS '11*, 2011.

[162] Michael S Wogalter, Dave DeJoy, and Kenneth R Laughery. *Warnings and risk communication*. CRC Press, 2005.

[163] Nan Xu, Fan Zhang, Yisha Luo, Weijia Jia, Dong Xuan, and Jin Teng. Stealthy video capturer: a new video-based spyware in 3g smartphones. In *Proc of Wisec 09*, pages 69–78. ACM, 2009.

[164] Li Zhang, Dhruv Gupta, and Prasant Mohapatra. How expensive are free smartphone apps? *SIGMOBILE Mob. Comput. Commun. Rev.*, 16(3):21–32, December 2012.

[165] Yajin Zhou, Xinwen. Zhang, Xuxian Jiang, and Vincent Freeh. Taming information-stealing smartphone applications (on Android). *Proc. of TRUST 2011*, pages 93–107, 2011.