

Choice Architecture and Smartphone Privacy: There's A Price for That

Serge Egelman, Adrienne Porter Felt, and David Wagner

University of California, Berkeley
{egelman, apf, daw}@cs.berkeley.edu

Abstract. Under certain circumstances, consumers are willing to pay a premium for privacy. We explore how choice architecture affects smartphone users' stated willingness to install applications that request varying permissions. We performed two experiments to gauge smartphone users' stated willingness to pay premiums to limit their personal information exposure when installing new applications. We found that when participants were comparison shopping between multiple applications that performed similar functionality, a quarter of our sample responded that they were willing to pay a \$1.50 premium for the application that requested the fewest permissions—though only when viewing the requested permissions of each application side-by-side. In a second experiment, we more closely simulated the user experience by asking them to value a single application that featured multiple sets of permissions based on five between-subjects conditions. In this scenario, the requested permissions had a much smaller impact on participants' responses. Our results suggest that many smartphone users are concerned with their privacy and are willing to pay premiums for applications that are less likely to request access to personal information. We propose improvements in choice architecture for smartphone application markets that could result in decreased satisficing and increased rational behavior.

Keywords: Smartphone Security, Privacy, Ubiquitous Computing

1 Introduction

Architecture starts when you carefully put two bricks together. There it begins.

—Ludwig Mies van der Rohe

Nearly 90% of U.S. adults own cellular telephones [23, 45], and it is estimated that over 40% of these are smartphones [36]. Smartphones pose a challenging information security problem: users need to regulate how applications access the private information that is stored on their phones. Their smartphones often store sensitive personal data, such as lists of contacts, financial information (e.g., mobile banking), location information (e.g., GPS, WiFi SSIDs, and cellular tower information), and sensor data (e.g., cameras, microphones, and accelerometers). Smartphones need to simultaneously protect this data and support the installation of a variety of third-party applications.

Google’s Android addresses this problem with user-granted *permissions*. Permissions govern an application’s ability to perform actions on a phone that make use of either personal data or sensor hardware. For example, an application can only read the user’s list of contacts if it has the `READ_CONTACTS` permission. When a user installs an application from the Android Market, the central application repository, he or she is shown a warning screen that displays the set of permissions that the respective application requires. In order to complete the installation, the user must consent to granting all of the requested permissions to the application. Currently, this notice-and-consent process is all-or-nothing; the user cannot selectively grant or decline a subset of the permissions (i.e., the user must opt to not install the application in order to deny the requested permissions).

In this paper, we examine whether Android’s notice-and-consent process influences users’ purchasing decisions. Specifically, we evaluate how the Android Market *choice architecture*¹ influences users’ abilities and desires to protect their privacy, as evidenced by their stated willingness to pay premiums for applications that request fewer permissions. To explore this topic, we performed two online experiments: one to examine the extent to which users will consider permissions when comparison shopping, and another to examine the role of permissions when users are valuating a specific application. We designed our experiments to study the two primary shopping behaviors supported by the Android Market: function-specific searches and application-specific searches.

During a *function-specific search*, users seek previously-unfamiliar applications to perform specific tasks. When performing function-specific searches, users do not have a particular application in mind and are therefore willing to consider several different applications to fulfill the desired function. For example, choosing one particular flashlight application amongst many is an example of a function-specific search. During an *application-specific search*, users seek a particular application that is known to them, such as through word of mouth, “popular” application lists, or advertisements. When performing application-specific searches, users are unlikely to compare features between alternate applications. For example, the decision of whether or not to install Angry Birds is an example of an application-specific search.

In our first experiment, we asked users to select an application from among a set of applications with similar functionality but different permission requests. This experiment tested whether participants were willing to pay a privacy premium for an application that requested fewer permissions than the cheaper alternatives, when those alternatives were presented side-by-side. We found that 25% of our participants stated a willingness to pay a \$1.50 premium on a \$0.49 application in order to grant the fewest permissions.

In our second experiment, we focused on application-specific searches. Under the guise of being a software company, we solicited users to participate in a private beta test of a new application. Participants submitted bids for the amount of compensation that they would require to test our fictitious application on their phones. These bids were proxies for participants’ willingness to install the application. We constructed several between-subjects conditions by varying the permissions that participants saw. We also

¹ The term “choice architecture” refers to the way in which options are presented to people, as these design decisions can have a profound impact on decision-making [39].

asked participants whether they would rather use a \$0.99 version of the application or a free version supported by behavioral advertising. We made it clear that the advertisements would be targeted based on data collected as a result of the requested permissions.

Unlike our first experiment, wherein privacy-conscious participants opted to pay the highest premium for the fewest permissions, we observed that participants were satisficing under the more realistic conditions of the second experiment. We observed that only the request for a user's list of contacts had a significant effect on their stated willingness to install the application; requests for location data or access to the user's photos had no observable difference over the control condition. Additionally, around 80% of participants stated that they would be willing to receive targeted advertisements, regardless of the permissions used for the targeting, if it would save them \$0.99.

We present the following findings and contributions:

- Most prior work has focused on smartphone users' preferences for sharing location data. We examine several other types of data stored on smartphones, and we find that users are less concerned about location than other types of data. Observable changes in participant behavior were significantly more likely to be attributed to concerns over address book data than location information.
- We contribute to the growing body of work on willingness to pay for privacy by examining smartphone users' decisions holistically: we measure privacy behaviors as part of a larger value proposition. We show that 25% of participants in our first experiment were willing to pay a \$1.50 premium on a \$0.49 application in order to grant it permission to the least amount of personal data, when the options were presented side-by-side for easy comparison. However, the choice architectures of current smartphone application markets do not actually allow for such comparisons. Our second experiment better approximated these current choice architectures. We found that when users are considering a particular application, they satisfice by downplaying their privacy concerns in favor of other considerations until those concerns reach a threshold, only then do they consider privacy.
- Our results lead to two suggestions. First, privacy-conscious users may be willing to spend more money for an application if the choice architecture supported comparison shopping for privacy, such as by annotating search results with privacy icons. Second, users may be less likely to satisfice if the decision to install a particular application were decoupled from the decision to grant it a set of permissions. Specifically, our results indicate that users may be better served by presenting permission requests individually, when the data is actually needed, rather than requiring an entire set of permissions to be granted at install-time.

Finally, while we performed our experiments using the Android platform, we believe our results are generalizable to other smartphone platforms.

2 Background

In this section we provide an overview of application permissions on the Android platform, previous research on smartphone privacy that has focused on location sharing, and previous research on willingness to pay for privacy.

2.1 Android Permissions

Users find applications in the Android Market by searching by name, searching by keyword, or browsing lists of popular and featured applications. When the user selects an application from a search result or list, he or she arrives at the application's description page. The description page includes developer information, screenshots, user reviews, the price, and a "Download" button. When the user presses "Download," the Market displays a warning screen that lists the application's permissions (Figure 1).

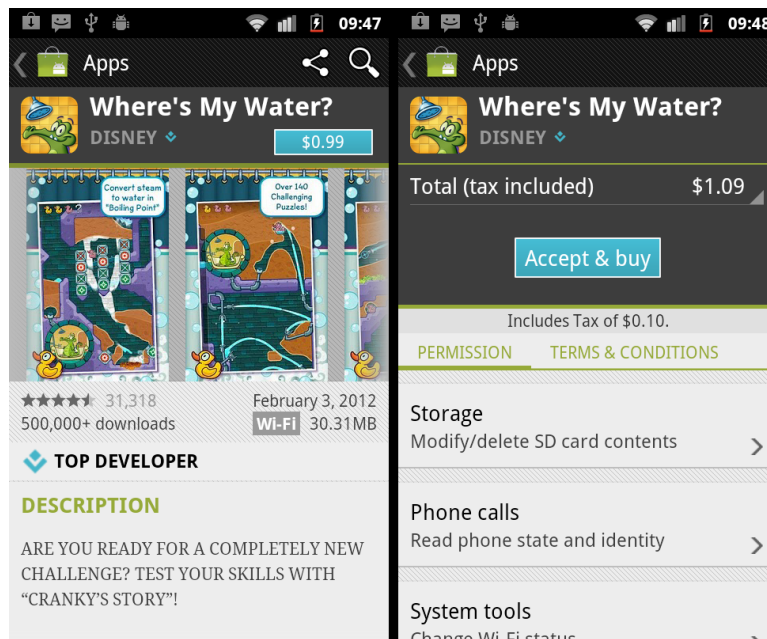


Fig. 1. The left screenshot shows an application description page in the Android Market, and the right screenshot shows the warning screen with the requested permissions.

Permissions govern access to privacy- and security-relevant actions on the phone, such as reading contacts, sending text messages, reading call history, connecting to the Internet, and turning WiFi on and off. Developers specify which permissions their applications require, and the user must agree to all of the application's permissions on the warning screen in order to install the application. If the user does not consent to any of the permission requests, he or she should cancel the installation. Users cannot selectively grant or deny permissions on a per-permission basis.

Applications often need permissions to implement their core features. However, smartphone applications also widely request access to private information for advertising, marketing campaigns, tracking users across applications, and other secondary uses [4, 17, 18]. At times, these practices have generated public outrage; for example,

consumer complaints forced Apple to promise to restrict how iOS applications can access contact lists [37]. Past research has established that free applications request more permissions than paid applications because many free applications share user data with advertising networks to generate revenue [11, 35]. Consequently, users who compare free and paid applications can choose to withhold personal information from advertisers by paying for applications. This motivates our exploration into whether users are willing to pay to protect the personal information that is on their smartphones.

In past work, we evaluated whether Android users pay attention to, comprehend, or act on permissions [22]. However, the previous study was limited to experimentally determining attention and comprehension rates, and we did not consider users' installation tradeoffs. We extend the past work by empirically testing the influence of permissions on application selection processes. We also designed this study's experiments to be robust to the low attention and comprehension rates that our prior study reported. In our experiments, we showed participants static screenshots so that they did not ignore the permissions by rapidly clicking through them, and in our second experiment we modified the permission warnings to avoid the comprehension problems that their study described. As such, our results represent the decisions of fully-informed users.

Researchers have built several tools to help users control how applications use their private information. Appfence [27] is a privacy tool that helps users protect their private data from exfiltration; it substitutes shadow data in place of private data and blocks network communication that contains user data. Kirin [19] operates under the assumption that users do not understand permissions and provides security rules to automatically accept or reject sets of permissions. Apex [34] lets users selectively grant or reject applications' permission requests. We focus on evaluating users' willingness to share their personal information with applications, rather than on building new tools.

2.2 Location Privacy

Most research on smartphone privacy has focused on users' willingness to share location data with their social contacts [8, 13, 29]. Three independent studies [13, 31, 44] found that the identity of the person with whom the user was sharing was the most important factor that influenced users' sharing decisions, whereas Barkhuus [7] found that the user's current location matters more. Anthony et al. [6] observed 25 undergraduates for a week and found that they fell into three categories: people who never shared their location, people who always shared their location, and people who are willing to share depending on both the requester and the current location. This past research thoroughly establishes that many users have concerns about sharing location with social contacts.

Other research has examined users' comfort with behavioral advertising, both in the context of mobile devices and on the web. Targeted advertising often makes use of personal information like age, gender, and location. Kelley et al. [30] report that 19 of 24 people "strongly agreed" that online companies should never share personal information unless it has been authorized by the user. McDonald and Cranor [32] found that many users are uncomfortable with behavioral advertising, and a large-scale 2009 survey similarly found that 50.5% of respondents were uncomfortable with online advertisers using browsing history [40]. However, despite this professed discomfort, people report liking and wanting relevant advertisements [40].

Our work is novel in that we explore users' willingness to share personal data stored on smartphones beyond location information. We examined several types of smartphone data, such as contacts, audio, and photos. Also, we did not direct participants to think about sharing with social contacts or advertisers. As in current smartphone application markets, study participants determined on their own how they thought the requested permissions would be used by applications. Different users might be concerned about social contacts, advertisers, police, governments, health insurance companies, etc.

2.3 Willingness to Pay for Privacy

Users do not always act in accordance with their professed privacy concerns [2, 38]. People are sometimes willing to trade private information for convenience, functionality, or financial gain, even when the gains are very small [26]. Previous studies have found that participants will sell their location information to researchers for 10 GBP and companies for 20 GBP [14, 15]. Good et al. asked people to install applications after viewing privacy statements; regardless of privacy, they found that people will install applications if they think the utility is high enough [25]. Another study found that many people would execute code from a stranger for \$1.00 or less [12]. Additionally, people's privacy tradeoffs are not always self-consistent, and vary based on how the choices are framed. Two studies reported that users value their privacy differently when asked to pay to protect their privacy than when asked to accept payment for personal information, although most people only placed a small financial value on their privacy in both cases [3, 26].

Acquisti [1] hypothesized that privacy-concerned people are willing to trade privacy for small gains because they are not economically rational agents with respect to privacy. He attributed users' actions to three factors that reduce economic rationality: incomplete information (i.e., users are not aware of the risks associated with privacy invasions), bounded rationality (i.e., people cannot calculate all of the payoffs associated with privacy-preserving strategies), and psychological distortions (e.g., hyperbolic discounting, self-control problems, and immediate gratification). A survey of 119 people supported this hypothesis: many survey participants greatly overestimated or underestimated the likelihood and magnitude of privacy abuses, were unable to remember all of the parties involved in standard financial transactions, and were less likely to use privacy-enhancing technologies if the perceived risk was in the distant future [2].

A corollary to Acquisti's hypothesis is that people will act more like economically rational agents if they operate within a system that mitigates the effects of incomplete information, bounded rationality, and psychological distortions. Several studies have explored whether this is true. Gideon et al. [24] and Tsai et al. [41] asked users to purchase items using Privacy Finder, a search engine that included privacy ratings in search results. Privacy ratings in search results provide users with additional information and make the tradeoffs easier to compute. Both studies found that participants were willing to pay premiums to purchase privacy-sensitive goods from merchants with better privacy policies when privacy ratings were embedded in search results. Good et al. [25] similarly found that privacy and security are important factors when choosing between two applications with similar features, but not when considering an application on its own. This indicates that the timing and placement of privacy information is crucial. In a

related study, Egelman et al. [16] tested the timing and placement of privacy indicators in Internet shopping and found that even non-privacy-conscious shoppers will pay more for privacy when indicators are presented before visiting websites rather than after.

Like past research on the economics of privacy, we aim to measure users' willingness to trade privacy for financial gain. However, we focus specifically on Android applications and smartphone data. We performed two experiments: the first showed participants side-by-side privacy information similar to that of Privacy Finder, and the second showed participants an individual application similar to the choice architecture of the Android Market and other smartphone application repositories. In the second experiment, we asked users to bid in a reverse auction for a chance to participate in a beta test of a particular smartphone application. We used these bids as proxies for willingness to pay and modeled our reverse auction on past studies that have used reverse auctions to gauge users' willingness to share information [14, 15, 28].

3 Privacy and Comparison Shopping

In October 2011, we deployed a survey to test whether smartphone users would be willing to pay more for an application if it requested fewer permissions than less-expensive applications that offered similar functionality. We asked participants to view screenshots from the Android Market of four fictitious applications. While each respondent saw what appeared to be four different applications, we counter-balanced the names, descriptions, and imagery, while controlling for price and requested permissions. Overall, we observed that when comparison shopping between applications, a quarter of participants stated a willingness to pay four times as much for an application that accessed less of their data, when that application was shown alongside screenshots of alternative applications that cost less but requested access to more data.

3.1 Methodology

For our survey, we created screenshots of four fictitious news aggregation applications as they might appear in the Android Market, but with one important difference: current smartphone application markets only allow application permissions to be viewed serially; we showed the four applications side-by-side to aid participants in contrasting their differences. We asked participants to choose which of the four they would be most willing to purchase. The purpose of this experiment was to examine whether participants would be willing to pay a privacy premium, when that option was apparent to them. The amount of personal information collected by each application was signaled by a permission request screen. Each fictitious application featured one of four possible prices: \$0.49, \$0.99, \$1.49, and \$1.99. These prices corresponded to four sets of requested permissions:

1. **\$1.99**—INTERNET
2. **\$1.49**—INTERNET and ACCESS_FINE_LOCATION
3. **\$0.99**—INTERNET and RECORD_AUDIO
4. **\$0.49**—INTERNET, RECORD_AUDIO, and ACCESS_FINE_LOCATION

We chose to focus on these three permissions for the following reasons:

- `INTERNET` controls Internet access, and it is the most frequently requested permission [21]. It also would be needed by a news reader to serve its intended purpose.
- GPS location data, represented by the `ACCESS_FINE_LOCATION` permission, has been heretofore the focus of most smartphone privacy research (see Section 2).
- Recent research has found that the ability to record audio (i.e., the `RECORD_AUDIO` permission) may be one of the most concerning permissions to users [20, 22].

We paired permissions to prices such that the least privacy-invasive application, which only requested `INTERNET`, had the highest price. The most privacy-invasive application, which requested all three permissions, had the lowest price. Participants with privacy concerns would need to pay a premium of \$1.50 over the base price of \$0.49 for the least privacy-invasive application. Since a previous study suggested that users are less concerned with location privacy than with an application’s ability to make audio recordings [20], we set the price of the application with the `INTERNET` and `RECORD_AUDIO` permissions to be the second least expensive. Finally, the application with the `INTERNET` and `ACCESS_FINE_LOCATION` permissions cost \$1.49, the penultimate price.

If participants had viewed four identical applications that only differed based on price and permissions, the purpose of the study would be obvious, which might have an impact on participants’ responses. To minimize the potential for the Hawthorne effect, we created four seemingly-unique applications with different names, manufacturers, screenshots, descriptions, and icons. We counterbalanced these features such that each of the four price and privacy combinations was equally likely to be assigned to each application. A second concern was that users would anchor on the first price. To compensate for anchoring effects, we also counterbalanced the order in which the price and privacy conditions were presented; the conditions were either ordered from the lowest to the highest price, or from the highest to the lowest. Figure 2 depicts an example of what participants saw.

We were initially concerned that participants’ responses would differ based on their general interest level in applications of this type. In particular, participants who did not have any interest in news readers might not put much thought into choosing one of the four news readers they were shown. To test whether this occurred, we used the first page of the survey to randomly display one of the four applications and we asked participants to indicate their willingness to purchase it using a 5-point Likert scale (i.e., “extremely unlikely,” “unlikely,” “indifferent,” “likely,” or “extremely likely”). In order to equally distribute anchoring effects, we also randomized the price that was shown alongside the application. (For this preliminary question, we showed the application’s description page but not the requested permissions, since those were irrelevant to this question.) If participants’ interest levels correlated with their choice of application in the experiment, we planned to remove all participants who indicated that they were uninterested. However, using Pearson’s correlation, we observed no statistically significant correlations between participants’ stated willingness to install an application and their selection in the subsequent experiment. As such, we concluded that we did not need to remove any participants due to lack of interest.

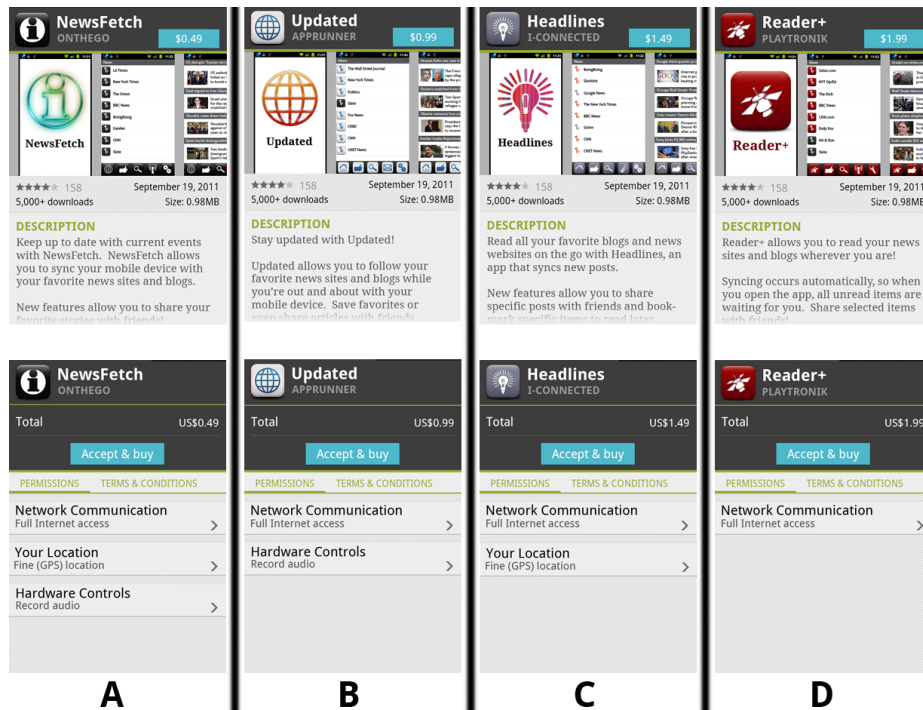


Fig. 2. We asked participants to choose the application they would be most willing to purchase. Privacy and price were inversely proportional, such that the most expensive application requested the fewest permissions. The other distinguishing features were counterbalanced.

We recruited participants using Amazon’s Mechanical Turk. We limited participation to U.S. residents over the age of 18. We did not limit participation to smartphone users, but we asked about smartphone usage in the demographics section of the survey in order to perform post-hoc analysis. In this manner, we accumulated a total of 483 valid survey responses, after screening out ten incomplete and questionable responses.²

3.2 Results

We considered three hypotheses:

$H_0 =$ Each price/privacy variant will be chosen with equal probability

$H_1 =$ Cost-sensitive participants will choose the least-expensive option

$H_2 =$ Privacy-sensitive participants will choose the high-privacy option

² We identified invalid results based on two factors. First, we included several questions that required free text responses, such as, “why or why not would you purchase this application.” Using these questions, we deleted surveys that contained nonsensical responses. Second, in addition to asking participants to select the application that they were most willing to purchase, we also asked them to select the application that they were least willing to purchase. We removed participants who gave the same answer to both questions.

Price	Permissions	Total
\$1.99	INTERNET	120 (24.84%)
\$1.49	INTERNET ACCESS_FINE_LOCATION	74 (15.32%)
\$0.99	INTERNET RECORD_AUDIO	77 (15.94%)
\$0.49	INTERNET ACCESS_FINE_LOCATION RECORD_AUDIO	212 (43.89%)

Table 1. The four price/privacy variants, sorted by decreasing price and privacy. The third column depicts the number of participants who chose each variant.

The null hypothesis (H_0) was that participants would select from the four variants with equal probability because they did not have any real financial stake in the selection. We found that this was not the case; a chi-square test showed that selections significantly diverged from uniformity ($\chi_3 = 102.9, p < 0.0005$). Thus, H_0 is rejected.

Overall, we observed that 43.9% of participants—a plurality—selected the cheapest application with the greatest number of requested permissions. At the same time, though, we observed that the second most popular variant was the most expensive one (24.8% of 483), which also afforded the greatest privacy by granting access to the least amount of data. The complete results are depicted in Table 1. Upon performing post-hoc testing using the Bonferroni correction to account for multiple tests ($\alpha = 0.01$), we observed statistically significant differences between the high privacy/high cost variant and each of the others (\$0.49: $\chi_1 = 25.49, p < 0.0005$; \$0.99: $\chi_1 = 9.39, p < 0.002$; \$1.49: $\chi_1 = 10.91, p < 0.001$). This indicates that significantly more participants chose the cheapest variant when compared to each of the others. However, significantly more participants chose the high-privacy variant than the two mid-price/privacy variants.

Influencing Factors In order to test H_1 and H_2 , we needed to know whether participants were cost-sensitive or privacy-sensitive. After participants chose an application, we asked them to rate how the following factors influenced their decisions using a 5-point Likert scale, ranging from “no influence” to “strong influence”:

- Number of Downloads
- Icon
- Size of App
- Permissions Requested
- Description
- Name of App
- Rating/Reviews
- Familiarity with App
- Cost
- Manufacturer of App

	\$1.99	\$1.49	\$0.99	\$0.49
1.	Permissions (46%)	Description (41%)	Cost (33%)	Cost (62%)
2.	Description (23%)	Permissions (18%)	Description (15%)	Description (15%)
3.	Icon (9%)	Cost (16%)	Icon (14%)	Rating/Reviews (6%)
<i>n</i>	120	74	77	212

Table 2. Participants reported the factor that most influenced their decision to select a particular price/privacy variant. Each columns list the top three factors listed for each variant.

We performed Pearson correlations between the prices of the variants that participants selected and each of the factors above. After correcting for multiple testing using the Bonferroni correction ($\alpha = 0.005$), we observed statistically significant correlations between the degree to which participants reported being influenced by price and the price of the variant they actually selected ($r = -0.39, p < 0.0005$). We also observed a statistically significant correlation between the degree to which participants reported being influenced by permissions and the price of the variant they selected ($r = 0.33, p < 0.0005$). These correlations support H_1 and H_2 : participants who chose lower-price variants were more concerned about price, and participants who chose higher-price variants were more concerned about the permissions.

As previously mentioned, we counterbalanced factors across all four price and privacy combinations to make the applications appear unique. Despite this, we noticed a marginally significant correlation between the price of the variants that participants selected and their reported influence from the variant’s icon ($r = .13, p < 0.004$). In other words, the participants who chose a more expensive application were more likely to say that the icon influenced their choice. This was not because one icon was more appealing than the others; each of the four icons appeared next to each price with equal probability. Instead, we believe that the most likely explanation is that the attention participants gave to each of these factors can be modeled as a zero-sum game: participants who were less concerned with price were more willing to base their selections on additional factors (e.g., privacy & icons), whereas participants who were most concerned with price gave less attention to other factors.

Finally, we asked participants to select the primary factor that influenced their application selection from the ten above. As expected, the majority of participants who chose the \$0.49 low-privacy variant claimed that cost was the primary factor behind their decisions (62.3% of 212), whereas a plurality of the participants who chose the \$1.99 high-privacy variant claimed that the permissions requested were the primary factor (45.8% of 120). Table 2 lists the top three primary factors participants considered, separated by the price/privacy variant they selected.

Permission Necessity While each application variant mentioned similar functionality in the descriptions, our results may be confounded if participants believed that permission requests beyond Internet access reflected differences in application functionality. If so, they may have viewed the less-expensive alternatives as more functional rather than more privacy-invasive. To test this, we provided participants with a list of possible

permissions along with a picture of one of the four previous applications, *Headlines*,³ and asked them to select all of the permissions that they believed would be required for the application to function as described. The description read:

Read all your favorite blogs and news websites on the go with Headlines, an app that syncs new posts. New features allow you to share specific posts with friends and bookmark specific items to read later.

We asked participants to choose which of the following permissions they would expect the application to request in order for it to behave as described:

- Modify and/or delete your accounts
- Determine your physical location
- Read incoming or outgoing text messages (SMS)
- Read incoming or outgoing email messages
- Access the Internet
- Read your list of contacts
- Read your web browsing history
- Determine your phone number
- Determine which other apps are running
- Record audio
- Record video
- Send email messages
- Send text messages (SMS)
- Prevent device from sleeping
- Modify your storage card contents
- None of the above

We performed Phi correlations, corrected for multiple testing ($\alpha = 0.0125$), between whether participants thought a particular permission was required for the application's functionality and whether they previously selected an application that requested that particular permission. We observed no significant correlations with regard to `RECORD_AUDIO`, since very few participants believed that the application needed this ability to function (7.9% of 483). However, we did observe a statistically significant correlation with regard to `ACCESS_FINE_LOCATION` ($\phi = 0.21, p < 0.0005$). This indicates that some participants may have chosen the \$0.49 and \$1.49 variants because they believed that the `ACCESS_FINE_LOCATION` permission signaled desirable location-based features. We cannot test the correlation between users' perception of the `INTERNET` permission and their selection of applications because it was requested by all four variants. However, 91.3% of 483 participants correctly understood that Internet access would be required for these applications to function as desired.

³ We only showed participants the first screen with the application description and did not show them a permission request screen. Because we were concerned that participants may have been primed from the previous tasks to a particular set of permissions, all participants viewed the \$1.99 version, which only requested `INTERNET`. This is also the only permission required to perform the functionality in the application description.

We conclude that nearly all participants understood that the RECORD_AUDIO permission was unnecessary (92.1% of 483). However, 59.3% of these 445 participants were unwilling to pay a premium to deny the application this extraneous data. Although more people thought that the ACCESS_FINE_LOCATION permission was required, 51.9% of the 316 participants who thought that the permission was unnecessary were unwilling to pay a premium to deny the application this data (even when it was as little as \$0.50).

Demographics We collected demographic information regarding age, gender, type of phone, and general privacy sensitivity. Our sample from Mechanical Turk was 52.6% female with an average age of 31.58 ($\sigma = 10.25$). Upon performing a Mann-Whitney U test, we found no observable differences between gender and the price of the variant selected. Based on our sample, there was no evidence that age or gender were correlated with willingness to pay for privacy.

A total of 372 respondents (77.0% of 483) reported owning a smartphone.⁴ Android users represented 42.7% of our smartphone users. One potential confound is that existing Android users may understand the significance of permissions better than others. However, while Android users were significantly more likely to report that permissions influenced their decisions than non-Android users ($U = 19,848.5, p < 0.0005$), we observed no statistically significant differences with regard to which price/privacy variant they ultimately selected. Thus, while Android users may have been more familiar with the UI or the word “permissions,” they were no more likely to factor them into their decisions than non-Android users.

To gauge general privacy sensitivity, we asked participants to rate three statements using a 5-point Likert scale (from “I strongly disagree” to “I strongly agree”) so that we could categorize them along the Westin privacy index [43]:

1. Consumers have lost all control over how personal information is collected and used by companies.
2. Most businesses handle the personal information they collect about consumers in a proper and confidential way.
3. Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.

We classified users according to Westin’s metric [43]. Participants who agreed with the first statement and disagreed with the second and third statements were classified as Privacy Fundamentalists (26.1% of 483). Participants who disagreed with the first statement while agreeing with the second and third statements were classified as Privacy Unconcerned (5.6% of 483), while the remaining participants were classified as Privacy Pragmatists (68.3% of 483). While 30.2% of Privacy Fundamentalists were willing to purchase the high-privacy variant, compared to 23.3% and 18.5% of Privacy Pragmatists and Privacy Unconcerned, respectively, these differences were not statistically significant. Thus, we did not find a correlation between the Westin privacy index and users’ privacy-preserving smartphone selection behaviors.

⁴ One respondent reported not having a smartphone, but provided the make/model of an Android phone. We therefore added her to the 371 respondents who self-reported having smartphones.

4 Privacy in Context

In our first experiment, we observed that a quarter of our participants were willing to pay a premium for a smartphone application that did not request extraneous permissions. Much like previous work (e.g., [16, 24, 42]), the products that participants were comparing were tightly controlled: participants made side-by-side comparisons between price and effective privacy. This side-by-side scenario is an idealized choice architecture that is not representative of any current smartphone application markets. We therefore performed a second experiment to test a more realistic scenario: seeking out a specific application based on previous knowledge and then using additional information (e.g., price and permissions) to decide whether to proceed with installation.

In our second experiment, we examined how permission requests impacted participant behavior during application-specific searching. In order to increase ecological validity, we deceived participants by telling them that we were a smartphone application developer soliciting participants for a private beta. We asked participants to provide a bid for the amount of compensation that they would need to participate in this beta and to recommend a price at which we should sell our application in the Android Market. We used these bids as proxies for willingness to pay.

4.1 Methodology

The goal of our second experiment was to quantify the effect of permission requests on participants' valuations of a particular application. We posed as a company named "AirZoom" and recruited participants from Amazon's Mechanical Turk to participate in a "private beta test" of an application. We required participants to be at least 18 years of age, in the U.S., and current Android users. We performed a reverse Vickrey auction, modeled after Danezis et al.'s study on willingness to pay for location privacy [15]. In their study, they explicitly asked participants to bid on the amount they would need to be compensated in order to be tracked. Our belief is that in practice, users are never explicitly asked to give up privacy; a request for information is almost exclusively part of a larger value proposition (e.g., desirable location-based features, cheaper applications due to advertising subsidies, etc.). We wanted to examine the impact of permission requests on decision-making within the context of a larger value proposition: the amount of compensation participants would demand to install a given application. We asked participants how much they would need to be compensated to install and use a fictitious application for a month, as well as to suggest a reasonable price for us to charge for this application in the Android Market. We constructed several between-subjects treatments that differed based on the permissions requested.

To increase ecological validity, we used deception: we registered a new domain, `airzoom.net`, in order to convince participants that we were a software company that was conducting a private beta test of an Android application. We displayed screenshots of our application, shown in Figure 3, and asked participants to provide a bid for how much they would need to be compensated to participate. We also asked them to provide a suggested price for us to charge in the Android Market. We explained that this was a reverse Vickrey auction using language similar to Danezis et al.'s experiment [15]:

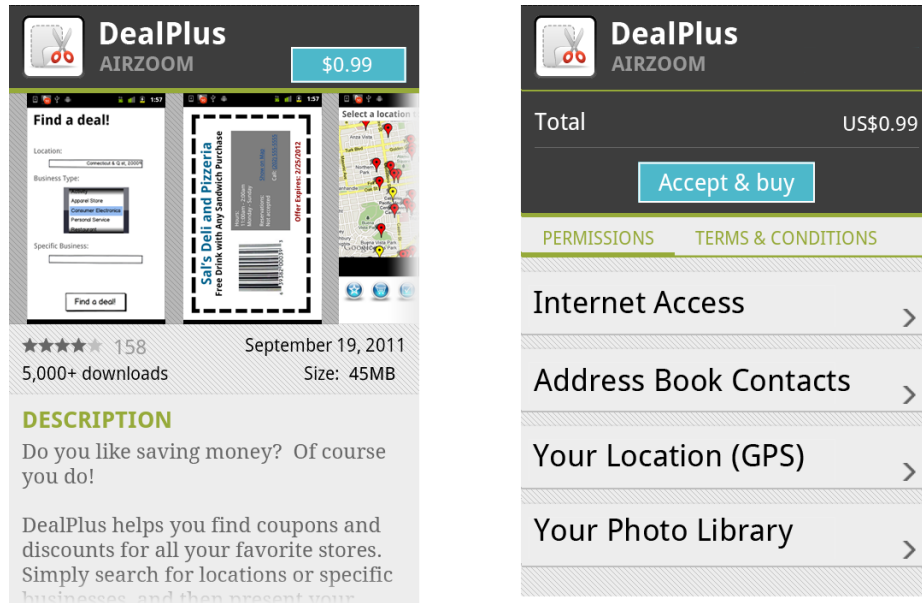


Fig. 3. Screenshot of our fictitious application that participants were asked to value. The permissions requested were randomly assigned from a set of five conditions.

We are recruiting current Android users to participate in a private beta test of this app. If you are selected to participate, you must purchase this app from the Android Market for \$0.99 and install it onto your smartphone. You will be expected to use the app for at least one hour per week over the course of a month.

Each person who is selected to participate in the beta will receive monetary compensation. We are running an auction to select those who will take part. We invite you to submit a bid for the amount of money you require to take part in the beta. Successful bidders will be those who bid the lowest amounts, and each will be paid the amount of compensation demanded by the lowest unsuccessful bidder. (We have yet to decide how many participants we will require.)

The fictitious application screenshots contained a request for up to four possible permissions: Internet access (`INTERNET`), location (`ACCESS_FINE_LOCATION`), address book contacts (`READ_CONTACTS`), and photos from the photo library (`PHOTOS`).⁵ We chose the latter two because they were unrelated to the application’s core functionality, as it was described to participants. We also chose them because our previous research indicated that they would be likely to raise concerns [22]. As before, we included a request for the `ACCESS_FINE_LOCATION` permission because location data has been the subject of extensive privacy research, and we wanted to compare participants’

⁵ This Android permission does not actually exist; no permission is needed to access stored photos. For consistency, we refer to this permission as “PHOTOS.”

concerns about location to their concerns about other permissions. We also included the INTERNET permission because it would have been required for the “DealPlus” application to function as described.

Each user saw one of five sets of permission requests: the INTERNET permission paired with one of the three other permissions, all four permissions at once, or the INTERNET permission alone. We selected these five sets of permissions so that we could isolate the effects of individual permissions and observe the synergistic effect of the combination. (We did not test every combination of the permissions in order to limit the number of experimental conditions.) The variant with only the INTERNET permission alone was our control condition, and the other four permission sets were our experimental conditions. We conducted a pilot study on a separate sample of 320 participants to determine whether participants differentiated between the INTERNET permission and the absence of any permission requests, using a fictitious game entitled “Zombies vs. Wookies.” We did not observe any significant differences attributable to the INTERNET permission, and therefore used it as a control condition.

Android users often misunderstand permission requests because they only read the permission category rather than the actual permission descriptions [22]. To eliminate this confound, we altered the appearance of the permission requests by removing the category information and increasing the font size of each permission description. We also altered the text of the permissions to remove other ambiguities that we observed in our previous study [22]. Since we expect that users understood our modified permission warnings better than the actual permission warnings, participants’ responses to our permission requests are likely an upper bound.

We told participants that they would need to purchase the application from the Android Market for \$0.99. Consequently, they needed to account for this in their bids. We added this caveat for two reasons: to minimize cheating and to limit variance by anchoring participants to an initial point. In our pilot, we included a condition that featured no price, which resulted in extremely divergent suggested prices.⁶ Therefore, to reduce variance, we decided to intentionally anchor participants to a default price of \$0.99. We also hoped to make participants feel more invested in the application by forcing them to consider paying for it, and therefore more willing to pay attention to its details.

Our survey was short to maximize the participation rate. We asked a total of seven questions on four pages. The first two questions solicited participants’ bid amounts and price suggestions. On the second page of the survey, we asked participants to sort several factors that may have influenced their perceptions of the application (“description,” “icon,” “manufacturer,” “permissions requested,” “name of application,” “cost,” and “size of application”) from “most influential” to “least influential.”

The purpose of the survey was to examine how participants would change their valuations of the application based on whether it was collecting information for secondary purposes. The bids were our primary experimental measure, but we also asked an explicit question about advertising. We told participants:

⁶ When we made the price “free,” skewness and kurtosis were 8.36 and 71.03, respectively ($n = 159$). Whereas when we set the price to “\$0.99,” skewness and kurtosis were 1.72 and 5.74 ($n = 163$). This anchoring effect was statistically significant: $U = 10078.5$, $p < 0.0005$, $\mu_{free} = \$2.94$ ($\sigma = 11.09$), $\mu_{\$0.99} = \1.11 ($\sigma = 0.57$).

We're also considering making a free version of the app. It will have targeted advertisements that will be relevant to you, based on how you use the app and your mobile device. For instance, ads may be selected based on:

- How you use the app (e.g., the specific deals you view)*
- Your address book contacts (e.g., the ads viewed by friends who also use the app)*
- Your location (e.g., the businesses you visit while carrying your smartphone)*
- Your photo library (e.g., activities depicted in your photos)*

Each participant saw a subset of the bullets, based on his or her assigned experimental condition. The first bullet was present for every participant because all participants saw the INTERNET permission. The second bullet was displayed for participants who saw the READ_CONTACTS permission, the third bullet was displayed for participants who saw the ACCESS_FINE_LOCATION permission, and the fourth bullet was displayed for participants who saw the PHOTOS permission. Participants who saw all four permission requests also saw all four bullets. We then asked participants to select the version they would be *more* likely to install: the “\$0.99 version with no advertisements” or the “free version with targeted advertisements.”

On the last page of the survey, we asked participants for demographic information, including their Android device make/model, age, and gender.

We collected survey responses during February 2012. After screening out 26 responses due to obvious cheating, non-Android users, and incomplete responses, we were left with 368 responses. These responses corresponded to 139 females (37.8% of 368) and 227 males (61.7%), while two respondents declined to disclose their genders. We observed no statistically significant differences with regard to gender, nor age ($\mu = 29.3$, $\sigma = 8.57$), and therefore did not further analyze demographic factors.

4.2 Results

We considered five hypotheses:

H_{0a} = *Bids will not change based on the permissions*

H_{1a} = *Bids will positively correlate with permission requests*

H_{0b} = *Suggested prices will not change based on the permissions*

H_{1b} = *Suggested prices will negatively correlate with number of permissions*

H_{0c} = *Popularity of the ad-supported version will not change with permissions*

We noticed several clear outliers for the open-ended bids and suggested prices (e.g., one participant bid \$10,000). We compensated for these outliers by excluding every data point above the 95th percentile.⁷ Thus, we analyzed 353 responses.

⁷ This corresponded to bids over \$100 and suggested prices over \$2.99. Prior to removing outliers, the skewness and kurtosis for the bids were 18.65 and 353.15, respectively. After removing outliers, they became 2.15 and 4.10. Regarding the suggested prices, the original skewness and kurtosis were 5.87 and 50.27, but were reduced to 0.63 and 1.79, after removing outliers.

	β	t	Significance
(Constant)		7.435	$p < 0.001$
READ_CONTACTS	0.168	3.104	$p < 0.002$
ACCESS_FINE_LOCATION	0.011	0.215	$p < 0.830$
PHOTOS	-0.027	-0.494	$p < 0.622$
$F_3 = 3.294, p < 0.021$			

Table 3. Regression of participants’ bids as a function of which permissions they were shown.

INTERNET	INTERNET ACCESS_FINE_LOCATION	INTERNET READ_CONTACTS	INTERNET PHOTOS	INTERNET ACCESS_FINE_LOCATION READ_CONTACTS PHOTOS
μ (σ)	μ (σ)	μ (σ)	μ (σ)	μ (σ)
\$13.03 (13.85) $n = 69$	\$16.84 (25.14) $n = 81$	\$24.82 (29.70) $n = 71$	\$15.17 (17.17) $n = 64$	\$21.77 (30.32) $n = 68$

Table 4. The five columns indicate the permission requests that participants saw. Within each column, we include the average bid amounts, the standard deviations, and the number of participants randomly assigned to each condition.

Bids as a Proxy for Privacy Concerns We performed a linear regression between participants’ bids and which of the three permissions they were shown (i.e., PHOTOS, READ_CONTACTS, or ACCESS_FINE_LOCATION). Our results were statistically significant, though not to the degree we had expected: READ_CONTACTS was the only permission that had a statistically significant impact on participants’ bid amounts (see Table 3). The 139 participants who were exposed to this permission demanded significantly more compensation to participate in the beta than the remaining 214 participants ($\mu_0 = \$15.11, \mu_1 = \$23.32, U = 12900.0, p < 0.034$). In summary, we reject H_{0a} , and H_{1a} is supported for contact data.

We performed the same regression with regard to participants’ suggested prices, but it did not yield statistically significant results. Nor did we observe a significant correlation between the suggested prices and participants’ bids. These results suggest that the anchoring effect of including a price in the screenshot overshadowed any effect that the permission request had on participants’ perceptions of what a reasonable price for the application should be; participants reported an average suggested price of \$0.98 ($\sigma = 0.56$), which was similar to the \$0.99 anchor in the screenshot. As a result, we conclude that the suggested prices are an inadequate proxy for participants’ willingness to pay for privacy (i.e., H_{0b} cannot be rejected and H_{1b} cannot be accepted). Thus, our focus remains on participants’ bids.

We performed a Pearson correlation between the number of permissions each condition requested and participants’ bids and found this to be statistically significant ($r = 0.10, p < 0.031$, one-tailed); participants requested more compensation as the application requested more permissions. Table 4 features participants’ average bids, separated into the five between-subjects conditions.

Influential Factors We asked participants to perform a sorting exercise wherein they ranked the factors that influenced their bids. The seven factors that they ranked, in order of decreasing influence, were:

1. Cost ($\mu = 2.39, \sigma = 1.54$)
2. Description ($\mu = 2.64, \sigma = 1.58$)
3. Name ($\mu = 3.36, \sigma = 1.73$)
4. Icon ($\mu = 4.41, \sigma = 1.85$)
5. Permissions Requested ($\mu = 4.64, \sigma = 1.81$)
6. Size of Application ($\mu = 4.83, \sigma = 1.76$)
7. Manufacturer ($\mu = 5.37, \sigma = 1.59$)

With the exception of the permissions requested, the other factors did not appear to change rankings across the five conditions, although we did not perform statistical tests to compare them because we held these factors constant across the five conditions. Regarding the requested permissions, we observed that participants in the control condition (i.e., those who only received the request for Internet access) ranked the requested permissions as the 5th largest factor that influenced their bids, whereas participants in the experimental conditions ranked it as 4th. While this difference was statistically significant ($U = 6227.0, p < 0.015$), participants still ranked permissions as a low fourth.

Willingness to See Targeted Ads We asked participants whether they would prefer a free, advertising-supported version of the application instead of the \$0.99 version that they previously saw. We indicated that the advertisements would be targeted using data from the requested permissions (e.g., participants in conditions requesting the `ACCESS_FINE_LOCATION` permission were led to believe that the advertisements would be targeted by location). Overall, we found that 22.3% of our 368 participants indicated a preference for paying \$0.99 to avoid having advertisements. However, chi-square tests with regard to the specific permissions to which participants were exposed yielded no statistically significant results, and we fail to reject H_{0c} . Thus, this 22.3%'s aversion to advertising does not appear to be based on what data is collected to support targeted advertising. The corollary to this is that 77.7% of our participants would prefer advertisements if it meant saving \$0.99, regardless of the personal data that is collected and used to target those advertisements.

5 Implications

The results of our two experiments indicate that the choice architecture of the application marketplace can have a profound impact on users' privacy decisions regarding their mobile devices. Our results indicate that Android users weigh privacy more heavily when they can easily compare applications' permission requests. In this section, we explore the various factors that compete for users' attention during application selection and installation in the Android Market. We conclude with suggestions for future work to help improve the Android Market to better support users' privacy concerns, though we believe that these suggestions are generalizable to other platforms.

5.1 Decision Factors

In our first experiment (Section 3), we observed that there were three factors that influenced the price/privacy variant that participants ultimately selected: whether they believed the `ACCESS_FINE_LOCATION` permission was relevant to the application’s functionality, the amount of consideration given to permissions in general, and the amount of attention paid to application cost. We performed a linear regression with these factors, which we observed to be highly statistically significant (see Table 5).

	β	t	Significance
(Constant)		16.488	$p < 0.0005$
<code>ACCESS_FINE_LOCATION</code>	-0.140	-3.581	$p < 0.0005$
Permissions	0.315	8.071	$p < 0.0005$
Cost	-0.367	-9.373	$p < 0.0005$
$F_3 = 59.843, p < 0.0005$			

Table 5. Regression of participants’ price selections based on three factors: whether they believed the `ACCESS_FINE_LOCATION` permission was appropriate, their perceived importance of permission requests in general, and their perceived importance of price.

The coefficients in Table 5 indicate that when participants were able to compare similar applications side-by-side, cost was the primary decision factor, but the set of permissions requested were a near second. As a result of these factors, we observed that when the choice architecture allowed them to compare multiple applications of similar functionality, a quarter of participants indicated a willingness to pay a 300% premium for an application that collected the least amount of data. This effect was pronounced when it was clear that the data was extraneous to the application’s core functionality. In this choice architecture, in which participants were able to directly compare permissions between applications, participants who selected the high-privacy variant indicated that the permissions were the primary factor behind their decisions.

The choice architecture that participants encountered in our second experiment (Section 4) did not allow them to view the permissions of multiple applications. Instead, we displayed varying permission requests and asked participants to indicate their willingness to install the application. We observed that their stated willingness was only correlated with the request for one particular permission, `READ_CONTACTS`. Participants were not observably less willing to install the application when it requested access to a user’s photo library (`PHOTOS`) or the location reported by onboard GPS hardware (`ACCESS_FINE_LOCATION`). More importantly, even though only one particular permission triggered privacy concerns, its presence did not force privacy concerns to the forefront of the decision process. Privacy-concerned participants indicated that permissions ranked a meager fourth in terms of the factors they considered when installing an application, as compared to fifth for participants who were not exposed to extraneous permission requests. Our results suggest that because this choice architecture does not allow participants to contextualize the appropriateness of applications’ permission requests, they give less weight to their privacy concerns.

We hypothesize that participants' devaluation of privacy in the second experiment was as a result of bounded rationality. In the first experiment, the side-by-side display of varying permission requests made it obvious to participants that they can directly choose between multiple applications that afford varying levels of privacy, as well as the fact that some applications may not actually *need* some of the requested permissions. Therefore, for those concerned with privacy, they had the option to simply avoid applications that they believed posed a conflict to their privacy preferences. Our data indicate that 25% of our participants exercised this option. In the second experiment, participants were not exposed to multiple sets of permission requests and therefore even if they believed some permissions may be extraneous to the application's functionality, they may have felt they had no choice but to grant them, because a choice was not apparent to them—though they could have reflected this apprehension in their bids. It was only when the permission requests crossed a particular “privacy threshold” that participants demanded additional financial incentives if they were to grant them (i.e., when the `READ_CONTACTS` permission was requested). Thus, they satisfied by accepting extraneous permissions that did not rise to this threshold.

Location, Location, Location? Most previous smartphone privacy research has focused on protecting users' location data. However, our data suggest that users view the location permission as an indicator of desirable functionality rather than an indicator of privacy risk. In our first experiment, we observed that some participants chose an application variant that requested the `ACCESS_FINE_LOCATION` permission because they believed that this permission was relevant to the application's functionality. This indicates that the value proposition offered by sharing location data with applications is generally seen as a net positive: users are more likely to see it as a signal for desirable location-aware features than an inappropriate intrusion upon their personal privacy. In the second experiment, participants did not significantly alter their bids in the presence of the `ACCESS_FINE_LOCATION` permission, nor did it prompt an observable change in participants' decisions in the context of behavioral advertising. This speaks to the acceptance of ubiquitous location-aware applications in the marketplace.

At the same time, we observed that the `RECORD_AUDIO` permission signaled privacy concerns across our first experiment's participants. Examining the Likert data used to report concern levels, we found that the mean concern level over the ability to record audio was significantly higher than concerns over determining location (`RECORD_AUDIO`: $\mu = 4.91$, $\sigma = 2.23$; `ACCESS_FINE_LOCATION`: $\mu = 4.62$, $\sigma = 2.05$; $Z = -2.69$, $p < 0.007$). This contrast did not change based on participants' selections: regardless of price/privacy variant chosen, participants were more concerned with the ability to record audio than with location.

Previous research has suggested that user attention is a finite resource [9, 10]. Therefore, prompting users to approve requests for access to data that does not concern them is likely to increase habituation, which could result in a failure to notice requests that are likely to be more concerning. Our results suggest that a more effective choice architecture needs to account for the relative levels of concern for the varying permissions.

Users and Behavioral Advertising Previously, McDonald and Cranor reported that only 20% of survey respondents would be interested in targeted advertising and that 11% would pay to avoid advertisements altogether [33]. However, in our second experiment, we observed that 77.7% of our 368 participants (95%CI: 73.1% to 81.9%) stated that they were unwilling to pay \$0.99 for an application in order to avoid targeted advertising. This divergence may be due to their study being performed three years prior to ours; privacy attitudes may have changed during the interim as users become accustomed to location-based services. Another possibility is that users are more accepting of behavioral advertising when it is part of a much larger value proposition. In the previous work on acceptance of behavioral advertising, participants were simply asked whether they would prefer it, whereas we specified an exact cost (\$0.99).

We were surprised that participants' support for behavioral advertising did not change as a function of the permissions that were requested in order to target advertisements. This may indicate that participants take an all-or-nothing approach to behavioral advertising: if users are forced to look at advertisements, it may be preferable for those advertisements to be relevant. We did not gather data to explicitly test this.

5.2 Open Problems and Future Work

Our findings suggest that it may be possible to improve the Android choice architecture to better address users' privacy and security concerns. The data that we presented in this paper leads to several logical directions for future work. Our study also suffered from several limitations, which future studies should address.

Privacy Annotations In our first experiment, we observed that privacy-conscious participants were willing to pay a \$1.50 premium over an initial \$0.49 for an application that requested the fewest permissions. In our second experiment, participants satisfied, perhaps because they could not explicitly choose an application based on privacy, even though they may have otherwise been willing to pay a premium for increased privacy. Our results suggest that a choice architecture that allows users to compare permission requests from similar applications side-by-side encourages privacy-preserving behavior. Currently, the Android Market does not support this.

This result, and other prior work on willingness to pay for privacy, suggests that some smartphone users would like the ability to comparison shop for applications based on privacy. We believe that such an architecture is in the best interest of all stakeholders. Many users desire greater privacy features and are willing to pay premiums for them. Such premiums are in the interest of the platform owner, as they increase total marketplace revenue.

We expect to empirically test this in the near future in the laboratory and field by modifying participants' smartphones to support similar search result annotations. Given the limited screen real estate, and that many of the current permissions are un concerning to users, users are likely to comparison shop based on permissions for only a small subset of permissions. We are currently designing icons to represent the permissions that users find most concerning. We expect that if these icons appear alongside search results, users will be more likely to install applications that are aligned with their privacy and security preferences, even if this means paying a premium.

Runtime Permissioning In our second experiment, we observed that participants' privacy-preserving behaviors were much more nuanced than in our first experiment: when participants viewed a single application, their willingness to install it only changed significantly when it requested one particular permission, `READ_CONTACTS`, which previous participants ranked as one of the most concerning. Requests for other permissions, while extraneous to the application's stated functionality, did not concern participants enough for them to significantly alter their valuations of the application; similar extraneous requests in the first experiment resulted in privacy-concerned participants choosing alternate applications with fewer permission requests. This result suggests that smartphone users are likely to install desirable applications regardless of whether or not they request extraneous permissions, even when these requests conflict with their stated privacy preferences—the desire to install the applications outweighs their privacy concerns. This may be due in part to hyperbolic discounting, where the immediate desire for the application results in devaluation of future privacy concerns. Regardless of the exact cause, we believe this result points to another limitation of the choice architecture: when participants are evaluating an application's entire value proposition, privacy concerns are but one aspect; many other factors may overshadow even a privacy-conscientious user's apprehension to disclose personal data.

The current choice architecture under-values privacy because it frames the choice as one between installing the application (and accepting the privacy risks) or not installing the application (and not benefitting from its functionality)—without providing other options at the same time (e.g., other applications that may provide similar functionality at less privacy cost). It is possible that these shortcomings could be addressed by decoupling the decision of whether or not to install an application from the decision of whether or not to grant it a particular permission. Studies are needed to validate this hypothesis. We expect to perform a field study using a modified version of the Android OS to examine whether or not users make the same decisions regarding whether or not to grant applications permissions when those permissions are requested at runtime, when the data is actually needed by the application. Such a choice architecture will need to consider many factors: for example, how often to prompt the user for particular types of data, why the data is being requested, and how to phrase the requests.

Limitations Both of our studies were based on users' stated preferences, rather than observing their actions in the Market. That is, unlike real interactions with the Market, where users are paying actual money and disclosing their actual personal data, our study did not expose them to these costs. At the same time, in our second experiment, we led users to believe that they would be paying for and installing an actual application, and they had no reason to disbelieve us. In fact, our statistically significant results suggest that they weighed these costs in their decisions.

Furthermore, our first experiment displayed four applications side-by-side, which allowed participants to directly compare the full set of permissions requested between all four applications. Due to screen size limitations on most mobile devices, this scenario is a best case for the ability to do side-by-side comparisons. Therefore our survey results likely represent upper bounds.

In both of our experiments, we exposed participants to varying permission requests. However, we only collected data on their hypothetical behaviors for five different permissions. These permissions represent just 4% of the 124 permissions available in the most recent release of Android [5]. It is likely that there are many permissions that users find even more concerning than the ones we examined in this study, as well as many more that users universally find un concerning. A improved choice architecture would need to account for the full spectrum of permissions. Likewise, because of this and the aforementioned limitations stemming from the realism of the tasks, we cannot make generalizations about how much users may be willing to pay to avoid granting particular permissions.

Finally, another limitation of our study was that the results from our first experiment cannot be quantitatively compared with the results from our second experiment, since they used different metrics and were performed over different periods of time (not to mention that they involved completely different methodologies). Instead, we qualitatively compare the results of the two studies to show how changes to the choice architecture can have profound impacts on users' decisions. In future work, we expect to directly address these limitations by conducting laboratory and field experiments wherein participants face real financial and privacy risks.

6 Acknowledgments

The authors would like to thank Jaeyeon Jung and Stuart Schechter for their feedback. This work was supported by Intel, through the ISTC for Secure Computing.

References

1. A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Electronic Commerce Conference (EC '04)*, pages 21–29, New York, NY, 2004. ACM Press. <http://www.heinz.cmu.edu/~acquisti/papers/privacy-gratification.pdf>.
2. A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, pages 24–30, January/February 2005. <http://www.dtc.umn.edu/weis2004/acquisti.pdf>.
3. A. Acquisti, L. John, and G. Loewenstein. What is privacy worth? In *Twenty First Workshop on Information Systems and Economics (WISE)*, 2009.
4. M. Agele, C. Kruegel, E. Kirda, and G. Vigna. Pios: Detecting privacy leaks in ios applications. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2011.
5. Android Developers. Manifest.permission. <http://developer.android.com/reference/android/Manifest.permission.html>, December 2011. Accessed: December 28, 2011.
6. D. Anthony, D. Kotz, and T. Henderson. Privacy in location-aware computing environments. *IEEE Pervasive Computing*, 6(4):64–72, 2007.
7. L. Barkhuus. Privacy in location-based services, concern vs. coolness. In *Workshop on Location System Privacy and Control at MobileHCI '04*, Glasgow, Scotland, 2004.
8. L. Barkhuus and A. Dey. Location-based services for mobile telephony: a study of users' privacy concerns. In *INTERACT'03*, pages 702–712, 2003.

9. A. Beauteament, M. A. Sasse, and M. Wonham. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 workshop on New security paradigms*, NSPW '08, pages 47–58, New York, NY, USA, 2008. ACM.
10. R. Böhme and J. Grossklags. The security cost of cheap user interaction. In *Proceedings of the 2011 workshop on New security paradigms workshop*, NSPW '11, pages 67–82, New York, NY, USA, 2011. ACM.
11. P. H. Chia, Y. Yamamoto, and N. Asokan. Is this App Safe? A Large Scale Study on Application Permissions and Risk Signals. In *World Wide Web Conference*, 2012.
12. N. Christin, S. Egelman, T. Vidas, and J. Grossklags. It's all about the Benjamins: Incentivizing users to ignore security advice. In *Proceedings of IFCA Financial Cryptography'11*, Saint Lucia, Mar. 2011. To appear.
13. S. Consolvo, I. E. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '05, pages 81–90, New York, NY, USA, 2005. ACM.
14. D. Cvrcek, M. Kumpost, V. Matyas, and G. Danezis. A study on the value of location privacy. *Proceedings of the 2006 Workshop on Privacy in an Electronic Society (WPES'06)*, 2006.
15. G. Danezis, S. Lewis, and R. Anderson. How much is location privacy worth? In *Proceedings of the Workshop on the Economics of Information Security Series (WEIS 2005)*, 2005.
16. S. Egelman, J. Tsai, L. F. Cranor, and A. Acquisti. 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*, CHI '09, pages 319–328, New York, NY, USA, 2009. ACM.
17. W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, OSDI'10, pages 1–6, Berkeley, CA, USA, 2010. USENIX Association.
18. W. Enck, D. Ocateau, P. McDaniel, and S. Chaudhuri. A study of android application security. In *Proceedings of the 20th USENIX conference on Security*, SEC'11, pages 21–21, Berkeley, CA, USA, 2011. USENIX Association.
19. W. Enck, M. Ongtang, and P. McDaniel. On lightweight mobile phone application certification. In *Proceedings of the 16th ACM conference on Computer and communications security*, CCS '09, pages 235–245, New York, NY, USA, 2009. ACM.
20. A. P. Felt, S. Egelman, and D. Wagner. I've got 99 problems, but vibration ain't one: A survey of smartphone users' concerns. Technical Report UCB/EECS-2012-70, EECS Department, University of California, Berkeley, May 2012.
21. A. P. Felt, K. Greenwood, and D. Wagner. The effectiveness of application permissions. In *Proceedings of the 2nd USENIX conference on Web application development*, WebApps'11, pages 7–7, Berkeley, CA, USA, 2011. USENIX Association.
22. 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 2012 Symposium on Usable Privacy and Security (SOUPS)*, 2012.
23. C. Foresman. Wireless survey: 91% of americans use cell phones, March 2010. <http://arstechnica.com/telecom/news/2010/03/wireless-survey-91-of-americans-have-cell-phones.ars>.
24. J. Gideon, S. Egelman, L. Cranor, and A. Acquisti. Power Strips, Prophylactics, and Privacy, Oh My! In *Proceedings of the 2006 Symposium on Usable Privacy and Security*, pages 133–144, July 2006.
25. N. Good, R. Dhamija, J. Grossklags, S. Aronovitz, D. Thaw, D. Mulligan, and J. Konstan. Stopping spyware at the gate: A user study of privacy, notice and spyware. In *Proceedings*

- of the *Symposium On Usable Privacy and Security (SOUPS 2005)*, pages 43–52, Pittsburgh, PA, July 2005.
26. 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 (online) of the Sixth Workshop on Economics of Information Security (WEIS)*, Pittsburgh, PA, 2007.
 27. P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall. These aren't the droids you're looking for: retrofitting android to protect data from imperious applications. In *Proceedings of the 18th ACM conference on Computer and communications security, CCS '11*, pages 639–652, New York, NY, USA, 2011. ACM.
 28. B. Huberman, E. Adar, and L. Fine. Valuating privacy. *IEEE Security & Privacy*, 3(5):22–25, September-October 2005.
 29. G. Iachello, I. Smith, S. Consolvo, M. Chen, and G. D. Abowd. Developing privacy guidelines for social location disclosure applications and services. In *Proceedings of the 2005 symposium on Usable privacy and security, SOUPS '05*, pages 65–76, New York, NY, USA, 2005. ACM.
 30. P. G. Kelley, M. Benisch, L. F. Cranor, and N. Sadeh. When are users comfortable sharing locations with advertisers? In *Proceedings of the 2011 annual conference on Human factors in computing systems, CHI '11*, pages 2449–2452, New York, NY, USA, 2011. ACM.
 31. S. Lederer, J. Mankoff, and A. K. Dey. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *CHI '03 extended abstracts on Human factors in computing systems, CHI EA '03*, pages 724–725, New York, NY, USA, 2003. ACM.
 32. A. M. McDonald and L. F. Cranor. Americans' attitudes about internet behavioral advertising practices. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society, WPES '10*, pages 63–72, New York, NY, USA, 2010. ACM.
 33. A. M. McDonald and L. F. Cranor. Beliefs and behaviors: Internet users' understanding of behavioral advertising. In *38th Research Conference on Communication, Information and Internet Policy (Telecommunications Policy Research Conference)*, October 2 2010.
 34. M. Nauman, S. Khan, and X. Zhang. Apex: extending android permission model and enforcement with user-defined runtime constraints. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security, ASIACCS '10*, pages 328–332, New York, NY, USA, 2010. ACM.
 35. P. Pearce, A. P. Felt, G. Nunez, and D. Wagner. Addroid: Privilege separation for applications and advertisers in android. In *Proceedings of the ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2012.
 36. K. Purcell. Half of adult cell phone owners have apps on their phones. Pew Internet & American Life Project, November 2 2011. <http://pewinternet.org/Reports/2011/Apps-update.aspx>.
 37. T. Simonite. Apple ignored warning on address-book access. Technology Review (MIT), February 16 2012. <http://www.technologyreview.com/communications/39746/>.
 38. S. Spiekermann, J. Grossklags, and B. Berendt. E-Privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior. In *Proceedings of EC'01: Third ACM Conference on Electronic Commerce*, pages 38–47, Tampa, Florida, 2001. http://www.sims.berkeley.edu/~jensg/research/eprivacy_acm.html.
 39. R. Thaler and C. Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, New Haven and London, 2008.
 40. TNS. 2009 Study: Consumer Attitudes About Behavioral Targeting. TRUSTe Technical Report, March 2009.
 41. J. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. In *Proceedings of the 2007 Workshop on the Economics of Information Security (WEIS'07)*, Pittsburgh, PA, USA, 2007.

42. J. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The impact of privacy indicators on search engine browsing patterns. *Information Systems Research*, 22(2):254–268, June 2011. <http://www.guanotronic.com/~serge/papers/isr10.pdf>.
43. A. F. Westin. *E-Commerce & Privacy: What Net Users Want*. Privacy & American Business, Hackensack, NJ, 1998. <http://www.pwcglobal.com/gx/eng/svcs/privacy/images/E-Commerce.pdf>.
44. J. Wiese, P. G. Kelley, L. F. Cranor, L. Dabbish, J. I. Hong, and J. Zimmerman. Are you close with me? are you nearby?: investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th international conference on Ubiquitous computing*, UbiComp '11, pages 197–206, New York, NY, USA, 2011. ACM.
45. K. Zickuhr. Generations and their gadgets. Pew Internet & American Life Project, February 2011. <http://pewinternet.org/Reports/2011/Generations-and-gadgets/Report/Cell-phones.aspx>.