Author's Note: Updated Version (Last revised March 2013)

"How Come I'm Allowing Strangers To Go Through My Phone?"— Smartphones and Privacy Expectations

Jennifer King
jenking@ischool.berkeley.edu
University of California, Berkeley
School of Information
102 South Hall
Berkeley, CA 94720-4600

ABSTRACT

This study examines the privacy expectations of smartphone users by exploring two specific dimensions to smartphone privacy: participants' concerns with *other people* accessing the personal information stored on their smartphones, and *applications* accessing this information via platform APIs. We interviewed 24 Apple iPhone and Google Android users about their smartphone usage, using Altman's theory of boundary regulation and Nissenbaum's theory of contextual integrity to guide our inquiry. Our contribution is a contextually-situated examination of smartphone users' privacy preferences and expectations based upon real world usage. Overall, we found that the default flows of smartphone APIs defy users' privacy expectations. In contradiction to the assumptions made by many mobile privacy studies, we found that our participants were far less concerned with sharing their location compared to other types of information available through the platforms' APIs. Further, we found that not only did some of our participants not understand the capabilities of applications, they also relied upon a number of assurance structures (sometimes inaccurately) to assuage their privacy concerns when selecting applications. We conclude with suggestions for platforms and application developers to make smartphone APIs and applications function in a manner that supports users' privacy expectations, as well as a call to use theoretically grounded methods for mobile privacy research.

Author Keywords

Smartphones; touch phones; privacy; mental models, applications; apps; user expectations; iPhone; iOS; Apple; Android; Google.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Security; Human Factors; Legal Aspects.

Symposium On Usable Privacy and Security (SOUPS) 2012, July 11-13, 2012, Washington, DC, USA.

1. INTRODUCTION

In 2012, the threshold of smartphone ownership in the U.S. surpassed 50 percent of all mobile subscribers. Google's Android, the most popular smartphone operating system in the U.S., claimed over fifty percent of the market share, followed by Apple's iOS (32 percent).[23] Both platforms claimed to have over a half-million applications created by third parties in their respective online stores. While smartphones allowing the installation of third party applications predate the iPhone, the introduction of Apple's iOS in 2007 ushered in a new era of application adoption by popularizing "apps" and attracting thousands of developers. The relative ease of developing applications for both Android and iOS has led to an unprecedented number of developers creating them globally. Much of the smartphone's popularity is due to its all-purpose usefulness, with users routinely using their devices to bridge diverse facets of their lives. This boundary-crossing usage creates a broad swath of personal information on smartphones that spans multiple life dimensions.

There are key elements of mobile platforms that increase the risks to the privacy of the personal information users store on their smartphones. We define information privacy risk in this paper as *access to one's personal information without express knowledge or consent.* The amount and type of personal information made accessible to and collectible by applications by default through the platforms' APIs increases users' exposure. Some applications need access to various categories of personal information accessible through the APIs to provide desired functionality; others do not but access it regardless. Developers can access and copy user contact lists, text messages, photos, and more directly through the APIs on both platforms. Both Android and iOS do seek to protect customers' information through contractual agreements with developers that limit developers' use. But the effectiveness of these agreements relies upon enforcement, and many application developers, particularly independent and/or inexperienced ones may lack experience handling customer information in a manner that mitigates security and privacy risks. Further, many application developers

partner with third party advertising services, which often involves sharing their customer's information.

The negative public reaction to incidents where developers accessed personal information without users' explicit consent reveals that both smartphone users and policymakers are concerned with its increased risk of collection and misuse. It also suggests that current platform designs not only do not capture users' expectations about applications' access to personal information, but also contradicts them. In this paper we seek to broadly understand smartphone users' privacy concerns by exploring the expectations that form the basis for their preferences. We draw upon two theories of privacy to examine smartphone users' concerns with *other people* and *applications* accessing the personal information stored on their smartphones. Using Irwin Altman's theory of boundary regulation and Helen Nissenbaum's theory of contextual integrity to guide our inquiry, we explore the divergence between participants' privacy expectations and management of access to their smartphones and their personal information. Using two qualitative methods—structured interviews and card sorting exercises—we explore the information privacy expectations of 24 Apple (iPhone) and Google (Android) smartphone users.

Most studies of smartphone privacy have focused primarily on concerns with location information, explored privacy issues within the context of device security, or limited their inquiry to narrow aspects of smartphone usage. In contrast, we cast a broad net by exploring participants' use and relationship to their smartphones, and participants' privacy concerns with others' (both other people and applications') access to varied categories of personal information entrusted to their smartphones. Our contribution is a contextually-situated examination of smartphone users' privacy preferences and expectations based upon their real world usage. Overall, we found that the default flows of smartphone APIs defy users' privacy expectations. In contradiction to the assumptions made by many mobile privacy studies, we found that our participants were far less concerned with sharing their location compared to other types of information available through the platforms' APIs. Further, we found that not only did some of our participants not understand the capabilities of applications, they also relied upon a number of assurance structures (sometimes inaccurately) to assuage their privacy concerns when selecting applications. We conclude with suggestions for platforms and application developers to make smartphone APIs and applications function in a manner that supports users' privacy expectations, as well as a call to use theoretically grounded methods for mobile privacy research.

2. BACKGROUND

In this section, we briefly review differences in application management between the two platforms and how they directly affect the accessibility of users' personal information. The two platforms' APIs allow developer access to similar information. Space constrains our review and comparison of the two platforms, but some of the information types that many users would consider sensitive—contacts, text messages, photos—are accessible to applications under the default settings in both platforms. A primary difference between the two APIs lies in how they grant and manage access to data. It differs at two levels: how applications are given access to the platform, and how users are presented with information about access.

Figure 1: iOS Privacy Preferences Screen

Apple reviews each application for content, quality, and security before allowing it into their App Store (the only approved method for obtaining applications on the iPhone). Application developers must state in their terms of service and privacy policy (if they have one; many do not [34][1]) what information they collect and how they use it. Apple claims to reject applications with data collection policies that are "inconsistent" with the intent of the application. iOS 5 required user consent at runtime for an application to access a user's present location (given through a dialog box), but otherwise all other data access occurred in the background.

This changed in iOS 6; user consent is now required for access to additional data types, and Apple provides an interface for reviewing or managing application access to data beyond location.[10] Prior to this change, if iOS users wanted information about which information an application collects they had to resort to reading the application's terms

[1]The lack of privacy policies is changing; as of 10/30/12, California Attorney General Kamala Harris began enforcing California's law requiring the posting of privacy policies on websites, arguing the law also applied to applications available in online stores. Please see: http://oag.ca.gov/news/press-releases/attorney-general-kamala-d-harris-notifies-mobile-app-developers-non-compliance

of service (TOS) or privacy policy (if the application had one). These documents are often dozens of pages long on an iPhone's small screen and written in legal language that is challenging for most users to comprehend.

In contrast, Google does not proactively review applications for content, quality, or data collection practices, though they do scan applications for malware. Anyone can submit an application to the Google Play store (previously the Android Market), and Google also allows applications to be downloaded from external sources. During the installation process Android presents users with a non-skippable screen displaying the data categories that the application is requesting to access. The installation and consent process is binary; a user must either accept all the requested permissions or forego installing the application.

Figure 2: iOS 6 Privacy Access Dialog

Applications running on iOS and Android routinely access and collect both personal information stored on users' phones via their APIs as well as collect information about user behavior within applications.[32] Both the platforms and application developers have come under fire for the use of both types of information without obtaining clear consent from users. For example, in 2011 *The Guardian* reported that Facebook's mobile application copied a phone's address book to one's Facebook account, leading to confusion and anger from users surprised by the default export of information from their phone to Facebook's application and ultimately their online profiles.[2] Similarly, in February 2012, a software developer blogged that the social networking application Path was transmitting his entire iPhone address book to their servers without consent, noting "I feel quite violated that my address book is being held remotely on a third-party service."[31] The observation resulted in a media firestorm. Under pressure from the U.S. Congress, Apple announced a change in practice requiring a runtime opt-in for application access to a user's address book (as well as other data types) in iOS 6.[26] In February 2013, the Federal Trade Commission announced a settlement with Path after investigating the company for violating the FTC Act by failing to disclose the collection of users' phone contacts in either the app's UI or in its privacy policy. [33]

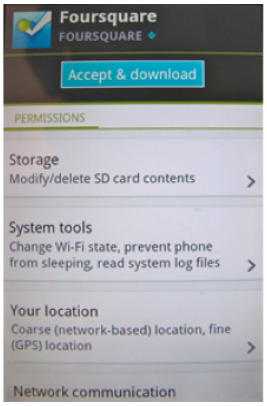

Figure 3: Android Permission Screen

3. RELATED WORK

While there are a handful of studies that examine either general smartphone usage or specific privacy issues (primarily with respect to identifying and/or sharing one's location), there is a paucity of work examining users' mobile privacy expectations across multiple contexts of use. Our research seeks to fill that gap.

Presently, there are two published qualitative studies that richly document the varied contexts of smartphone use, though neither explicitly examines users' privacy concerns or expectations. Matthews *et al*.[19] and Barkhuus and Polichar [4] provide qualitative examinations of smartphone use in everyday life. Matthews found smartphone use was highly contextual, with both situation and place strongly influencing when and how people used their phones. Barkhuus focused on smartphones as ubiquitous computing devices and how "seamfully"—or not—the devices supported users in their everyday tasks.

Häkkilä and Chatfield's 2005 study of mobile phone privacy focuses on mobile telephone (not smartphone) use generally and SMS messaging specifically.[13] Using both survey and interview techniques, the authors found that their subjects considered mobile phones to be "private and personal devices." They found that privacy is protected by widely held social norms about the confidentiality of phones and messages.

Until recently, most mobile privacy studies focused on location sharing (we do not include cites here due to length), under the assumption that this represented the key privacy issue with smartphones, though a few studies focused primarily on security have also examined users' privacy perspectives on smartphones beyond location issues. Karlson *et al*. in a qualitative study examined users' preferences when sharing mobile phones with other people in order to model users' security preferences.[14] They found that their 12 participants shared their phones with a variety of different "guest" groups, and that concerns with sharing were tied to concerns about access to the information stored on the phones. Mushlukov *et al*. explored users' information protection requirements on smartphones, finding that while users stored "sensitive and valuable data" on their devices, they typically did nothing to secure them from intrusion by other people.[22]

Felt, Egelman, Chin, Wagner and others have conducted a series of studies examining user perceptions, comprehension, and preferences with smartphone APIs. In [7], Felt *et al*. studied user comprehension and preferences with Android application permissions. The study found that only 17 percent of participants paid attention to the permission screens, and of those only 21 percent understood their content. Forty-two percent of participants were completely unaware of permissions, leading the authors to conclude, "current Android permissions warnings do not help most users make good security decisions." Based on these findings, Felt *et al*. proposed a set of guidelines in [8] to help designers in determining more effective permissions granting mechanisms in order to avoid habituation effects and alert users to potential privacy risks.

Next, Felt *et al.* conducted a survey asking smartphone users to rank their level of concern with 99 distinct risks associated with 54 Android API permissions.[9] They constructed a ranking of the risks, finding that users' concerns about a specific permission and its risk were dependent upon the context of usage by an application. They also found that concern about access to location information was low compared to other permission types. Finally, in [6] Chin *et al*. conducted qualitative interviews and surveys with 60 participants ascertaining users' comfort levels with conducting specific tasks on smartphones to test the hypothesis that users elect not to use smartphones for some tasks due to privacy or security concerns. They also explored the reasons why users selected applications for download. Using willingness to perform the same tasks on personal laptops as a basis for comparison, they found that participants were less likely to perform privacy-sensitive tasks, such as financial and health-related transactions, on their smartphones, with over 60 percent citing security-related concerns. Participants also expressed more concern with personal privacy issues on their phones than on their laptops.

Finally, Lin *et al*. specifically examined users' privacy expectations with regards to resource usage by applications.[16] Using a survey deployed on Amazon's Mechanical Turk, they limited their exploration to four types of sensitive resources (device ID, contact list, network and GPS-based location) and asked participants to rate their comfort level with access by specific apps to these resources. They found that participants were the most surprised by access requests they could not explain, and there was a fairly high level of misunderstanding by participants of how applications used resources. Concluding that providing greater rationale behind application requests increased participant comfort, they make specific interface recommendations for communicating the purpose of a request more clearly to users.

4. METHODS

Our study consisted of 24 structured qualitative interviews and a card sorting activity with 11 iOS and 13 Android users representing a range of demographic groups living in the San Francisco Bay Area. While interviews support detailed and nuanced investigation, because of the small sample size and exploratory nature of the inquiry we caution against extrapolating our findings to the broader population of smartphone users without further exploration and testing or concurring research.

The interviews took place over a three-week period in August 2011. We recruited a convenience sample primarily via Craigslist, and respondents took a brief online survey to aid in screening for demographic criteria and to ensure they owned Apple or Android phones and used applications. We randomly selected respondents from pools grouped by gender, age, and phone type. We wanted to interview users of both platforms because although the platforms' APIs provide developers with similar access to information, the manner in which users are made aware of information sharing, and the oversight practiced by the platforms themselves over applications, differ.

Respondents were told they were participating in a general smartphone study and received $40 upon completion. The interviews were approximately one hour in length, were audio recorded and followed a structured interview instrument and later transcribed and coded by theme. They consisted of a series of questions about participants' behaviors, sharing preferences, and experiences with their smartphones. Participants completed card-sorting exercises that allowed them to visually represent the way they thought about mobile applications and information sharing. We interviewed nine men and fifteen women for the study. Our participants ranged in age from 18 to 64, with nearly sixty percent under the age of 34. Education levels varied, with the majority (18) reporting completion of four years of college. The largest ethnic group (15; 63 percent) identified as White, followed by Asian-American (4), African-American (2), Hispanic (2), and mixed race (1). In a 2011 study of smartphone users in the United States, the Pew Internet & American Life Project found that in younger age brackets, higher percentages of Americans own smartphones, with the greatest between ages 18 to 29, and they are owned at greater rates in groups with higher levels of education.[30] Coincidentally, our participant pool mirrored the distribution found in the Pew study.

Interviewees had a range of work and school arrangements, and a subset of our participants worked in technology professions. None specifically worked with mobile applications or had specialized knowledge of the interview topics, though some knew developers who created mobile applications. We were sensitive to the potential bias introduced by asking participants about privacy. To avoid influencing their responses or heightening their concern, we avoided the use of the term "privacy" in our questions,

phrased all descriptions about information access and sharing as factually and neutrally as possible, and placed questions that dealt the most directly with risks and concerns towards the end of the interview.

4.1 Theoretical Privacy Frameworks

Both platforms' notice and consent frameworks reflect the goal of facilitating individual control over the disclosure and management of personal information dominant in U.S. information privacy law. As Section 2.3 notes, the platforms' approaches to privacy protection have not quelled the concerns of users or policymakers. We believe this is at least in part due to their reliance on this failed conceptual framework. We argue that designing technical systems, including mobile systems, that afford privacy requires attention to concepts of privacy that emphasize privacy as an iterative and situated practice rather than a formalistic and procedural mechanism. Thus, we framed our inquiry broadly, seeking to understand how our participants' use of their smartphones fit into the varying context of their daily lives. For our inquiry into privacy expectations we drew upon the theories of Irwin Altman—privacy as *boundary regulation*—and Helen Nissenbaum—privacy as *contextually appropriate information flows,* to situate our line of questioning. We examined privacy expectations from the perspective of both *other people* and *applications* accessing one's phone in order to provide a basis for comparison for understanding more concretely whether users' privacy expectations differed between the two cases (and why). We felt that exploring these two perspectives could yield additional insights into how users conceptualized privacy if they demonstrated substantive differences.

To elaborate on the theories we utilized, boundary regulation, as conceived by psychologist Irwin Altman to explain privacy practices in physical contexts, is "an interpersonal boundary process by which an individual or group regulates interactions with others."[1] Individuals manage their privacy through a dynamic and contextual process of regulating personal boundaries. Leysia Palen and Paul Dourish's framework for understanding privacy as applied to HCI is based upon Altman's work, where they also draw on his theory to grapple with the privacy issues that arise from mediating interpersonal interactions with information technology.[27] We used Altman's theory to structure our exploration of our users' decisions to allow physical access to their smartphones and potentially their personal information by creating questions that focused on participants' formation and organization of boundaries regulating access to their devices. Contextual integrity, focuses on the contextual norms—specifically appropriateness and flow—that govern personal information.[25] We used Nissenbaum's theory to structure our exploration of our participants' level of comfort with the flow of their personal information from their phone to applications, the applications' websites, and to third parties, and whether these flows supported or violated their expectations. Louise Barkhuus recently advocated for the use of contextual integrity by HCI privacy researchers, and specifically studies of mobile systems, in order to better understand the "underlying contextually grounded reasons for people's privacy concern or lack thereof."[3]

5. FINDINGS

We begin with a brief overview of our participants' descriptions of their relationship with and use of their phones, which documents our participants' intense connection to and reliance on this inanimate object. We then offer a detailed account of participants' privacy-related concerns and expectations about their phones and the information stored on them, centering on other people's physical, and application developers' remote, access to them.

5.1 Usage: "You panic if the phone is not with you."

Our participants love their smartphones. As one phrased it, "I am extremely addicted to my iPhone. I use it all the time. I use it for everything, e-mail, directions, apps, reading books, listening to music. It is my life." Smartphones are deeply integrated into their daily lives. Several expressed feeling lost without their phones: "You panic if the phone is not with you." Twenty participants said they carried their phones at all times, and seventeen reported never turning them off: "It's the first thing I do in the morning, it's the last thing I check before I go to bed." Participants used a broad range of applications and enthused about their utility and usefulness. Most reported using their phones for work and/or school in addition to their personal use, blending information and supporting activities from different life domains. The boundary crossing nature of the activities they support offers one explanation for participants' strong attachment to their phones.

Participants found and chose applications through a variety of methods, including recommendations by friends, searching the App Store or Google Play store, through companies' websites, and top application lists from magazines, websites, and blogs. They use crowd-sourced information, such as comments/reviews and ratings, to choose applications. Some picked applications to support a particular need or desire. Participants were generally experienced with deleting applications, with the most common reason for deleting was disuse: "cleaning house." A few mentioned privacy or security concerns as a basis for application removal.

5.1.1 Apps: Website shortcuts? Icons?

We were curious about how people conceived of applications. Technically, they are software programs; some are task specific (*e.g.,* flashlight apps, or games), while others recreate and extend functionality from existing websites (*i.e.*, Yelp!, Google Maps). Because users are often introduced to applications through existing websites, and many websites offer applications that mimic the functionality of their sites, we hypothesized that some users

might confuse an application's functionality with a browser's. This is a crucial distinction, as websites do not have access via the browser to the same resources (either on smartphones or on traditional computers) that applications do through a smartphone's operating system.

To explore their conceptions, we asked our participants how they would describe an application to someone who had never used one before. The results were varied; only about half described an application as a software program, with more Android than iPhone users doing so. One participant explained: "I didn't understand that icons like what we have on our desktop computers and laptops, icons or quick launch, [were] different from an app. The concept of app as a mini program written for a phone didn't sink in until a little later until I actually got the phone and I started playing extensively with downloading apps and looking at comments and seeing, oh, there [are] developers that are creating these apps." Those who understood what applications were realized there was a difference between accessing a website through an application versus a browser, though their descriptions of those differences focused on the limited features and graphical interface of the applications. The substantive difference noted was that an application was optimized for mobile display; no mention was made by any participants of the API as a difference between the two. The other half of our participants either confused applications with websites, describing applications as "icons on my phone" or "shortcuts to a website," or provided answers based on applications' functionality: "It allows you to do all sorts of things;" "It makes life a lot easier."

5.2 Privacy Expectations

We explored participants' expectations of privacy from two dimensions: the risks and concerns associated with *other people's* physical access to the user's smartphone and the personal information within it, and *applications* gaining remote access to personal information through the API. Our theoretical frameworks guided our inquiry; we explored whether, how, and why participants regulated other individuals' access to their smartphone (boundary regulation); and, whether expectations of appropriateness with data types and information flows (contextual integrity) informed participants' decisions about applications. We also examined whether differences between the two platforms, such as the effect of reviewing applications prior to distribution in a platform's store, as well as whether notice and consent mechanisms affected expectations.

We used device sharing as the entry point into our investigation of interpersonal privacy concerns because it is (generally) a voluntary, knowingly undertaken activity that nonetheless makes the sharer's personal information vulnerable to access by another person. It offers a useful comparator case for exploring privacy concerns with the access and collection of information by application developers through the API. For our purposes electively downloading and installing third-party applications is conceptually similar to a user's decision about whether to allow another person to use their device as both activities pose a risk to personal information residing on the device. At the same time, we also hypothesized that the activities may be experientially distinct in ways that matter to the user's risk perception and risk avoidance, as the consequences for exposure in each case can be substantively different.

5.2.1 Smartphone Access by Others: "It's like going into somebody's computer."

Participants considered their smartphones to be highly private as well as deeply personal devices, and they articulated strong opinions about when and who they allowed to access them. "For me, it's a very private, personal thing. There is an intimate relationship with your phone." Another drew upon the fact that smartphones are no longer just phones: "This is not really a phone, it's a computer. So it's like going into somebody's computer."

First, we asked about voluntary sharing: whom did participants trust to share their phones? We found norms governing both *with whom* and *when* to share access. Like in [13], most participants said they only shared their phones with people they knew, and expected that the person with whom they shared would only use it to make a call, look up information, or play a game—not read their emails or text messages, or look for other personal information. Participants shared their phones with family members, friends, or other trusted individuals, though a few noted that they would occasionally share a phone with a stranger or an acquaintance that needed to make a call. Some mentioned they shared their phone rarely (and unwillingly), and others claimed not to share their phone at all, even with those close to them. As one participant described it: "I don't mind if they're not looking at my information. If they're just using it for the browser, [or] if it's just the nearest phone, I don't care at all. But if they're getting on there to look at my emails, obviously that's much different. But I trust my friends and family and husband enough to where I wouldn't be worried about that."

Next, we asked about involuntary access—someone using their phone without their permission. Similar to [13], [14] and [22], nearly every participant explicitly articulated their concern in terms of access to the *information* on their phone, most commonly mentioning text messages, photos, email, and applications with pre-enabled logins, such as webmail or financial accounts—not the use of the phone or charges that might result. We also asked what would concern them more: a stranger accessing their phone, or someone they knew? Overwhelmingly participants were more concerned with strangers: "I'd say it'd be better off in the friend's hands than the stranger's hands, although personally I would keep my phone away from either." Several volunteered that access by a stranger could only occur if their phone were lost or stolen, which gave rise to

specific concerns (*e.g.*, a stranger or thief might access one's bank accounts, email, or social media accounts).

When considering people the participants knew, several expressed qualifications regarding their level of concern, noting that it depended on the context: who the person was and why they would access their phone. For a necessary need (e.g. making an urgent phone call) they may not mind, though others said they would be upset if someone they knew used their phone without permission no matter what the circumstance. Users appeared to rely on these norms to protect the privacy of information on their devices, as only two locked their phones with a passcode. Of the four who reported having had their phones stolen, none could directly trace any deleterious privacy consequences from the theft.

Overall, participants viewed their phones as private and personal closely held devices that functioned much like an extension of their selves, containing highly personal, detailed information about their lives. Accordingly, they relied on shared norms to protect their privacy in the limited contexts in which they shared their devices with friends and family. One participant summarized these boundary regulations concretely: "I have a lot of friends who have phones or other kinds of smartphones so we have a culture around what is permissible use because they're at least aware of the kinds of things you can do with it and the kinds of information they store and you wouldn't want other people to read."

5.2.2 Information Access by Applications

We asked a series of questions to probe participants' understanding and expectations about applications' access to the personal information on their smartphones, exploring their privacy concerns by focusing on trust relationships and their expectations around the access, storage, and sharing of their personal information on their smartphones. We avoided direct questions about the API because they assumed technical knowledge.

5.2.3 Trusting Applications and Platforms

We asked participants whether they trusted any applications more than others. The open-ended construct of the question was chosen to allow participants to interpret the term "trust" as they wished, and we observed three separate interpretations of trust. A minority of participants explicitly tied their perception of an application's trustworthiness to whether it could be relied upon to access and manage their personal information *fairly and respectfully*: "When you engage in a relationship with this phone, with all these applications and what not, there's an understanding here that you're going to respect me and I'm going to respect you back." In contrast, most participants tied trustworthiness to *functional reliability*, anchoring their answers in terms of how reliable they found applications (*e.g.,* if the information an application returned was accurate, or if it was stable and didn't crash the phone). Based on participants' comments, this perception appeared to be at least partially influenced by their prior experiences with the reduction in performance caused by malware on desktop systems, where poorly functioning software is perceived as a sign of a potential system compromise. Additionally, several mentioned *name or brand recognition* as a factor influencing whether or not they found an application to be trustworthy enough to install it. Name recognition was tied to expectations around *assurance structures*: *i.e.*, the "organizational or institutional mechanisms [that] exist to protect individuals from harm." [5] In this instance, the participants' expectation was that a large company wouldn't jeopardize a good reputation or risk a lawsuit by creating something that potentially ran afoul of the law or risked public disapproval.

5.2.4 Perceptions of Security

We asked questions about smartphone information security in order to ascertain what effect practices by the platform maker might have on privacy expectations. First, we asked participants what they thought Google or Apple did to protect the information on their phones, and whether they were ever concerned about the security of their information when they used an application. Two participants expressed a belief that their phone was secure, though based on their responses to other questions they did not appear to have any technological expertise to support their beliefs. Another participant based his opinion in terms of reputation-based effects: "I think Apple will protect their users. . . if they don't protect it, the users will probably leave." Most participants said they had no idea if their phones incorporated any type of security. A few noted they hoped it did, though several said they expected that the platform makers had placed the responsibility on users to make their own smartphones secure.

Some participants told us they were not concerned about the security of their phone information; reasons included that they had little information to lose on the phone, and that an application's popularity provided them *de facto* protection, reasoning that successful application makers would not jeopardize their popularity. "As long as the app's pretty popular . . . if people found out that this app is being corrupt, then everyone's going to uninstall it and it'll just be done, and so they'll lose money." Given the relatively low incidences of malware on phones to date, and the fact that none of our participants reported having had any security-related incidents with their smartphones, information security issues did not appear to be currently exerting much influence on participants, though some discussed security concerns on their smartphones with clear references to negative experiences with desktop computers.

5.2.5 Application Review – Taming The Wild West

We explored whether the platforms' different approaches to curating applications influenced user expectations of privacy. We hypothesized that review of applications by the platform creator would be perceived as an assurance structure, and therefore potentially influence users' perceptions of privacy risks and thus application use. While

Apple's review provides a form of assurance about the applications in their store, it's difficult for users to discern exactly what is being vouched for, since Apple is notoriously vague about the standards it uses to review and approve applications. Regardless, participant awareness of Apple's application review was high; only two of our iPhone users were unaware of it. Strikingly, all but two of our Android users believed that Google also reviewed applications before allowing them in the Google Play store, with several mistakenly attributing Apple's review policies to Google. A few participants mentioned that they would only install applications from the Google Play store based on their (incorrect) assumption that they were reviewed.

Apple's review policy increased iPhone participants' trust in the applications running on their phone. As one put it, "I feel like if it's in the store, then it's fine." A few expressed reservations about the process, focusing on issues of content censorship by the company. When the misinformed Android users were told that Google *did not* review applications prior to their availability at the Google Play store, about half expressed a wish that Google would do so, primarily for security and privacy reasons. "I want them to review because I want them to protect me from unscrupulous data collectors. It's a Wild West data collection issue."

5.2.6 Disclosures – "I always click yes."

Next, we examined the disclosures the platforms make to users. Research on online privacy policies concludes such disclosures are largely ignored.[21] We wanted to know whether smartphone users notice privacy policies and other legal disclosures related to applications, and if so, do they read them? Further, we sought to discover whether Android users notice, read, and understand the permissions presented to them during the installation process. Nearly all the participants recalled seeing a terms of service (TOS) or privacy policy in some form on their phones, but most suffered from notice fatigue: only one claimed to have read one. The majority reported clicking through or ignoring such notices, while a few reported skimming them: "I never read them but I always just click yes." Only one participant reported the contents of a TOS or a privacy policy deterring him from installation of an application.

As discussed above, Apple requires that iOS application developers disclose information collection practices in their TOS or privacy policy (if they have one). Notices are not uniform—in language or presentation—across applications, nor do applications have a uniform process for obtaining user consent. Some applications present a runtime dialog or a checkbox asking users to consent to their TOS, while others do not collect affirmative consent. Prior research found users to be habituated to clicking through such disclosures without reading.[11] Given the abysmal read rate of these documents, reliance on them as the sole means of conveying privacy risks to users to empower them to manage their privacy risk is an unsuccessful strategy, and thus long-term iPhone users were likely highly unaware of the extent to which information access by applications took place prior to iOS6. While nearly all the Android users recalled seeing a permission screen, two participants reported never reading the permissions, two thought permissions were a TOS agreement (and consequently ignored them), and two said they didn't understand what permissions were.

The language used to describe Android permissions was (at the time of the study) quite technical. Only two participants felt they understood what an application was allowed to do after reading permissions, and another four reported a general understanding but stated they didn't understand a few of them. Four participants reported, at least once, not installing an application based on the permissions it requested. Two participants noted that the language needed improvement, and one pointed out that the yes/no nature of permissions presented a non-negotiable choice that left her frustrated: "I like that I'm being asked permission but I also don't feel like it's really that much of a choice because you either accept it or you don't get to use the application." Overall, our findings with respect to permissions complement those in [7].

5.2.7 Impressions of Information Access, Use, and Sharing by Applications

Our participants completed a card sorting exercise structured, following contextual integrity theory, to explore their expectations for how information flows from their phone to applications, application's websites, and third parties with whom application developers might share or sell their information. We used a card sort in order to make tangible the concept of information flow and to ensure that we would have a consistent method for evaluating and comparing preferences across participants. In each instance we chose two applications already installed on participants' phones, selecting one that was account-based (requiring a login) and one that was not. We used participants' own applications to limit any friction caused by confusion about application functionality. We selected login and non-login based applications to explore potential differences between the privacy concerns within relationships with companies that required some amount of personal information (e.g., an email address) or personalization to function, and those that did not.

We showed participants twelve cards (11 for iPhone users), one for each of the following information types: phone number, text messages, location data history, real-time location from GPS, browser cookies, browser history, photos from camera, address book, device (phone) ID, phone call logs, Apple or Google ID, and files on SD card (Android only). We explained that these cards represented different data types that were stored on their phones. We did not include the information collected individually by applications about customer usage (such as individual or aggregated usage statistics). We advised participants that if

there was a data type that they were unfamiliar with to set it aside, though we provided basic definitions if asked. For each of their two applications, we conducted three card sorts (for a total of six card sorts) with each participant:

Sort 1: *Which of the following types of phone data do you think this application needs in order for it to work (function) on your phone?*

Sort 2: *Which of the following types of data would you be comfortable with the application collecting and storing off of your phone and on the developer's website?*

Sort 3: *Which of the following types of data would you be comfortable with the application developer sharing or selling with other companies?*

We asked participants to select the cards with the data types they thought were relevant for each sort. Results were tallied for each exercise. Not all of the cards represent data available to applications through respective phone APIs, and between each platform there are some variations in how applications can access data. For example, we included browser cookies, location history, and Google and Apple IDs even though applications cannot access these data types because we thought participants would be familiar with them and we were interested whether participants would over or under-assume the amount of data their applications could access. We must note that at the time we conducted the study, iOS 5 was not yet available; subsequently, Apple depreciated the use of the device ID (UDID), ostensibly to discourage tracking by ad networks and data aggregators, requiring applications to create their own unique user IDs.[29]

5.2.7.1 Card Sort Results

We analyzed the results by comparing aggregate sorts by application type (login vs. non-login) and by platform type. Given that our participant pool is a small, non-random convenience sample no statistical inferences were drawn. Differences by application type (login vs. non-login) were minor but there were pronounced differences between platforms. iPhone users thought applications needed access to more information, were more comfortable with applications storing data off their phones, and were more comfortable with their applications sharing or selling their data than Android users. Because of the overall similarity between the application types used by both groups and the generally favorable comments iPhone users had regarding Apple's application curation policies and processes, we conjecture that iPhone users' greater familiarity with Apple's store review provides some rationale for these findings. Both iOS and Android users mistakenly believed that applications could access a phone's browser cookies. They cannot, though the Android API does allow access to browser history, while iOS does not. We believe the cookie misunderstanding reflects general confusion about how cookies work [20] and, as noted earlier, the blurry understanding many users have in distinguishing applications from websites. Understanding some users' lack of differentiation between accessing a website via an application, or accessing it via a mobile or desktop browser may offer a crucial insight for disentangling users' mental models about applications' functionality.

We tallied and analyzed the selection frequencies for each data type. We found there was a hierarchy of comfort with information types: overall, participants indicated more comfort with sharing their real-time location, device ID, and location history, and less sharing their photos, address book, call logs, text messages, and files stored on SD cards (Android only). These findings were consistent with comments during the interviews tying comfort with sharing and to perceptions about the level of privacy sensitivity of each of these data types, echoing the findings in [9] and [16]. Participants linked the use of information by an application to its function, demonstrating that context shaped expectations of privacy consistent with the theory of contextual integrity. Several rooted their level of comfort with an application receiving their information in the specific context of the application's function or request, such as a banking application asking for access to one's current location in order to locate the nearest ATM. Again, there were differences between the platforms; across all of the card sorts, for example, iPhone users selected real-time location almost twice as often as Android users. One reason might be the long-term inclusion in iOS of a runtime prompt asking for user approval prior to accessing location has made iPhone users more aware of (and perhaps less sensitive to) location requests.

Participants were more comfortable with applications alone receiving personal information than they were with those applications in turn sharing it with other entities. While participants could understand the rationale for and were generally tolerant of (when it was contextually relevant) an application developer transferring information from their phones to an application's website, there was little tolerance of third-party information sharing—a flow of information that violated their sense of contextual integrity. During interviews, seventeen participants told us they did not want their applications to share or sell *any* information to third parties. One participant described his sense of violation: "I always thought these things were personal devices that we use for ourselves, for our own benefit, but apparently, people have other ideas. I guess we have to share even our own lives with these people." Of the quarter of participants who were comfortable with third party information sharing, location-related information was the type most frequently viewed as legitimate to share. Notably, in all instances where participants selected location information, the application was one where location sharing was relevant to the application's function (*e.g.*, using Yelp for local recommendations), reinforcing participants' reliance on context to inform expectations of information sharing

consistent with contextual integrity. The next most popular information type selected for sharing was device ID.[2]

6. DISCUSSION

The love that our participants felt for their smartphones sets them apart from other computing devices. Their unique combination of usefulness and convenience encouraged many to treat them with the intimacy reserved for a diary or personal journal. And as our findings demonstrate, this close attachment in turn engenders a strong set of privacy expectations surrounding the personal information users store on these devices.

6.1 Contextual Expectations: People vs. Applications

Overall, our participants demonstrated significant privacy expectations with the personal information stored on their smartphones, whether the target of access in question was a person or an application. We discovered that while the *norms* governing these expectations differed between cases, ultimately the *outcome* was similar: participants' comfort level with access was directly related to the context of use in both situations.

While in many cases our participants were comfortable with people they knew accessing their devices, most claimed there would still have to be a justifiable reason for that person to view personal information such as text messages or emails. In the case of strangers, usually no justifiable reason could be conjured. In contrast, applications (and their developers) occupy an interesting middle ground between an entity that users may "know" (*e.g.* through brand familiarity, previous customer experience, or even a personal relationship to the developer) or not (*e.g.* no past history or knowledge of the developer or company). While we did not probe these relationships deeply, it was evident that our participants were willing to allow access by applications when the need was contextually relevant (*i.e.* tied to functionality) and were highly uncomfortable when, similar to the findings in [16], the request appeared to be out of context, or made with the intent of sharing or selling their personal information outside of the original context of use.

The relevance of context isn't isolated to privacy; in [19], the authors claim that no smartphone is an island: "context affects nearly every aspect of next generation mobile phone use, from when participants employ them to what they do with them." Charting the broader landscape of where, when, and why people use their smartphones is key for understanding multiple aspects of smartphone usage, including users' privacy expectations and preferences. Without this backdrop, we may glean that people have preferences and expectations but not necessarily *why* they have them. Lacking this insight, we risk building privacy solutions that fail to grasp the depth of the problem or the true need (as elaborated on next in Section 6.2).

6.2 The Hierarchy of Comfort

The bulk of mobile privacy work to date has focused on location issues, presumably under the assumption that the trackability of smartphones would lead to this issue being at the forefront of users' privacy concerns. Yet, we found that location ranked low in our participants' "hierarchy of comfort" when considering the sensitivity of information types made available to applications by the API, such as text messages, photos, or emails. This finding corroborates complimentary work by Felt [9] and Lin [16].

We argue this finding is important not only because of the specific insight it provides regarding which types of personal information users prioritize on mobile platforms, but also to highlight the need for privacy research more generally that seeks to uncover user expectations from the ground up, using theoretically based methods and definitions of privacy. Certainly location tracking is of concern to many smartphone users, but in our study it was conditional. Participants were uncomfortable with applications requesting access to location if that access wasn't contextually relevant; the access in and of itself was not a locus of concern. One area for future research we did not explore in depth was that of applications requesting and storing a history of a user's location; again, this is an issue where context would likely exert substantial influence. We would expect that most users would be uncomfortable with the use and long-term storage of their location history if any feature of the application didn't justify it.

6.3 Application and Website Confusion

Some of our participants did not distinguish between applications and websites. This was particularly true when using an application created by a website (*e.g.*, Facebook.com and the Facebook mobile application); some thought an application was simply a "shortcut" to a website. Others believed that the only substantive difference between applications and websites was the optimization of the application's user interface for the mobile platform. The author found similar confusion with applications in survey research examining user perceptions of Facebook applications, where approximately 20 percent of survey respondents did not understand that Facebook did not create the applications hosted on Facebook Platform.[15] Even among the participants in this study who understood that mobile applications were not the same as websites, most thought that applications could access their browser's cookies, demonstrating a fuzzy understanding of the differences between browser and application functionality.

Over the years, users have been trained to look for signals that suggest browser-based privacy or information security violations. Subsequently, we found that participants who assessed an application's trustworthiness based on its

[2]When we asked participants what a device ID was, it appeared few understood that it could be used to track their phones by third parties. Most indicated they believed it was a serial number of use only to their service provider.

functional validity referenced experiences with suspicious websites as well as anti-virus software in their explanations, drawing upon the expectation that websites or desktop applications that deliver a low-quality experience or "act strange" signal a lack of credibility or potential malice. While this experience could be useful for evaluating applications that are overtly questionable, the reference point offers no help to users faced with credible, well-functioning applications that nonetheless are engaged in information access and sharing practices enabled by the API yet that violate their expectations.

If, as our findings suggest, the existing default accessibility of information to applications is inconsistent with users' privacy expectations, then ensuring that users understand the difference between interacting with applications, versus visiting web sites, is an essential—if insufficient—step to alert them to the risk that personal information may be automatically disclosed in the background. While the changes introduced by Apple to iOS 6 may help educate users about this difference, on Android there is currently (other than location) no signaling to users that an application is accessing their personal information in real time. This leaves users unable to manage access to personal information in a manner consistent with their privacy needs.

While presenting an endless number of runtime prompts to users may not be a desirable solution, given that users may train themselves to tune them out over time, making application requests for personal information at least more visible and prominent in some form may offer promise. The challenge is to do so creatively and effectively without contributing to notice fatigue by overwhelming or desensitizing users to potential risk. At the same time, visibility without agency is not meaningful; to be effective, users would need to be given choices over the use of their information, such as the ability to deny its use by an application, rather than simply clicking a non-negotiable consent dialog more frequently. Again, future work validating the effectiveness of the iOS 6 privacy settings or similar innovations (such as BlackBerry 10's Trusted Application Status[3]) would be helpful to assess the success of this approach in the wild. More desirably, but less realistic, is the hope that the platform makers might reconsider the level of access they've allowed developers by restricting access in some form to the most sensitive information types.

6.4 Trustworthiness, Assurance Structures, and Disclosures

Most of our participants believed that a variety of assurance structures—e.g., obtaining applications only through official stores, or a developer's positive reputation—protected them from privacy violations by applications. Notably, participants did not or only minimally made reference to an application's TOS, privacy policies, or with Android, the permission interface to evaluate the trustworthiness or privacy risks posed by applications. Though many mobile applications currently lack privacy policies [34], nearly all of our respondents reported that they usually don't read them anyway. Further, as several of our participants pointed out, the non-negotiable, take it or leave it "consent" demanded by these tools is disempowering, especially when it is predicated upon having read and understood the aforementioned unreadable policies. This regime of false choice provides users with few meaningful options. Our participants generally chose applications hoping that either they selected wisely or that their belief in the assurance structures they trusted offered them protection from potential violations.

Users rely upon the design choices and business models of platform providers to make privacy-related decisions—and, rightly or wrongly, to protect their privacy. Although our participants used multiple features of the application marketplace—*e.g.*, reviews, publisher descriptions, and ratings—to make privacy-related decisions, these features were at best proxies and were not optimized for this purpose. But they could be; there is an opportunity for the application stores to offer tools to inform and guide users in making selection choices that support their customers' privacy interests. Recent work by Kelley *et al*. [18] tested a privacy-focused pre-download screen (and an improved permissions screen) on Android devices designed to replace the current interface in the Google Play store and found it had an effect on application decision-making (particularly among those who were interested in privacy). Our research suggests that users expect application stores to police privacy to some extent. Given that Google Play engages in no curation and Apple's curation is not aimed at protecting privacy in a manner consistent with the expectations of our participants, users' reliance on curation to protect their privacy places them at risk. The belief that application stores provide affirmative privacy protection may contribute to the privacy gap, where users' behavior undermines their stated privacy interests.[12] The gap contributes to the self-reinforcing notion that existing designs adequately address users' privacy needs simply because they're being used without recognizing the limited choices users face.

It also wouldn't be difficult for designers to improve mobile notices by standardizing their format, making them visually comprehendible and reader friendly, though this approach may not solve the larger problem of privacy policies remaining unread.[15,28] Better might be to introduce *meaningful choice* to consumers by allowing them to easily and clearly set their preferences for information collection and sharing with applications (illustrated by both the iOS 6 privacy settings and BlackBerry 10's model). However,

[3]The documentation for BlackBerry's Trusted Application Status is available at: http://docs.blackberry.com/en/smartphone_users/deliverables/48900/psm1344876967771.jsp#psm1344877000107.

improving the implementation of the notice and consent paradigm alone is insufficient to address the privacy concerns we document; as we have highlighted, attention to design that meets a broader understanding of privacy expectations and needs is required.

6.5 Theoretical Explanations

In keeping with Altman's theory of boundary regulation, many of our participants viewed their devices as extensions of their selves, regulating access to them accordingly. This was true not only with the types of tasks they engaged in and domains (personal vs. work/school) in which they used their smartphones, but also in the types of information they stored on them. Given the intense personal connections they reported with their smartphones, it is unsurprising that participants controlled and managed access to their devices as much as other researchers have found individuals control and manage access to the self. [24] Some research [14, 22] has proposed improvements to existing platforms in recognition of this boundary blending, such as adding partitions to phones to separate work from personal use (BlackBerry 10 offers a work/personal partition, and Windows Phone 8 includes a "kid's corner" for use by children). However, not all users who merge divergent life roles on their smartphones can or wish to neatly subdivide their lives into multiple contexts.

Rather than arguing for a specific design solution based on these findings, we propose that it highlights the general need for designers to recognize the deep personal connections users may have to both these devices and the information they store on them. We think that they ignore the connections at the peril of undermining user confidence and trust in smartphones and their applications. Arguably, personal computing has become even more personal as focus shifts from the desktop PC, to laptops, and now smartphones and tablets. Ours and others' research suggests that simply assuming that the smartphone is a smaller scale laptop may be misguided. As Matthews *et al.* argues, "be careful transferring user experiences from larger devices." [19] We would expand on that advice to caution transferring assumptions about users' usage patterns and expectations from unrelated platforms, including porting website functionality from desktop browsers to mobile applications. It may be simpler to assume that mobile is just a scaled-down version of the desktop/laptop computing experience, but as we and others have uncovered, it appears instead to be personal computing at its most personal to date, with its own vagaries.

Guided by Nissenbaum's theory of contextual integrity, we found that existing information flows enabled by platform APIs also violated participants' privacy expectations, in particular when information moves in ways that disrupt the established context. As discussed earlier, this finding is confirmed by [16], who also observed confusion or discomfort by their participants when applications made information requests that violated the context of use. Our application of contextual integrity also supports the argument made by Barkhuus that the theory offers "a more nuanced treatment of privacy" in HCI and ubicomp research generally, but with respect to mobile research specifically, through the investigation of an "actual sharing situation." [3]

Our results demonstrate that users are open to granting contextually relevant access requests when the benefits are clear and circumspect when not. In particular, the proposition of sharing personal information with third parties when no contextual justification exists was non-negotiable for the majority of our participants. Doubtless there are multiple approaches designers and researchers can take to improve current models of access requests and to suggest new models; our goal here again is not to propose specific design solutions but again to drive the broader point home that no matter what specific design solution is tried, incorporating a contextual inquiry into the process may produce a result that better reflects users' privacy expectations.

7. CONCLUSION

The default flows of smartphone APIs defy users' privacy expectations. Our participants held consistent expectations of privacy in the personal information on their smartphones whether they were concerned with other people or applications accessing it. They expected that the use of their personal information by applications would hew closely to the minimum required for an application's functionality. Most strongly objected to the transfer of information to third parties. These findings suggest that platform developers should restrict API defaults and application developers should design explicitly for privacy in ways that better align flows of personal information to users' expectations.

Further, user confusion about how applications function, as well as false or unsupported beliefs in assurance structures contribute to an environment where users act in ways that may ultimately belie their privacy interests. The behavior our participants reported demonstrated that many were concerned about privacy issues but were forced to navigate an ecosystem that was not supporting their interests. Incorporating users' privacy concerns across the mobile ecosystem would reduce the need for policing and curation by platform providers, reduce the burden on users who don't want to trade privacy for functionality, and ease policymaker and public concern about smartphone risk.

Finally, the overemphasis on location-related concerns by mobile privacy researchers suggests that incorporating (or supporting) theoretically and contextually grounded research into users' privacy concerns and expectations may yield findings that are better aligned with users' needs. The two theoretical approaches we used here, boundary regulation and contextual integrity, provided a useful framework for structuring our research and yielded findings we argue better match actual usage.

8. ACKNOWLEDGEMENTS

Emily Barbaras and Maydha Basho provided invaluable interviewing and research assistance and contributed to earlier drafts of this paper. Thanks to Deirdre Mulligan, Coye Cheshire, Liz Goodman, the I-School Ph.D seminar, and my workshop participants at the 2012 Privacy Law Scholars Conference for feedback and support. This research was supported by the U.S. Department of Homeland Security, under grant award #2006-CS-001-000001, and the National Institute of Standards and Technology, under grant award #60NANB1D0127, under the auspices of the Institute for Information Infrastructure Protection (I3P) research program.

9. REFERENCES

1. Altman, I. The environment and social behavior. Brooks/Cole: Monterey, CA, 1975.
2. Arthur, Charles. "Is your private phone number on Facebook? Probably. And so are your friends'." *The Guardian*, Oct. 6, 2011. http://www.guardian.co.uk/technology/blog/2010/oct/06/facebook-privacy-phone-numbers-upload
3. Louise Barkhuus. The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI. In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems* (CHI '12). ACM, New York, NY, USA, 367-376.
4. Barkhuus, L. and Polichar, V. E. Empowerment Through Seamfulness: Smartphones in Everyday Life. Personal and Ubiquitous Computing. August 2011, Volume 15, Issue 6, pp 629-639.
5. Cheshire, Coye (2011). "Online Trust, Trustworthiness, or Assurance?" Daedalus. Vol. 140, Issue 4: 49-58.
6. E. Chin, A. P. Felt, V. Sekar, and D. Wagner. Measuring User Confidence in Smartphone Security and Privacy. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS '12)*.
7. A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner, "Android Permissions: User Attention, Comprehension, and Behavior," In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS '12)*.
8. A. P. Felt, S. Egelman, M. Finifter, D. Akhawe, and D. Wagner. How To Ask For Permission. USENIX Workshop on Hot Topics in Security (HotSec) 2012.
9. A. P. Felt, S. Egelman, and D. Wagner. I've Got 99 Problems, But Vibration Ain't One: A Survey of Smartphone Users' Concerns. CCS Workshop on Security and Privacy in Smartphones and Mobile Devices, 2012.
10. Golson, Jordan. Apple Requires User Permission Before Apps Can Access Personal Data in iOS 6. MacRumors.com, June 14, 2012. http://www.macrumors.com/2012/06/14/apple-requires-user-permission-before-apps-can-access-personal-data-in-ios-6/
11. N. S. Good, J. Grossklags, D. K. Mulligan, and J. A. Konstan. 2007. Noticing notice: a large-scale experiment on the timing of software license agreements. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (CHI '07). ACM, New York, NY, USA, 607-616.
12. J. Grossklags, and A. Acquisti. When 25 Cents is Too Much: An Experiment on Willingness to Sell and Willingness to Protect Personal Information. In *Proceedings of Sixth Workshop on the Economics of Information Security (WEIS 2007)*, Pittsburgh, PA, June 7- 8, 2007.
13. Häkkilä, Jonna and Chatfield, Craig. 'It's Like if you Opened Someone Else's Letter'–User Perceived Privacy and Social Practices with SMS Communication. In *Proceedings of the 7th International Conference on Human Computer Interaction With Mobile Devices & Services* (MobileHCI '05). ACM, New York, NY, USA, 219-222.
14. A. K. Karlson, A.J. Bernheim Brush, and S. Schechter. 2009. Can I Borrow Your Phone?: Understanding Concerns When Sharing Mobile Phones. In *Proceedings of the 27th International Conference on Human factors in Computing Systems* (CHI '09). ACM, New York, NY, USA, 1647-1650.
15. Jennifer King, Airi Lampinen, and Alex Smolen. 2011. Privacy: Is There An App For That?. In *Proceedings of the Seventh Symposium on Usable Privacy and Security* (SOUPS '11). ACM, New York, NY, USA, Article 12, 20 pages. DOI=10.1145/2078827.2078843
16. Jialiu Lin, Norman Sadeh, Shahriyar Amini, Janne Lindqvist, Jason I. Hong, and Joy Zhang. 2012. Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy Through Crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (UbiComp '12). ACM, New York, NY, USA, 501-510. DOI=10.1145/2370216.2370290
17. P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems* (CHI '10), ACM, New York, NY, USA, 1573-1582.
18. Patrick Gage Kelley, Lorrie Cranor, and Norman Sadeh. Privacy as Part of the App Decision-Making Process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13). ACM, New York, NY, USA (advance copy).
19. T. Matthews, J. Pierce, and J. Tang. No Smartphone is an Island: The Impact of Places, Situations, and Other

Devices on Smartphone Use. Research Report RJ10452, IBM, Sept. 2009.

20. McDonald, A. M. Cookie Confusion: Do Browser Interfaces Undermine Understanding? In *Proceedings of the 28th International Conference Extended Abstracts on Human Factors in Computing Systems* (2010). CHI EA '10.
21. Aleecia M. McDonald, Robert W. Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor. 2009. A Comparative Study of Online Privacy Policies and Formats. In *Proceedings of the 9th International Symposium on Privacy Enhancing Technologies* (PETS '09), Ian Goldberg and Mikhail J. Atallah (Eds.). Springer-Verlag, Berlin, Heidelberg, 37-55.
22. I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov. Understanding Users' Requirements for Data Protection in Smartphones. In Workshop on Secure Data Management on Smartphones and Mobiles, 2012.
23. NielsenWire. "America's New Mobile Majority: a Look at Smartphone Owners in the U.S." May 9, 2012. http://blog.nielsen.com/nielsenwire/online_mobile/who-owns-smartphones-in-the-us/
24. Nippert-Eng, Christena. *Islands of Privacy.* Chicago, IL: The University of Chicago Press, 2010.
25. Nissenbaum, Helen. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Palo Alto, CA: Stanford University Press, 2010.
26. Paczkowski, John. "Apple: App Access to Contact Data Will Require Explicit User Permission." AllThingsD, Feb. 15, 2012. http://allthingsd.com/20120215/apple-app-access-to-contact-data-will-require-explicit-user-permission/
27. L. Palen and P. Dourish. 2003. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '03). ACM, New York, NY, USA, 129-136.
28. PrivacyChoice PolicyMaker, available at: http://www.privacychoice.org/policymaker.
29. Schonfeld, Erick. "Apple Sneaks A Big Change Into iOS 5: Phasing Out Developer Access To The UDID" Techcrunch.com, August 19, 2011. http://techcrunch.com/2011/08/19/apple-ios-5-phasing-out-udid/
30. Smith, Aaron. Smartphone Adoption and Usage. Pew Internet & American Life Project. July 11, 2011. http://pewinternet.org/Reports/2011/Smartphones.aspx
31. Thampi, Arun. "Path uploads your entire iPhone address book to its servers." Feb. 8, 2012. http://mclov.in/2012/02/08/path-uploads-your-entire-address-book-to-their-servers.html
32. Thurm, S. and Kane, Y. I. "Your Apps Are Watching You." *The Wall Street Journal*, Dec. 17, 2010. http://online.wsj.com/article/SB10001424052748704694004576020083703574602.html
33. United States of America *v.* Path, Inc. Consent Decree and Order for Civil Penalties, Permanent Injunction, and Other Relief. Case No. C 13 0448. February 1, 2013. Available at: http://www.ftc.gov/os/caselist/1223158/130201pathincdo.pdf
34. Vora, S. and Dakin, S. FPF Finds Nearly Three-Quarters of Most Downloaded Mobile Apps Lack A Privacy Policy. Future of Privacy Forum, Dec. 2, 2011. http://www.futureofprivacy.org/2011/05/12/fpf-finds-nearly-three-quarters-of-most-downloaded-mobile-apps-lack-a-privacy-policy