

Privacy as Part of the App Decision-Making Process

Patrick Gage Kelley
University of New Mexico
pgk@cs.unm.edu

Lorrie Faith Cranor
Carnegie Mellon University
lorrie@cs.cmu.edu

Norman Sadeh
Carnegie Mellon University
sadeh@cs.cmu.edu

ABSTRACT

Smartphones have unprecedented access to sensitive personal information. While users report having privacy concerns, they may not actively consider privacy while downloading apps from smartphone application marketplaces. Currently, Android users have only the Android permissions display, which appears after they have selected an app to download, to help them understand how applications access their information. We investigate how permissions and privacy could play a more active role in app-selection decisions. We designed a short “Privacy Facts” display, which we tested in a 20-participant lab study and a 366-participant online experiment. We found that by bringing privacy information to the user when they were making the decision and by presenting it in a clearer fashion, we could assist users in choosing applications that request fewer permissions.

Author Keywords

Privacy; Android; Mobile; Interface; Decision-making

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces

INTRODUCTION

In the past five years Android and iOS, the two now-largest smartphone operating systems, have transformed phones from devices with which to call others into true pocket computers. This has largely been accomplished through smartphone applications, often small, task-focused, executables that users can install on their phones from software markets. However, with each application a user downloads they may be sharing new types of information with additional app developers and third parties. Easy access to hundreds of thousands of applications from a diverse and global set of developers and the large amount of personal and sensitive data stored on smartphones multiply the privacy risks.

In Google Play, the current Android application marketplace, users are shown a series of “permissions” only after they have elected to download an application. Previous research suggests that users are likely to ignore the permissions display because it appears after they have decided to download a particular app [4, 12]. Furthermore, even users who pay attention

to permissions displays have trouble using them because the screens are jargon-filled, provide confusing explanations, and lack explanations for why the data is collected.

Our research aims to provide an alternative permissions and privacy display that would better serve users. Specifically, we address the following research question: Can we affect users’ selection decisions by adding permissions/privacy information to the main app screen?

To answer this question, we created a simplified privacy checklist that fits on the main application display screen. We then tested it in two studies: a 20-participant laboratory exercise and a 366-participant Mechanical Turk study. In each study we asked our participants to role-play selecting applications for a friend who has just gotten their first Android phone. Participants were assigned to use either our new privacy checklist or the current permissions display found in the Android market. Our results suggest that our privacy checklist display does affect users’ app selection decisions, especially when they are choosing between otherwise similar apps. We also found that both the timing of the privacy information display and the content of the display may impact the extent to which users pay attention to the information.

RELATED WORK

We outline previous research on the security model of the Android operating system, the current permissions model, and users’ expectations regarding their phones. We focus on Android due to its historically more detailed permissions system and its large user base.

Android as a Major Application Provider

As of May 2012, Android has had over 15 billion application downloads, and over 500,000 applications, with both these numbers continuing to grow at an increasing rate [19].

Applications are not pre-screened for quality. Android app rating and recommendation site AppBrain reports that 33% of the applications in the Android Market are rated “low quality” by users. Additionally, a 2011 Juniper Networks report found “a 472% increase in Android malware samples” between July and November 2011 [11]. Similar studies from McAfee [16], Kaspersky Lab [20], and Symantec are all reporting continued exploits. The types and quality of this malware vary widely, ranging from attacks that collect user data (normally IMEI and other identifiers), to attacks that delete user data or send premium SMS messages.

To combat malicious applications Google internally developed a malware blocking tool codenamed Bouncer. Google announced that Bouncer had been checking “for malicious apps in Market for a while now,” and as a result malware was

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI’13, April 27–May 2, 2013, Paris, France.

Copyright 2013 ACM 978-1-4503-1899-0/13/04...\$15.00.

declining [18]. However, there are reports of Bouncer’s limitations, such as applications existing in the market for weeks without being noticed [21].

Android Security Research

While Android has only existed publicly since 2008, a significant amount of work has been conducted on studying the Android permissions/security model. Much of this work focuses on creating theoretical formalizations of how Android security works or presents improvements to system security, and is largely out of scope. Enck et al.’s TaintDroid has bridged the gap between system security and user-facing permissions, focusing on analyzing which applications are requesting information through permissions and then sending that data off phone [5].

Vidas et al. also studied how applications request permissions, finding prevalent “permissions creep,” due to “existing developer APIs [which] make it difficult for developers to align their permission requests with application functionality” [25]. Felt et al., in their Android Permissions Demystified paper, attempt to further explain permissions to developers [6]. However, neither of these papers explore end-users understanding of permissions.

There is also a growing body of work on the complexity of the current permissions schemes users must deal with. Researchers have discovered novel attack vectors for applications to make permission requests that are not reported to users [3]. Others who have looked at Android permissions have attempted to cluster applications that require similar permissions to simplify the current scheme [2] or have attempted a comparison of the differences between modern smartphone permission systems [1].

Android Permissions and Privacy Research

Android permissions are a system controlled by the Android OS to allow applications to request access to system functionality through an XML manifest. As these permissions are shown to the user at install time, this system as a whole forms a Computer-Supported Access Control (CSAC) system, as defined by Stevens and Wulf [24].

The majority of work done on user expectations related to this Android access control system has been done by our own group at Carnegie Mellon [12, 17] and two separate teams at Berkeley.

Felt and her colleagues have published a series of papers on the Android permission model, and how users understand it. They found that most users do not pay attention to the permissions screens at install time (83%) and that only three percent of their surveyed users had a good understanding of what the permissions were actually asking for access to [9]. They also performed a large risk-assessment survey of users’ attitudes towards possible security and privacy risks, and possible consequences of permission abuses [8]. These results influenced our selection of items to include in a privacy checklist. Felt also performed work detailing other possible methods for asking for permission, with a set of guidelines for presenting these privacy and security decisions to users [7].

Moving away from permissions, the work of King et al. has explored user expectations across the entire use of their smartphones. This broader work, which included interviews with both iPhone and Android users, highlighted difficulties in recognizing the difference between applications and websites, personal risk assessments of possible privacy faults, and how users select applications in the application marketplaces [14].

Research in privacy policies, financial privacy notices, and access control have all similarly shown that privacy-related concepts and terms are often not well understood by users expected to make privacy decisions [13, 15, 22]. No work we are currently aware of has proposed and tested alternative permissions displays, or other ways to help users select applications in Google Play, or other application markets, as we do here.

PRIVACY INFORMATION IN THE ANDROID MARKET

This section details how Google Play currently presents privacy information and other information to consumers to help them select new applications to download to their Android smartphone. We then discuss the privacy facts display we designed to make privacy- and security-related information more central to users’ selections.

Privacy currently in Google Play

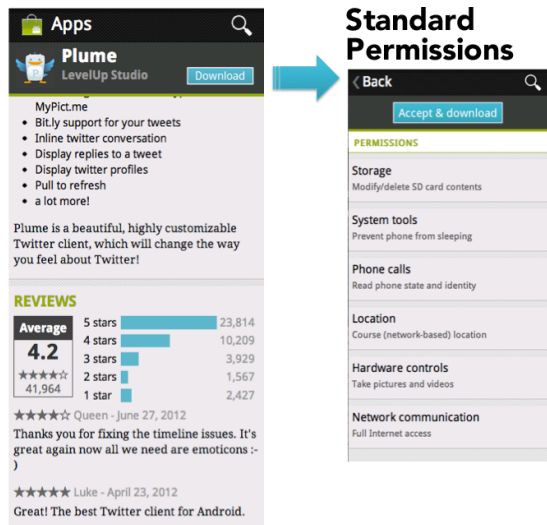
Google Play users are presented with a number of ways to search and browse for new applications. Featured applications, top charts, categories, a search tool, and similar application lists each direct users to a common “Application Display Screen” (Figure 1A. Standard Market).

This screen provides users with a long list of information about each application. This includes (but is not limited to), a series of navigational items, application information, screenshots, a series of market-assigned labels (top developer, editor’s choice), free-text descriptions, a series of reviews, and a series of other types of applications that users may have viewed or chosen. The current market application display screen is very long, yet completely lacks privacy information.

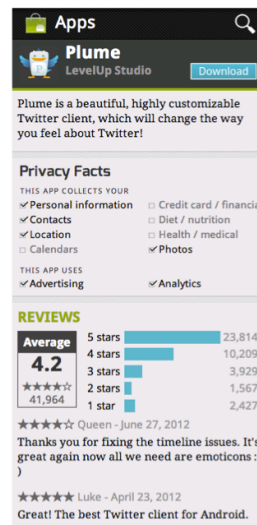
Privacy/security information appears on the above screens only when it is mentioned in free-form text by developers or when it appears in text reviews (almost always in a negative context). Market-provided (and by extension, system-verified) privacy/security information appears only on the secondary screen shown after a user has clicked the download button.

This secondary screen, where permissions are displayed (Figure 1, A. Standard Permissions), again displays the application name, icon, developer, and top developer status icon. This is followed by a very large accept button, which is followed (thus after the action target) by a list of grouped permissions. Only some permissions are shown initially, followed by a “See all” toggle that expands to display the remainder of the permissions an application requests. Each of these permission groups can be selected to see a pop-up window that contains the definitions for each of the permissions in the selected group. Because there may be several grouped

A Standard Market



B Privacy Facts



C Permissions Inline

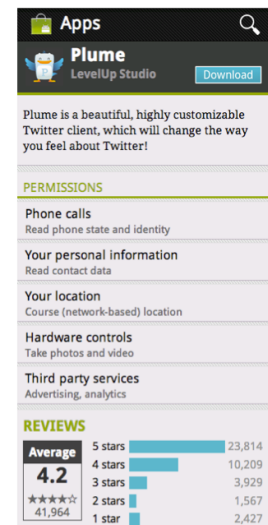


Figure 1. The three privacy/permissions display conditions we tested in our experiments.

permissions, the pop-ups may have to be scrolled to be read completely.

Reasons for modifying the Android application display

We posit that by the time a user selects to move forward by tapping the Download button, they have already made their purchase decision. We will see that this is true within our interview study below. For privacy information to be a salient part of the decision process, it must be presented to the user earlier in the process. Privacy information could be included in the long list of other application aspects on the standard application screen. Instead the current market places permissions on a secondary screen. While some might argue that placing permissions on their own screen draws users' attention to them, our results suggest that it actually does a disservice to users because they are unable to consider permissions as they consider other app characteristics.

Prototype privacy facts checklist design

We created a series of several possible locations and distinct styles of display in an ideation round. The custom privacy display that we decided to test is the Privacy Facts Checklist display shown in situ (Figure 1B. Privacy Facts). The display has several features:

Information— The display has two areas of information. The first with the header “THIS APP COLLECTS YOUR,” describes eight types of information the app may collect: Personal information, contacts, location, calendars, credit card/financial, diet/nutrition, health/medical, and photos. The second header specifies “THIS APP USES” and lists advertising and analytics. Each of these ten items has a checkbox next to it, indicating use.

Display Style— The display is 270 pixels tall and the full width of the device (matching other standard application display sections). For comparison, the rating histogram is 162 pixels tall and the screenshots are the same as our privacy

display at 270 pixels.¹ The display has a bold header “Privacy Facts” in a non-Android-standard type.² The remainder of the display is presented in the standard Android Market typeface. The items are each displayed at the standard size, with the headers in capital text in a lighter font color.

Location— The display is shown immediately after the Description section (and Video and What's New sections, if present, which they were not in our studies) and always immediately before the Reviews section. This means when participants first see each app screen there is no visual difference from the market as it is currently displayed, as the Privacy Facts section appears below the fold (as it would on most phone models).

Permission mapping— For this display we strayed from the current Android permissions by:

- Including types of information being collected that fall outside of the scope of the current permission model (health information, other financial information).
- Including the use of third-party modules, specifically advertising and analytics.
- Removing permissions that are nearly always used (Internet) and those that are irrelevant to most users such as networking protocols and rarely used permissions.
- Including photographs, which are currently accessible to applications.

The final selection of the checklist items we used was strongly influenced by the work of Felt et al. [8] as well as our earlier work [12], and a series of online pilots. The checklist includes

¹There is variation in screenshot size on different Android phone models. The measurements above, and throughout this paper, are from a Google Nexus One.

²The font used is Exo from the Google Font Library.

both Android permissions as well as user-provided information. We wanted this display to include both, for a more holistic privacy summary. Also, by including an item like photos, we create a display that is more in line with users' expectations (which universal accessibility of photos is not). A more complex form of this display could include information that explains how these permissions are used, what they are used for, or how frequently they are used.

METHODOLOGY

We will discuss two phases of experiments: a 20-participant laboratory exercise and interview study, and an online 366-participant MTurk app comparison survey.

In our studies we ask participants to actively consider how and why they download applications in the market, complete our application selection task, and then discuss that experience. In both studies, the core of the experiment was an application selection task using different market designs that vary in how privacy information is presented.

Our study design was based on a similar study run by a team of researchers at Berkeley. The researchers had participants decide whether to install applications on a computer to see whether people read license agreements at install time. Their users evaluated the software tools as complete packages, based on brand, design, functionality, and also End User License Agreements [10]. Similarly, we seek to understand whether people read the permissions display or our updated privacy facts display when installing software on an Android smartphone, and whether we can manipulate their decisions through improved design and information.

Application selection task

The main task asked participants to select one application from each of six pairs of applications we presented in our "custom Android market." We presented two applications for each of the six categories (below). All of the applications we used were real applications that could be found and downloaded in the market. Their names, screenshots, descriptions, features, ratings, and reviews were all authentic. However, we picked most applications in the 1,000 to 10,000 download range, such that the applications would not have been seen or used by most participants. We displayed three text reviews per application, one 2- or 3-star, one 4-star, and one 5-star review.

In four of the comparisons we tested applications that were roughly equivalent (Twitter, document scanning, word game, nutrition app). In each of these four cases participants were presented with two applications with different permissions requests, detailed in Table 2. In each of these choices one of the applications requested less access to permissions and personal information (low-requesting v. high-requesting).

We also tested two special-case comparisons, to begin to explore the effects of rating and brand. In the flight-tracking comparison, we modified one of the applications (Flight-Tracker, low-requesting), to have an average rating of 3-stars. All of the other applications in all categories had 4-star average ratings. In the case of streaming music apps, we tested Spotify, a highly-known (shown in pre-tests) application with

over 50 million downloads. Nearly all of our participants recognized this application.

Lab Study

To test the privacy facts display, and explore our research question, we conducted a series of semi-structured laboratory exercises in July 2012 with 20 participants. This was a between-subjects design. For the main application selection task ten participants saw the privacy facts checklist, and the other ten saw the current Android permissions display. We performed exploratory follow-up interviews seeking broad understanding of participants' interactions with their smartphones as well as diving deeply into issues surrounding the display of permissions, understanding of the terms in the checklist/permissions display, the safety of Google Play, and possible harms of information sharing.

We recruited participants through flyers and local Craigslist postings. Each candidate filled out a short pre-survey online before the exercise, which allowed us to confirm they used an Android-enabled smartphone. We performed the study in an on-campus lab and audio recorded the interviews. Participants were assigned randomly to conditions (without any balancing for gender, time-using android, technical knowledge, or age). They were paid \$20 for successful completion of the interview, in the form of their choice of Target, Starbucks, or Barnes & Noble gift cards.

Exercise and Interview focus

The lab study followed a semi-structured format, outlined here:

- *Android introduction:* Questioned participants about general Android experience
- *General new smartphone advice:* Asked for advice to give to a hypothetical friend and new smartphone owner
- *Specific new smartphone advice:* Requested advice framed around a desire for six specific types of apps
- *Application selection task:* Had participants select applications with a Google Nexus One smartphone on our modified market
- *Post task explanation:* Requested explanations for why each app was selected
- *Android in the news and malicious activity:* Inquired on awareness of Android and apps in the news or on the Internet, then on malicious apps
- *Android permissions and privacy displays:* Drilled down to the privacy and permissions issues, asking if they had noticed the new display or used the current permissions display, depending on condition

Online Study

We conducted an online survey, a 366-participant MTurk test of the same application selection task used in the laboratory study. Because this was performed on MTurk the application selection task had a more structured survey format, as well as some other methodological differences that will be discussed

	Gender	Age	Occupation	Phone Model	Time Using Android	# of Apps Downloaded	# of Apps Frequently Used
P1	Female	21	Student	Motorola Droid	1–2 years	11–25	1–5
P2	Male	21	Student	Motorola Photon	7 months–1 year	101+	6–20
P3	Female	29	Other	Motorola Droid	1–2 years	11–25	6–20
P4	Female	39	Non-Profit	T-Mobile MyTouch 4G Slide	More than 2 years	11–25	20+
P5	Female	44	Marketing	Pantech Breakout Droid	7 months–1 year	11–25	6–20
P6	Female	30	Research / Science	Motorola Droid	1–2 years	11–25	6–20
P7	Male	43	Other	Motorola Droid	1–6 months	1–10	None
P8	Male	20	Student	Motorola Defy	1–2 years	1–10	6–20
P9	Male	31	Healthcare / Medical	Motorola Droid	More than 2 years	1–10	1–5
P10	Female	23	Research / Science	Samsung Galaxy	1–2 years	11–25	1–5
A1	Female	20	Student	Motorola Droid	7 months–1 year	11–25	1–5
A2	Female	23	Don't work	T-Mobile G2/HTC Desire Z	More than 2 years	1–10	1–5
A3	Female	20	Student	LG Ally / Optimus	1–2 years	26–100	6–20
A4	Female	28	Student	Samsung Galaxy	More than 2 years	11–25	6–20
A5	Female	24	Student	HTC rezound	More than 2 years	11–25	1–5
A6	Female	24	Research / Science	LG Ally / Optimus	1–2 years	11–25	1–5
A7	Male	21	Student	Motorola Droid	1–2 years	26–100	6–20
A8	Female	23	Research / Science	T-Mobile HTC G2	1–2 years	11–25	1–5
A9	Female	26	Research / Science	HTC Status	1–2 years	26–100	6–20
A10	Female	44	Healthcare / Medical	Motorola Droid	1–2 years	26–100	6–20

Table 1. Basic demographics of our lab study participants. Participant numbers beginning with P saw the privacy facts checklist, those with A saw the standard Android system. All the information above is self reported.

below in the limitations section. We again used a between-subjects design, but with three conditions. Participants saw one of: the privacy facts checklist (Figure 1B); the current android permissions display (Figure 1A); or the current android permissions display style and terms, presented in the application display screen with additional terms to cover categories from the privacy facts display (Figure 1C). In each case they were asked to pick six from the same 12 applications that our participants in the lab study were given, and then were asked to write a short sentence explaining their choice. For successful completion of the survey turkers were paid \$0.30.

We used MTurk’s user filtering system (95% success required) and required English speakers and Android users. The survey was front loaded with questions about the turker’s Android device to discourage users who did not use Android phones. We manually inspected free-response questions to check for participants who were answering randomly, but removed no participants in that stage, only filtering (12) users who had not used the Android market.

LAB STUDY RESULTS

In this section we detail the results from our lab study. We cover the basic demographics of our participants, their experience with Android, their advice both general and specific to their hypothetical friend, the results of their application selection, and their post-task interview responses.

Demographics

As shown in Table 1, 25% of our 20 participants were male and 75% were female. Participants were between 20 and 44 years old, with an average of 28; 30% were undergraduates. All of our participants had downloaded Android applications from the market and were neutral or satisfied with the Google Play experience.

Application selection

The Privacy Facts display appears to have influenced participants in two of the four standard comparisons and in both of the special comparisons. Full selection percentages can be found in the first two columns of Table 3 (alongside the online study results).

In two of the four standard comparisons (word game, and Twitter) participants who saw the privacy facts display were, on average, more likely to pick the application that requested fewer permissions. In Document scanning, only one participant in each condition did not pick DroidScan Lite (the low-requesting app). In the diet application choice, no participants in the Android condition picked Doc’s Diet Diary (the high-requesting app), while three with the Privacy Facts display did. In both the two special comparisons more of the participants who saw the privacy facts display picked the low-requesting app.³

Participants placed substantial weight on the design and perceived simplicity of using the application. Participants continued to surprise us with ever more idiosyncratic reasons for selecting certain applications. One participant preferred applications with simplistic names, saying “I like to download the apps that have a name that I can easily find. So Calorie Counter, I know where that is gonna be on my phone. I don’t have to be like, oh, what is this called.”

Participants reported wanting to try the apps out, often saying they would download many and see which was the best (which our study prevented them from doing). One said “And I might try things out and see... I just kind of see how well it works, because some things are more glitchy.”

³Given the small numbers of participants we did not expect differences to be statistically significant and only the Twitter application choice was significant (Fisher’s Exact test, $p = 0.023$, the odds ratio is 11.64).

	Personal	Contacts	Location	Calendars	Financial	Diet/nutrition	Health/medical	Photos	Advertising	Analytics	Total
Wordoid!	-	-	-	-	-	-	-	-	-	-	0
Word Weasel	✓	-	✓	-	-	-	-	-	-	✓	3
Twidroid	✓	-	-	-	-	-	-	✓	-	-	2
Plume	✓	✓	✓	-	-	-	-	✓	✓	✓	6
DroidScan Lite	-	-	-	-	-	-	-	✓	-	-	1
M. Doc Scan Lite	✓	✓	-	-	-	-	-	✓	-	✓	4
Calorie Counter	✓	-	-	-	-	-	-	-	-	✓	2
Doc's Diet Diary	✓	✓	✓	-	-	✓	-	✓	-	✓	6
Rdio	✓	✓	-	-	-	-	-	-	-	-	2
Spotify (<i>brand</i>)	✓	✓	✓	-	✓	-	-	✓	✓	✓	7
Flight Tracker	✓	-	✓	-	-	-	-	-	-	-	2
iFlights (<i>rating</i>)	✓	-	✓	✓	✓	-	-	-	-	✓	5

Table 2. The boxes checked in the privacy facts checklist for each application are shown above. In each application category, one of the two applications requested access to fewer permissions (low-requesting always shown first).

Possible hidden costs also impacted application selection. Several participants noted that while the music streaming applications were free (as were all the applications we tested), they might have to purchase a subscription, or be unable to access certain functionality after a trial period ended. Participants generally wanted to avoid applications where features would expire or that would require later costs, but more importantly they expected the details of these arrangements to be extremely clear in the descriptions.

Android in the news and malicious activity

Most participants reported not seeing much about Android in the news, and most of what they did see being comparisons between Apple's iOS and Android. When we asked about reports of malicious apps, or apps doing unintended things, participants said they had not heard about this. Many believed that it could be hypothetically possible. One participant said "Like, I have wondered, oh could an app be a virus," another "I've heard about viruses, that they can actually shut your computer or phone down. Spyware."

Permissions and Privacy terminology

To test whether the terms we selected for the Privacy Facts display were understandable, we asked participants to explain what each term meant. While most were very clear, Personal Information and Analytics were the two that participants had the most trouble with. Personal Information answers were often too broad, encompassing things we did not intend. For example, one participant defined it as "That would mean like... interactions within the phone, Gmail, Messaging, Calling different people."

Participants generally preferred the checklist and its terminology. One participant said, "[Privacy Facts is] very straightforward to me. And that is something I noticed, I was thinking, Oh this is cool, is this what they are doing now. That is why I didn't say anything about it. I can immediately go: No, Yes, No, Yes."

Only two participants explicitly mentioned privacy information in their application selection decisions, both in the privacy facts checklist condition. One participant, said, "If this one is offering the same thing and they want less of your information, I would go with the one that wants less of your information." This comment shows her awareness of the privacy information, but also that the functionality must be matched between apps.

Task time and permission views

Overall, the entire laboratory exercise ranged from 29 minutes to 59 minutes (average 39:53). Participants spent between 3 minutes and 47 seconds to 25 minutes and 6 seconds on the application selection task. There was no statistically significant difference between conditions (two-tailed t-test, $p = 0.726$), although participants who saw the privacy facts checklist took on average 50 seconds more (11:40 v. 10:51) to complete the task.

Across all participants in the Android permissions condition, the permissions screen was used by participants for about half the selection decisions. Four participants decided which applications they would select without ever looking at any permissions screens. Another four participants looked at permissions for all the applications they selected. A6 looked at both Twitter applications permissions, but did not look at the permissions for either of the flight applications. A9 looked at only the permissions for the Twitter application she selected and no other applications.

Across all 31 permission screen views, participants spent between 1 and 11 seconds looking at the Android permissions display. On average they viewed the permissions display for 3.19 seconds (median 2 seconds), including page load time, a minuscule amount compared to time spent on the applications display screen.

ONLINE STUDY RESULTS

In our online study, the application selection task was conducted on MTurk through a participant's computer, not a smartphone. Participants saw the applications presented at smartphone size, side-by-side in iframes. Participants selected the application they thought was better for their friend, provided a short text reason, and then rated each of the two presented applications on the likelihood that they would personally acquire it.

With this study, we introduce a third condition, called Permissions Inline. This treatment was designed to separate the location of the privacy information from its format. It showed the standard Android Permissions Display, but positioned on the app display screen (where Privacy Facts is located) rather than in the standard location after the user tapped "Download." This condition tested whether it was only the existence of any privacy information on the application screen that changed behavior, or the checklist format and position.

We used the graphic design of the permissions display from the current Google Play store; however, we modified the labels to present the same information as our Privacy Facts display (including health, nutrition, advertising, and analytics).

	Lab Study		Online Study		Diff. from Android	p-value	Permissions Inline (n=123)	Diff. from android	p-value	Inline v. Facts
	Privacy Facts (n=10)	Android Display (n=10)	Android Display (n=120)	Privacy Facts (n=123)						
Wordoid!	60%	50%	40.8%	61.0%			49.6%			
Word Weasel	30%	50%	59.2%	39.0%	20%	0.002	50.4%	9%	0.198	0.095
Twidroyd	70%	20%	25.0%	52.9%			35.8%			
Plume	30%	80%	75.0%	47.2%	28%	< 0.001	64.2%	11%	0.051	0.014
DroidScan Lite	90%	90%	73.3%	60.2%			62.6%			
M. Doc Scan Lite	0%	10%	26.7%	39.8%	-13%	0.031	37.4%	-11%	0.076	0.784
Calorie Counter	70%	100%	55.8%	73.2%			73.2%			
Doc's Diet Diary	30%	0%	44.2%	26.8%	17%	0.005	26.8%	17%	0.005	1
Rdio	40%	30%	17.5%	28.5%			22.8%			
Spotify (<i>brand</i>)	60%	70%	82.5%	71.5%	11%	0.048	77.2%	5%	0.340	0.381
Flight Tracker	40%	20%	40.8%	35.0%			37.4%			
iFlights (<i>rating</i>)	50%	80%	59.2%	65.0%	-6%	0.358	62.6%	-3%	0.601	0.791

Table 3. Application selections in the laboratory and online studies. The application that requested access to fewer permissions (the privacy-protective choice) is always displayed on top. Statistics for the online study are comparisons to the base Android display. The right-most column shows the significance between the checklist and the inline permissions. Differences in bold, Fisher's Exact. Comparisons with the Android display were planned contrasts. The final comparison between the permissions inline and privacy facts display is Holm-corrected with an adjusted alpha of 0.01667.

An example of this is shown in Figure 1C.

Demographics

Of our 366 MTurk participants 59% were male and 41% were female (markedly different from our lab study). Our participants were between 18 and 63 years old, with an average of 28. All of our participants had experience downloading Android applications from the market (the 12 who did not were removed from this analysis).

Application selection

Overall the privacy facts display (changed format and position) had a stronger effect on participants application selections than only moving the permissions inline (changed position).

Privacy Facts display

In three of the four standard comparisons, significantly more privacy facts participants than Android participants chose the low-requesting app. Only for the document scanner did more participants in the standard Android condition choose the low-requesting app, and this difference was not significant.

For the Twitter choice, nearly three-quarters of the Android display participants chose Plume (high-requesting). One participant captured many of the common reasons for making this choice, reflecting, "Plume has 35,000 more reviews, which suggests to me that this is the more popular, more frequently used application. The description includes a list of everything you can do with the app and those all seem like useful features." However when presented with the privacy facts checklist, the two applications were selected at almost the same rate, with slightly more selecting Twidroyd. Here participants noted and cited the permissions information. One stated, "I picked the one that respects privacy more. The other gets too much personal info." Another participant wrote, "Plume collects too many personal facts."

For the special comparisons, rating and brand recognition outweighed privacy. However, even when one of the choices

was a well-known brand privacy facts participants were significantly more likely than Android participants to select the relatively-unknown, low-requesting choice. For the flight tracking choice, more participants chose iFlights (high-requesting) over Flight Tracker. Although participants thought iFlights "sounds like an iPhone port," many believed it had a cleaner UI, but the top reason given was the rating difference. Flight Tracker's 3-stars seems to have outweighed all other factors. For the streaming music choice, Spotify (high-requesting) had much higher brand recognition (although again, both are real services). In the Android permissions display condition over half of the people (66/104) who selected Spotify explicitly stated that they had already heard it was very good or that they or friends use Spotify. One participant said "Spotify is pretty popular and I have never heard of Rdio." Spotify collected much more information than Rdio. but in this case we see that brand information trumps privacy concerns, though there is still a significant shift (11%) in favor of Rdio in the privacy facts condition.

Permissions Inline

As shown in in Table 3, the permissions inline display, while in the same place and often more space-consuming than the checklist, did not have as large an effect on users' decisions. In only one of the four standard comparisons, the nutrition application, was this change significant, and in most cases it underperformed the checklist display (significantly underperforming for twitter apps). This suggests that in addition to moving privacy information to the application display screen, it is important to present that information in a holistic, clear, and simple way if it is to impact users' app selections.

Free responses

Across the free-text responses for why applications were selected by participants in the Android display conditions, privacy was only mentioned by one participant, and permissions were mentioned by four others. Across the privacy facts checklist condition privacy was mentioned by 15 participants,

and permissions were mentioned by seven more. Information or info were mentioned by 49 people in the privacy facts checklist condition, but by only six participants using the Android display. Based on these responses privacy and personal information seem to have factored more strongly into the decisions of those who saw the privacy facts checklist.

Similar to our lab study, many participants, when directly asked, said they did not notice the privacy facts checklist. Of the 125 participants who were shown the privacy facts checklist, 49 (39.2%) reported in a free-text response having not noticed or paid any attention to the display. Both those people who did and those who did not notice the display provided reasons for why they ignored it, or believed it was not necessary:

- “I noticed the Privacy Facts but it really didn’t influence me that much. I feel like with social networking it’s so much easier to get contacts, photos, or information of someone.”
- “It didn’t influence my decision even though i noticed it. I tend to pay more attention to ratings and usefulness then anything else.”
- “No, not really. It’s not the most important factor. I don’t keep a bunch of vital personal info on my phone, so no worries. I think people who do are really stupid.”

There were also users who found the privacy facts display helpful and made their decisions based on it:

- “Yes. I believe the privacy information is helpful. It would only bother me if I saw something that didn’t make sense for the app to use. However, I am not terribly concerned about privacy.”
- “Yes. It only influenced me if it seemed to be the only thing to distinguish between the two apps.”
- “Yeah, I always check that stuff. I want to know exactly what is happening to and with my data from that program when I use it. It was useful though I wish some apps would go into greater detail.”

Participants who both used and didn’t use the display still had misconceptions about companies, sharing information, and the market. Many assumed that all applications collect the same information. One participant who didn’t look at the display said she did not because, “I assume they always say the same thing....”

Participants also continued to believe external forces protect them. One said, “Yes I saw the privacy facts. That didn’t really affect my decision as companies are required to protect consumer’s information and companies don’t really wanna get sued for breach of security so I am not worried about all that.” Another stated the continued belief that the market is internally well-regulated, “I think it is trustworthy, I would assume google play keeps a tight leash on that stuff.”

Finally, one participant gave an answer that applies quite broadly, and mirrors work by Staddon et al. [23], “Yes, I noticed the privacy facts but it didn’t effect [sic] my decision because I don’t really know what the negative impacts of the




Figure 2. A series of factors users consider when decided on applications. Ranked by the number of users reporting a 4 or 5, where 5 is “Very important.”

information they obtain would be.” Understanding the potential harm in allowing access to certain types of data remains difficult for consumers both in smartphones and other digital domains.

Self-reported decision factors

We also asked our online participants to rank a series of factors in their personal application-selection process from “Not important” to “Very important.” The results of this are presented in Figure 2. Permissions ranked 8th (of 11), just below two metrics of popularity and just above the size of the application. 80% of participants said ratings and reviews were important or very important, compared with only 43% who said that permissions were important or very important.

This result seems to align with how often participants across our tests tended to ignore permissions.

Limitations

Our short checklist display had items that consumers were able to explain in most cases. Analytics and Personal Information were the most problematic. Participants were generally correct when defining Analytics, but often created more invasive definitions that were not intended. Personal Information was more difficult, as it was too vague and many participants listed other types of data that they then realized were covered by another item on the list. We will continue to further refine the terms and types of information that is most important to people.

One more significant design flaw with the display was that participants do not view permissions displays in the same way as they view privacy policies. They see this information only as items the phone can *take*, not things that they personally input. While we believe a complete privacy display should cover both user-provided information that is stored (i.e., medical or diet) *and* automatically collected information like location, this was not explained well by the current design.

Our lab study has many more female participants than male participants, and due to random condition assignment they were not evenly distributed. We note this as a potential limitation, though our results from the two studies are aligned, and we did not see such a similar gender imbalance in our online study.

Mechanical Turk also has its own set of limitations and biases, which we attempted to counter through a careful survey design. While we compared our two survey phases, they did not follow identical methodologies. Our lab study was more realistic, with users using actual cell phones, when on MTurk users saw the applications side by side, and could make direct and visual comparisons. While the reasoning and behavior given seems similar, it is possible that our online survey users had an easier time making decisions, not due to our improved permissions display, but due to the side-by-side display. The only evidence we have to counter this is the permissions-in-place display did not perform as well as the privacy facts display, implying that the side-by-side display alone is not responsible for all the improvements we saw.

Finally, we tested only 12 applications in the studies described above (and an additional 12 in early pilots). We picked applications that seemed similar, functional, and would be unrecognized, but we would like to expand this work in the future to consider larger application datasets.

DISCUSSION

Our goal was to better understand how users select Android applications, and to make privacy and permission information a salient part of that process.

We found that users did not use the current permissions display. By moving privacy/permissions information onto the main screen and presenting it clearly and simply we could affect user decisions in cases where applications were similar. Users mostly appreciated the new privacy facts display, said they would use it to help make their decisions or at least glance at it, and found comparing applications in the market to be a difficult task where better displays would assist them.

Can we affect users' decisions?

The short answer, is yes—the privacy information on the application display screen affected user behavior. In laboratory responses and our online test we saw behavioral differences as well as differences in quality and tone of responses relating to private information.

We also found most people do not consider permissions when downloading new applications. Even when instructed to download applications, most users made decisions without ever pushing the button that would take them to the permissions display. Both our lab participants and our online participants also self-reported that they were aware of the display, but did not look at it. This was confirmed by our lab study participants who, when they did fully “download” applications, spent a median time of 2 seconds on the permissions display. While this was expected based on other research and our own earlier work, we now have evidence that the permissions are, at least partially, disregarded due to their position in the application selection process.

In online testing we found that having a privacy display would in cases of equivalent applications change user application selection, and do so more strongly than simply moving the permissions display to the main screen. We also seem to have initial results that indicate even relatively sizable privacy differences cannot outweigh some other factors such as very

popular applications (significant change, but not a majority) or differences in average ratings (3-star vs. 4-star apps, almost no change).

All of our participants had never seen a privacy facts display before, but were immediately able to make comparisons when specifically instructed to do so after the selection task. However, some simply did not believe privacy information was important or relevant to their decision. Some said it would depend on how much their friend (as part of the role-play) cared about his or her own privacy.

These results are similar to those seen in other labeling efforts. Consumers who care more about privacy, whether they have had a credit card stolen or have started receiving spam text-messages, are more likely to take advantage of labeling information. Even if the impact is not drastic, we see the privacy information on the main screen having an affect on selection behavior.

Do users enjoy, notice, and trust permissions information?

Participants in our studies reported being familiar with permissions displays and being aware that there are differences between applications. While this may seem unimportant or obvious, leveraging the awareness of privacy differences means creating interfaces, like checklists, that help consumers identify and compare differences should benefit users who want to make privacy-preserving decisions.

The terms on the current Android Permissions display remain difficult to understand and participants believed that there was little they could do as most of their information was already exposed. Participants reported that they did not, in most cases, read the information in the displays, and they did not select the permission groupings to see more details or try to better understand the terms. Even when the display was moved to the main screen, it does not have the impact of the privacy facts display.

Participants continued to report not being concerned with data sharing generally, partially due to a belief that companies are following laws and a strong belief that Android/Google is watching out for their safety as a consumer. While this is accurate in a very general sense, the specifics are quite far off from reality. Correcting the ubiquitous idea of Google Play as a safe, protected marketplace, must necessarily be changed if consumers are to protect themselves through understanding privacy and security in their decision-making process.

From both the lab and online studies we found that participants continued to report that other characteristics of applications are as important or more important than permissions, including: cost, functionality, design, simplicity, rating, number of ratings, reviews, downloads, size, and others. Continuing to understand how much privacy can compete and offset other aspects is important future work as consumers battle with a crowded and complex market.

When asked why an application was collecting a type of information, participants most often stated they did not know, but would occasionally venture possibilities. All of our lab

study participants wanted to better understand why applications required the permissions they did.

Finally, participants overwhelmingly trusted the application in both the privacy facts display and the permissions display. The question of trusting the information was one most had never considered, and actually gave some participants pause as they realized for the first time that this information might not be accurate. Again, users believe this information is correct, is being verified, and will assume they misunderstand something before they would believe the displays are incorrect. Mistakes in the permissions are not recognized, even when directly discussed. Users will assume they themselves are wrong, not the policy.

CONCLUSION

Smartphones have unprecedented access to sensitive personal information. While users are aware of this, generally, they may not be considering privacy when they select applications to download in the application marketplace. Currently, users have only the Android permissions displays to help them make these application selection decisions, screens which are placed after the main decision occurs, and are not easily understood. We sought to investigate how we could make permissions and privacy play a true part in these decisions. We created a short “Privacy Facts” display, which we then tested in 20 in-lab exercises and an online test of 366 participants. We found that bringing information to the user *when* they are making the decision and by *presenting* it in a clearer fashion, we can assist users in making more privacy-protecting decisions.

ACKNOWLEDGEMENTS

We acknowledge our colleagues at Carnegie Mellon University, and would like to thank Alessandro Acquisti and Sunny Consolvo. This work was supported by the National Science Foundation under Grant DGE-0903659 (IGERT: Usable Privacy and Security). Additional support was provided by NSF grants CNS-1012763, CNS-0905562 and DGE 0903659, by CyLab at Carnegie Mellon under grants DAAD19-02-1-0389 and W911NF-09-1-0273 from the ARO as well as Google.

REFERENCES

1. Au, K., Zhou, Y., Huang, Z., Gill, P., and Lie, D. Short paper: a look at smartphone permission models. In *Proceedings of the 1st ACM workshop on Security and privacy in smartphones and mobile devices (SPSM '11)* (2011).
2. Barrera, B., Kayacik, H., van Oorschot, P., and Somayaji, A. A methodology for empirical analysis of permission-based security models and its application to android. In *In Proceedings of the 17th ACM conference on Computer and communications security (CCS '10)* (2010).
3. Barrera, D., Clark, J., McCarney, D., and van Oorschot, P. C. Understanding and improving app installation security mechanisms through empirical analysis of android. In *2nd Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)* (2012).
4. Egelman, S., Tsai, J., Cranor, L., and Acquisti, A. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proceedings of the 27th international conference on Human factors in computing systems*, ACM (2009), 319–328.
5. Enck, W., Gilbert, P., Chun, B., Cox, L., Jung, J., McDaniel, P., and Sheth, A. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *In Proceedings of the 9th*

USENIX conference on Operating systems design and implementation (OSDI'10) (2010).

6. Felt, A., Chin, E., Hanna, S., Song, D., and Wagner, D. Android permissions demystified. In *In Proceedings of the 18th ACM conference on Computer and communications security (CCS '11)* (2011).
7. Felt, A. P., Egelman, S., Finifter, M., Akhawe, D., and Wagner, D. How to ask for permission. In *USENIX Workshop on Hot Topics in Security (HotSec) 2012* (2012).
8. Felt, A. P., Egelman, S., and Wagner, D. I've got 99 problems, but vibration ain't one: A survey of smartphone users' concerns. In *2nd Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)* (2012).
9. Felt, A. P., Ha, E., Egelman, S., Haney, A., Chin, E., and Wagner, D. Android permissions: User attention, comprehension, and behavior. In *Symposium on Usable Privacy and Security (SOUPS) 2012* (2012).
10. Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D., and Konstan, J. Stopping spyware at the gate: A user study of privacy, notice and spyware. In *In Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS 05)* (2005).
11. Juniper Networks. Mobile malware development continues to rise, android leads the way, 2011. <http://globalthreatcenter.com/?p=2492>.
12. Kelley, P., Consolvo, S., Cranor, L., Jung, J., Sadeh, N., and Wetherall, D. A conundrum of permissions: Installing applications on an android smartphone. In *Financial Cryptography and Data Security*, vol. 7398. 2012, 68–79.
13. Kelley, P. G., Bresee, J., Cranor, L. F., and Reeder, R. W. A “Nutrition Label” for Privacy. In *Proceedings of the 2009 Symposium On Usable Privacy and Security (SOUPS)* (2009).
14. King, J. “How come i'm allowing strangers to go through my phone?– Smartphones and privacy expectations, 2013. <http://jenking.net/mobile/>.
15. Kleimann Communication Group Inc. Evolution of a prototype financial privacy notice., February 2006. <http://www.ftc.gov/privacy/privacyinitiatives/ftcfinalreport060228.pdf>.
16. Labs, M. McAfee threats report: Third quarter 2011, 2011. <http://www.mcafee.com/us/resources/reports/rp-quarterly-threat-q3-2011.pdf>.
17. Lin, J., Sadeh, N., Amini, S., Lindqvist, J., Hong, J. I., and Zhang, J. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. *UbiComp '12*, ACM (2012), 501–510.
18. Lockheimer, H. Android and security, 2012. <http://googlemobile.blogspot.com/2012/02/android-and-security.html>.
19. Lunden, I. Google play about to pass 15 billion app downloads? psst! it did that weeks ago, 2012. <http://techcrunch.com/2012/05/07/google-play-about-to-pass-15-billion-downloads-psst-it-did-that-weeks-ago/>.
20. Namestnikov, Y. It threat evolution: Q3 2011, 2011. http://www.securelist.com/en/analysis/204792201/IT_Threat_Evolution_Q3_2011.
21. Rashid, F. Y. Black hat: Researchers find way to “bounce” malware into google app store, 2012. <http://www.scmagazine.com/black-hat-researchers-find-way-to-bounce-malware-into-google-app-store/article/252098/>.
22. Smetters, D., and Good, N. How users use access control. In *In Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS 09)* (2009).
23. Staddon, J., Huffaker, D., Brown, L., and Sedley, A. Are privacy concerns a turn-off? engagement and privacy in social networks. In *Symposium on Usable Privacy and Security (SOUPS)* (2012).
24. Stevens, G., and Wulf, V. Computer-supported access control. *ACM Trans. Comput.-Hum. Interact.* 16, 3 (Sept. 2009), 12:1–12:26.
25. Vidas, T., Christin, N., and Cranor, L. F. Curbing android permission creep. In *W2SP 2011* (2011).