The Impact of Timing on the Salience of Smartphone App Privacy Notices

Rebecca Balebako[1], Florian Schaub[2], Idris Adjerid[3],
Alessandro Acquisti[2],Lorrie Faith Cranor[2]

[1]RAND Corporation
Pittsburgh, PA, USA
balebako@rand.org

[2]Carnegie Mellon University
Pittsburgh, PA, USA
{fschaub, lorrie }@cmu.edu
acquisti@andrew.cmu.edu

[3]University of Notre Dame
Notre Dame, IN, USA
Idris.Adjerid.1@nd.edu

ABSTRACT

In a series of experiments, we examined how the timing impacts the salience of smartphone app privacy notices. In a web survey and a field experiment, we isolated different timing conditions for displaying privacy notices: in the app store, when an app is started, during app use, and after app use. Participants installed and played a history quiz app, either virtually or on their phone. After a distraction or delay they were asked to recall the privacy notice's content. Recall was used as a proxy for the attention paid to and salience of the notice. Showing the notice during app use significantly increased recall rates over showing it in the app store. In a follow-up web survey, we tested alternative app store notices, which improved recall but did not perform as well as notices shown during app use. The results suggest that even if a notice contains information users care about, it is unlikely to be recalled if only shown in the app store.

Categories and Subject Descriptors

H.5.2 [**Information Systems Applications**]: Information Interfaces And PresentationUser Interfaces

General Terms

Mobile, Privacy, Privacy Decision Making

1. INTRODUCTION

Smartphone users are concerned about the privacy intrusions that may result from sharing information with apps [23]. Seemingly harmless apps (e.g. games, flashlight apps etc.) can ask for extensive permissions or collect sensitive information. Privacy concerns depend not only on the type of data collected by apps and but also with whom that information is shared. Smartphone users are tasked with managing privacy risks and may do so by selecting between apps at install time, uninstalling existing apps [23], or managing the privacy settings of smartphone apps.

SPSM'15, October 12, 2015, Denver, Colorado, USA.
ACM 978-1-4503-3819-6/15/10.
DOI: http://dx.doi.org/10.1145/2808117.2808119.

Currently, users of major smartphone platforms are informed about data collection practices through permissions dialogs or long privacy policies. Despite efforts to improve the content of privacy notices through standardization [29] or privacy metrics [19,25,27], it is not clear *when* privacy notices should be shown to users. While it is well-known that people often ignore notices such as computer security dialogs or End User License Agreements, which may be shown at install-time [11], it is less understood whether there are optimal times to show smartphone privacy notices to maximize attention and recall. Recall is relevant in order to support users' retention of accurate mental models of an app's privacy practices.

This paper assesses how timing can impact the recognition-based recall of an app privacy notice's content. Recall is an established measure of risk warning effectiveness [6] that provides an indication of the lasting salience of the notice, rather than just measuring the attention paid to the notice in the short term. We investigated whether privacy notices can be made more salient, holding content constant and focusing on the timing of notices, because even the best notice will be ineffective if shown at an inopportune time. We assess how small differences in the timing of the privacy notices (in some cases, only a few seconds difference) impacted memory twenty-four hours after users installed and used the mobile app.

We conducted a web survey and a field experiment to compare whether participants better recognized content of the privacy notice shown before installation, or before, during, or after app use. A follow-up web survey investigated app store notices in particular. We make the following contributions in this paper:

- We find that timing matters for smartphone privacy notices. Even if a notice is carefully designed to show information users care about, the notice's content is unlikely to be recognized when displayed in the app store. We show these effects in multiple settings - namely two web surveys and in a field experiment.
- We provide recommendations on how to integrate privacy notices into apps for improved recall. Since notices in the app store had low recognition-based recall rates, our results suggest that a notice should be shown at the beginning of app use or during app use, for example when the requested resource is accessed.
- Since there are other benefits to providing notices in the app store, we offer design guidelines for improving privacy notices shown in the app store based on a follow-up web survey.

Our web surveys and field study produced complementary results. The initial web survey sampled a diverse population and

found that users cared about the notices. The field experiment, conducted for its higher ecological validity, provided evidence of replicability and robustness. It also raised new questions about the app store condition, in particular about whether the size of the notice in the app store was comparable to the other conditions. Examining this question required modifications to the app store not possible in a field experiment. Thus we ran a second web survey, modifying a screenshot of the app store to test these changes.

2. BACKGROUND AND RELATED WORK

First, we discuss the role of privacy and security notices in human decision making. Then, we provide a brief overview of the current state of privacy notices on smartphone platforms. Finally, we discuss user studies examining novel smartphone privacy notices.

2.1 Privacy Notices and Decision Making

Informing users about privacy and security issues is an important step in involving humans in security and privacy decisions [3, 12]. However, it can be difficult to get users' attention and inform them for several reasons. One reason is that privacy and security management are usually not the user's primary task. Another reason is that the repercussions of decisions are not immediate (an unintended disclosure or misuse of information may occur days or weeks after the notice was shown) and users may not associate consequences with a privacy or security notice.

The Computer Human Information Processing (C-HIP) framework discusses the stages in which humans notice and process warnings [32], and has been expanded for secure systems [12]. The framework models the stages and variables a human may go through when presented with a security warning or notice, such as switching attention to the notice and encoding the information. It emphasizes that many aspects of notifications must work in concert to influence the behavior.

Salience of privacy notices and whether users switch attention to the notice is important, but difficult to measure. Memory has been used as a proxy for salience, because without salience one cannot remember. A meta-analysis of consumer warnings across a variety of products found that recall of the notice was one of five dimensions – including also attention, judgment of risks, comprehension, and compliance with the warning – that define the warnings' effectiveness [6]. Furthermore, as described by Argo and Main, "Consumers' recall of warning information influences their decisions of whether and how to use a product correctly" [6]. In the risk warning literature, the term recall (memory without cues) was not distinguished from recognition (such as picking the correct answer from a list).

Smartphone users are tasked with the on-going management of privacy and data sharing over the lifetime of their phones. While users may make privacy decisions when they are first shown a privacy notice, this is not the only opportunity to make privacy decisions. Users may be surprised by an app's data practices, becoming aware of or suspecting undesired data practices. They may then assert control in at least two additional cases: 1) by changing privacy settings in system options; 2) by deleting or stopping the use of installed apps [23].

Timing is one dimension of privacy notice design [30]. Experimental work in other contexts, such as web shopping or software installation, has demonstrated that poor timing of privacy or security notices may hamper attention to or ability to act on a privacy notice. People who are engrossed in the installation process may fail to pay attention to an install-time notice [20]. Furthermore, introducing only a 15-second delay between the presentation of privacy notices and privacy relevant choices can be enough to render notices ineffective at driving user behavior [4]. The timing of privacy indicators can impact online shoppers' willingness to pay a premium at websites with better privacy policies. A study found that users were more likely to act on web privacy information when it was shown in search results than when the same information was displayed with website content [14]. While these previous works indicate that timing is important in privacy notices, no one has, to our knowledge, investigated the role timing plays in the context of smartphone privacy notices.

2.2 Smartphone App Privacy Notices

Smartphones differ from computers in that they have several types of sensors and input mechanisms for collecting data passively. Furthermore, the owner tends to carry the smartphone with them at all times. They may choose to install a myriad of apps created by various developers. The privacy implications are that smartphone users can be exposed to sharing more data, from more locations, with many apps and app developers than a traditional desktop or laptop user.

The current major smartphone platforms — Google Android and Apple iOS — automatically display permission notices for apps based on an app's required data collection. However, they provide the notices at different times. At the time of writing, Android users are shown a list of requested permissions while the app is being installed, i.e., after the user has chosen to install the app. In contrast, iOS shows a dialog during app use, the first time a permission is requested by an app. This is also referred to as a "just-in-time" notification [30]. While these permission request include an option to control sharing, our privacy notices do not, in an effort to isolate the impact of timing.

U.S. policy-makers recognize that non-standardized permission requests may be confusing to users and may provide inadequate privacy information. The National Telecommunications and Information Administration (NTIA) developed a code of conduct for standardized short-form privacy notices for smartphone apps [29]. This notice includes a list of data elements and third-party entities about which users should be informed, and which should be shown in addition to the permissions requests. Industry associations have backed the code and have developed several examples of how these privacy notices could be implemented in the app store or within an app [21]. While a previous study showed that study participants did not understand all the terms used in the NTIA code [9], we chose to use a privacy notice based on this standard for three reasons: it was developed by a multi-stakeholder group, including consumer privacy advocates who considered consumers' needs and desires for a privacy notice, and it is backed by industry and it is not biased towards any platform.

2.3 Smartphone Privacy Notice Studies

Privacy information may influence users' choice of apps at install time, if the information is provided in a clear manner. Previous research has examined app store privacy notices designed to facilitate comparison of apps on the basis of privacy or sharing risks [18,24]. These papers used online studies with modified screenshots of app stores to measure users' reactions to privacy notices. Our web surveys use this method as well.

Previous research has examined users' immediate reactions to privacy notices when comparison shopping. When shown a 'privacy checklist' in the Google Play Store, users would select the app requesting fewer permissions [24]. Additionally, when asked to compare similar apps with different permission requests, users demonstrated that they were willing to pay more for apps requir-

ing fewer permissions [13]. A 'risk' score generated by examining an app's permissions was shown to be effective in helping users choose an app that requested fewer permissions [18]. However, users may follow different paths when selecting an app, and may not be comparison shopping between apps in the app store. They may select an app based on peer recommendations, and may forgo comparison shopping.

Previous studies examined just-in-time permission notices. Participants in a lab study were shown just-in-time notifications after a few minutes of playing Angry Birds and Toss It [7]. Participants in a field study were shown location notifications while using the same apps they normally used on their phones [5, 17]. In these studies, and in contrast to our study, the notices were shown multiple times if the permission was requested frequently. These studies found that users appreciated the notices, and the latter found that users did take privacy-protective actions, such as uninstalling apps, when they saw notices during app use.

There are several practices that are known to improve the design of privacy notices, which we take into account when providing recommendations. The design of notices and warnings needs to conserve user attention by easing decision-making and avoiding interruptions [16]. Privacy decision making may be overwhelming: the cognitive costs associated with considering potential ramifications of sharing data may hamper decision making [1, 2]. In addition, when notices are shown too frequently, users may become habituated. Habituation may lead to users disregarding warnings, often without reading or comprehending the notice [10]. To reduce habituation from app permission notices, Felt et al. identified a method to determine which permission requests should be emphasized [15].

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

3. METHODS

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

We used two methods – a web survey and a field experiment – in combination as they each have specific benefits. Web surveys allow quick access to a large, relatively diverse participant pool. Our field experiment aimed to be ecologically valid, as users were installing our app on their own phones on their own time, with the vagaries and distractions that may occur naturally. The follow-up web survey was conducted after the first web survey and the field experiment, and allowed us to examine some variations on the app store notice condition that are not available in the real app store. The web surveys and field experiments used the same questions and steps.

We measured participants' recall of the privacy notice, and their self-reported desire to see and to remember the notice within the context of the employed app. All participants completed five steps: 1) consent form and demographic questions, 2) install and play the app, 3) experience a distractor or delay, 4) answer recall questions, and 5) evaluate notice.

Web survey participants completed all five steps in a web browser using a survey powered by SurveyGizmo. In contrast, field study participants completed only steps 1 and 5 in a web browser survey; the rest of the study was performed on their phones. To transition cleanly from step 1 to 2, field study participants were redirected to the study app's page in the app store after completing the demographic questions. They were not asked to comparison shop or search for the app. Field experiment participants were sent an email after the delay (step 3) with a link to the on-line questionnaire that included steps 4 and 5.

In both web surveys, participants completed all steps in a web browser. Like several previous studies of mobile privacy notices [22, 24, 31], we displayed screenshots of the app store in a web survey. In particular, we displayed a "virtual smartphone" in the web survey, which was simply a web browser iFrame that had similar dimensions to a smartphone and a thick black border to distinguish it from the web page. Within the iFrame, web survey participants clicked on screenshots of the app store's page to "virtually install" the app, which led to mock-ups of the install process. After "installing," they played the app itself in the same "virtual smartphone." They completed the app before advancing to the next pages of the web survey to get to steps 3, 4, and 5 of the study.

All participants reviewed the instructions for participation in the first step. In the final step in which participants were asked to evaluate the notice, they were shown an image of the notice and debriefed on the purpose of the study.

3.0.1 Quiz app

We designed and deployed a simple quiz app. Our two objectives in selecting the app content were to create an app that is both entertaining and distracting, and that could be completed in a few minutes. Furthermore, we wanted the privacy notice to be the users' source of privacy information about the app, so we refrained from app design that clearly collected sensitive information or impinged on user privacy. Therefore, we developed an innocuous-seeming history quiz that asked eleven questions about the inventions of less famous inventors (see Figure 1).

Before beginning the quiz, the app showed two consecutive screens: first a paragraph of instructions, and second a page to enter an email address (field experiment) or code (web survey) to link participants to their consent form. After answering the eleven quiz questions, participants saw their score. The app was developed in HTML and JavaScript using PhoneGap.[1]

We created two similar Android app store entries; one entry showed the privacy notice (shown in Figure 2) and the other did not. In the first web survey, the privacy notice was the right-most image (as seen in Figure 3). For the field experiment, we hypothesized that the notice might be more salient as the first (leftmost) image. We did not find that this improved recall.

At the time of the field experiment, our apps were not rated, and had neither ratings nor comments from users. The description of the apps in the store stated that it was part of a research project and included a link to the consent form approved by our IRB. In the "Developer's Website" section of the app store, we included a link to our website detailing the steps of the experiment and including a link to the consent form. The privacy policy link in the store pointed to the image of the privacy notice shown in Figure 2.

3.0.2 Privacy Notice Design and Verification

We had two goals when creating the notice; first to create a realistic notice and second to create a notice that contained privacy information that people care about and would want to recall. To create a realistic privacy notice, our design was based on the aforementioned code of conduct for standardized privacy notices [29]. More

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

Figure 1: A quiz question from employed app.

specifically, we employed Private Parts,[2] an open source implementation of this standard, which we modified to match our app's color scheme (see Figure 2). Our notice informed users that the app collected "Browser History" and that this data would be shared with "Ad Networks." We selected only one data and one entity to avoid confounds or information overload for the participants. To insure that the notice was the same across all conditions, including in the app store where we could not show an interactive notice, we moved the explanatory text – which appears only after clicking a Private Parts element – next to the privacy icons.

Figure 2: The privacy notice.

We designed our notice based on previous research. We speculated that users may be inured to collection of location, as location data is collected by many apps [5, 8, 26]. On the other hand, previous research [27] has found that users are very uncomfortable with apps taking contact information without a clear purpose. We selected the middle ground of user concern: "Browser History." We further ascertained the appropriateness of these choices for the notice in our web survey.

3.1 Timing Conditions

We used a between-subjects design in which participants were randomly assigned by our survey software to one of 5 timing conditions. The conditions varied based on the moment in time at which the privacy notice was shown. The first web survey and the field experiment had five conditions, which represent privacy notice timings that occur in existing apps and platforms. The app store condition varied slightly from the other conditions: the notice in the app store did not occupy the full screen, whereas all the other notices were shown full screen. The timing conditions were:

Not Shown. The privacy notice was not shown to the app user. This is the no-treatment condition.

App Store. The notice is displayed as a screenshot in the app store (see Figure 3), similar to previous work on showing privacy indicators in the app store [18, 24]. This is the only way that Android allows privacy notices to be displayed at install time.[3]

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

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

After Use. The notice is shown after completing the last question of the app quiz. This would be the timing of a privacy notice shown to summarize data sharing and collection, after the app has been played, similar to the timing of a summary notice [7].

The 'App Store', 'During Use', and 'After Use' conditions have all been used in research or major platforms. For example, the 'during use' condition shows the privacy notice while the app is in use; this is similar to the iOS platform which displays permissions requests while the app is in use (and blocks data collection until permission is granted). The 'before use' resembles notices we have seen in some apps that describe data collection practices, such as the use of analytics, that are not covered in the permissions dialogs. The 'after use' condition resembles summary statistics shown in some research efforts [5, 7]. In the follow-up web survey, we investigated variations on the app store timing condition. The two additional conditions used in the second web survey are introduced in the Section 6.

In the 'Not Shown' condition, participants did not see a notice. All participants were given the option to select "I don't remember" for the recall questions, but some participants in the no-treatment condition still guessed, allowing us to establish a baseline of how many participants could correctly answer the recall questions by chance. In questions in which participants were asked to evaluate the notice, they were given an option of "Not Applicable/Don't Remember" so that they could honestly respond even when they had not seen the notice or didn't remember it.

Our two web surveys and field experiment were between-subjects experiments, and participants were randomly assigned to one con-

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

[3] Just before publication of this paper, Google announced that in upcoming versions of Android permission notices would be shown at run-time.

Figure 3: App store with the privacy notice.

dition. The app, privacy notice, recruitment, and all associated materials were identical across conditions.

3.2 Questions Following App Usage

In both web surveys and the field experiment, participants completed the same exit survey at the end, which is shown in Appendix A. The questions allowed us to evaluate their recall of the notice and app, and to evaluate the notice. The questions used to measure recall of the privacy notice were, "With whom does the app share data?" and "What information was collected by the app?" The questions were multiple-choice, with six possible answers, including "I don't remember." Participants were also asked two multiple-choice questions about the how well they remembered other aspects of the app, such as the contents of the quiz questions, the color of the app background, and whether they remembered seeing the privacy notice.

At the end, participants were shown the privacy notice again and were asked to evaluate it. These questions were used to measure whether participants perceived the privacy notice's content as important and whether they wanted to remember it. Participants were told, "This is the privacy notification for the app. Please note, we did *not* collect this information, but please imagine your reaction if this really occurred on your phone." Six 5-point Likert-scale questions about the notice included positively-biased questions such as, "The privacy notification gave me information I care about" and negatively-biased questions such as "This notification could be improved so I understand it better." Four 5-point Likert-scale questions were used to evaluate participants' opinions of the timing of the notice. The questions included whether the timing was disruptive, unexpected, allowed them to make decisions, and whether they could pay attention to the notice. These questions also included a "Not Applicable/Don't Remember" option, as participants in the no-treatment group did not see the notice and therefore were not positioned to evaluate the timing.

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

Table 1: Number of participants in web survey, and correct recall of both the data and entity described in privacy notice, by condition. Values significantly different from "not shown" are marked with * (Mann-Whitney U with Bonferroni correction)

4. WEB SURVEY RESULTS

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

4.1 Web Survey Participants

Web survey participants were paid $1.01 and were recruited via Amazon MTurk.[4] To ensure quality of MTurk workers we allowed only MTurkers with an acceptance rate of >=89%, we required completion in 30 minutes, and we included two attention check questions regarding the instructions. Two hundred and seventy-seven U.S. participants completed the survey. Participants completed the survey in a median of 9.08 minutes (range 2.82-29.6). The participant group was diverse. Nearly half of the participants (49%) were female (1 declined to state gender). Almost half (48%) had a bachelors degree or graduate degree. While the ages ranged from 18 to 69 years, the median age was 29 years. Forty-five out of 50 U.S. states were represented. Most of our participants owned and used a smartphone (95%), although we did not recruit for smartphone owners, and specifically stated that owning a smartphone was not a prerequisite for the web survey. There were no significant differences in the following demographics across timing conditions: age (ANOVA F=1.67 p=16), gender ($\chi^2(8)$=12.4, p=.135), and smartphone type owned ($\chi^2(20)$=19, p=.524) respectively.

The timing conditions were randomly assigned by SurveyGizmo, which initially distributed participants quite unevenly. There were between 39 and 67 participants who completed each condition, as seen in Table 1.

4.2 Web Survey Analysis

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

4.2.1 Recall of the Privacy Notice

Most participants did not feel confident in their recall of the privacy notice when asked, "Do you remember seeing the privacy notice?" after the distraction. Only 36.5% responded either, "I remember most of it," or "Yes, I remember it well," while the remainder responded that they did not remember it at all, or only remembered it vaguely.

Only 24.5% of participants correctly remembered both the data (*browser history*) and entity (*ad networks*) shown in the privacy notice. More participants remembered the data (40.8%) than the entity (31.4%). Both were recalled better than simple chance of selecting one of the six options (16.6%). Self-reports of remember-

[4] `www.mturk.com`

ing the notice did positively correlate with the ability to correctly identify the elements on the notice (r_Φ=.546 p=.001 for data and r_Φ=.602 p=.001 for entity). There was a positive correlation between correct recall of the data and the entity (r_Φ=.515 p=.001). We used an ordinal variable "RecallCorrect" with three levels: 1) did not remember any part of notice, 2) remembered at least one part of notice, 3) remembered both data and entity from the notice correctly.

When the notice was shown before, during, or after app use, participants remembered it more accurately than when shown in the app store, see Table 1. A Kruskal Wallis test revealed a significant effect of timing condition on RecallCorrect (KW $\chi^2(4)$=70.2, p=0.001).[5] Post-hoc tests (Mann-Whitney U with Bonferroni correction) showed significant differences between 'not shown' and each of the three app use conditions (before/during/after), as seen in Table 2. The three app use conditions were not significantly different from each other. Participants who saw the notice in the app store were less likely to remember the notice than those who saw it during app use. This indicates that notices shown at the time of app use are most beneficial for retention and later recall of the notice.

condition	not shown		app store		before use		during use	
	r	p	r	p	r	p	r	p
app store	0.1	.27						
before use	0.5	.001*	0.3	.005*				
during use	0.6	.001*	0.4	.001*	0	1.0		
after use	0.5	.001*	0.2	.006*	0	1.0	0.04	0.93

Table 2: Web survey: r (effect size) and p-values of pairwise comparisons on recallCorrect using Mann-Whitney U with Bonferroni correction. Significant results marked with * α=.05.

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

In the web survey, the distractor was a set of IUIPC web privacy concern questions [28]. Responses to the IUIPC scale [28] in the categories "Control" and "Collection" were not correlated with RecallCorrect, although "Awareness" was a weak predictor (one-way ANOVA, F=5.97, p=.015). Participants recall was not affected by the following demographics: age (ANOVA, F=.38, p=.54), education (KW $\chi^2(2)$=.267, p=.875), gender (KW $\chi^2(2)$=1.06, p=.590), and owning a smartphone (KW $\chi^2(2)$=2.31, p=.315). Overall, previous preferences and timing are the main predictors of whether the participants remember the privacy notice.

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

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

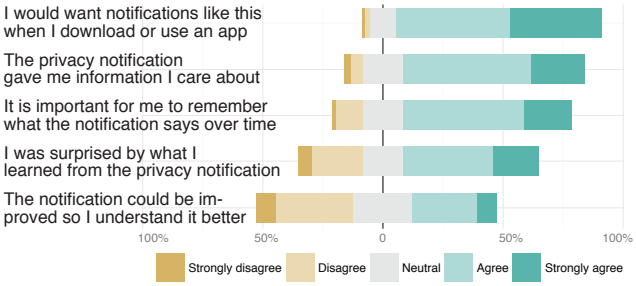

Figure 4: Web survey participants' responses when asked to review the notice after completing the survey. Participants want the notification and want to remember it. No sig. diff. between conditions (KW).

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

4.2.2 *Evaluation of the Privacy Notice*

We would not expect participants to remember a notice unless they care about it and would want to remember it. Figure 4 shows the results of the Likert-scale questions used to evaluate the privacy notice. Our results validate the notice's relevance, showing that, overall, participants wanted to remember it and felt it had information they cared about. We note that liking the notice does not imply that they liked the data collection described in the notice. The responses to these Likert-scale questions did not significantly depend on the timing condition (KW test with Bonferroni correction, α=.01).

We also evaluated participants' reactions to the timing of the notice. We did not remind participants what timing condition they were in. We included a "Don't Remember/Not Applicable" option. We show this response option in Figure 5, as it is deepens our understanding of how many participants in specific conditions recognized that they did not remember a notice.

The timing condition significantly impacted participants' responses to two questions about the timing of the notice: "The privacy notification was shown at a time when I could make decisions about whether to allow the data collection" (KW $\chi^2(4)$=32.4, p=.001) and "The privacy notification occurred at an unexpected time" (KW $\chi^2(4)$=44.2, p=.001) (see Figure 5). Participants in the "after use" condition reacted negatively to the timing of the notice, and said more frequently that the timing was unexpected and that they could not make decisions about the data collection.

5. FIELD EXPERIMENT RESULTS

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

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

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

5.1 Field Experiment Participants

We recruited 126 participants from three university participant pools: Phone-Lab at SUNY Buffalo[6] (n=29), Notre Dame University[7] (n=37), and CBDR at Carnegie Mellon University [8] (n=42). We also posted ads on craigslist and reddit, which yielded 18 additional participants. Our participant pool skewed young. While the range of ages was 18–55, 80% of our participants were 30 or younger (median=23.5). Our participants were well educated, as 57% had a bachelors or graduate degree; 46.8% were female, and the rest were male. Participants were based in 24 different U.S. states. There were no significant differences between conditions in age (ANOVA, F=1.67, p=.16), gender ($\chi^2(4)$=.716, p=.949), U.S. state ($\chi^2(100)$=113, p=.171), or education level ($\chi^2(12)$=14.1, p =.297). The field experiment was conducted only on Android. The dropout rate between conditions was not significant ($\chi^2(8)$=3.02, p =.933). Compared to the online survey, participants were slightly younger, all used Android phones, and resided in fewer US states, but otherwise the participant groups were similar. Table 3 shows the number of participants in each condition.

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

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

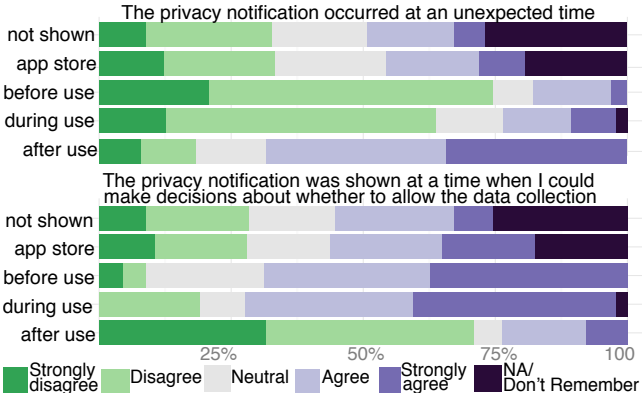

Figure 5: Web survey responses about timing of privacy notice. Participants in after-use condition were more negative about timing.

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

Table 3: Number of participants in field experiment, and correct recall of notice by condition. Values significantly different from "not shown" are marked with * (Mann-Whitney U with Bonferroni correction).

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

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

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

5.2 Field Experiment Analysis

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

5.2.1 Recall of the Privacy Notice

Participants did not feel confident that they remembered seeing the privacy notice. When asked, "Do you remember seeing the privacy notice?" 54% of all participants said they only remembered it "vaguely," while 21% said they did not remember it at all. Only 5% said they remembered the notice well. Unlike the web survey, self-reported response of remembering the notice did not correlate with

the ability to correctly identify the elements on the notice (r_Φ=.207 p=.021 for data and r_Φ=.121 p=.184 for entity).

While the recall rate is lower than that of the web survey – likely due to the longer delay – the trends are similar. More participants remembered the data than the entity. Overall, just over one-third (37.3%) of field experiment participants correctly identified that the privacy notice said data was shared with the entity *Ad Networks*. A smaller percentage (26.2%) correctly identified that the privacy notice stated that it would collect *Browser History* data. About one-fifth of participants correctly remembered both aspects of the privacy notice (21.4%). Both of these percentages are better than if multiple choice answers had been selected randomly (16.6%). Correct recall of the two aspects of the privacy notice was positively correlated (r_Φ=.548, p=.001); for example, 81% of participants who remembered the entity type also remembered the data.

The timing condition was a significant predictor of recall, with all conditions during app use yielding better recall rates than the app store or no-treatment conditions. The percentage of participants who correctly recalled both aspects of the notice is shown in Table 3. Of the participants who saw the notice, those in the app store condition were the least likely to remember it. Overall, timing had a significant impact on "RecallCorrect" (KW $\chi^2(4)$=24.1, p=.001). Post-hoc tests (Mann-Whitney U with Bonferroni correction) showed significant differences between "not shown" and all three conditions during app use (before, during, and after use), but the difference between "app store" and "not shown" was not significant. Differences between the three within-app conditions were also not statistically significant. These pairwise comparisons are shown in Table 4.

A statistically significant difference in RecallCorrect exists between participants who self-reported to read privacy policies frequently versus those that did not (KW $\chi^2(3)$=16.1, p=.001). Of participants who indicated that they read privacy policies 'Always' or 'Sometimes,' 30% correctly recalled the notice, while of those who selected 'Never' or 'Rarely,' only 11% correctly recalled the notice. This indicates that pre-existing preferences and behaviors impact the user's ability to remember the notice.

One additional variable impacted RecallCorrect. The self-reported frequency of installing apps on their phone impacted RecallCorrect (KW $\chi^2(3)$=11.2, p=.010). The more frequently they installed apps (e.g. 'Often' or 'Daily'), the more likely they were to correctly remember the notice.

There were no statistically significant differences between demographic groups in RecallCorrect. Region (KW $\chi^2(25)$=28.6, p =.282), gender (KW $\chi^2(1)$=2.09, p=.148), age (ANOVA, F=.01, p=.91), or education level (KW $\chi^2(8)$=5.67, p=.683) did not affect RecallCorrect. As in the web survey, previous preferences and timing are the main indicators of whether the participants remember the privacy notice. Thus, the results of the web survey also hold in the more ecologically valid setting of the field study.

Field study participants had a much better recall of other aspects of the app than the privacy notice. The majority correctly identified the background color (74%), and were able to identify two inventors described in the quiz in two questions (86.6% and 57.1%). Recall of these aspects of the app did not correlate to the timing condition, indicating that the ability to recall the app in general was evenly distributed between conditions. In general, participants' recall of the privacy notice was not correlated with their memory of the other aspects of the app. That is, correctly identifying the people in the app quiz or the background color did not correlate to correctly remembering the data or entity in the privacy notice (χ^2-test, corrected α=.017). This further suggests that any ability to remember the notice, or not, was not simply a matter of remembering the app overall, but isolates the effect of timing as an impact. This also indicates that privacy is treated as a secondary task as expected, as the primary task (history quiz) was better retained.

condition	not shown		app store		before use		during use	
	r	p	r	p	r	p	r	p
app store	0	1.0						
before use	0.3	.001*	0.2	0.05*				
during use	0.3	.004*	0.1	0.13	0	1.0		
after use	0.2	.006*	0.1	0.12	0	1.0	0	1.0

Table 4: Field experiment: r (effect size) and p-values of pairwise comparisons on recallCorrect using Mann-Whitney U with Bonferroni correction. Significant results marked with * (α=.05)

.

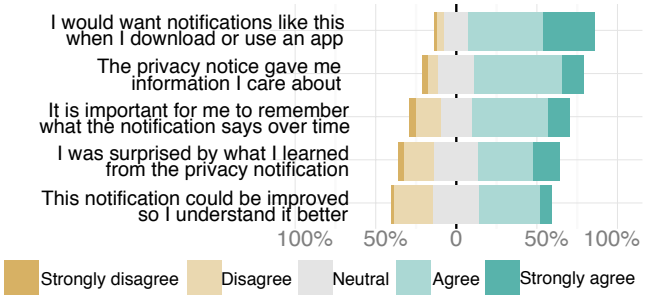

Figure 6: Field experiment participants want the notification and want to remember it.

5.2.2 *Evaluation of Privacy Notice*

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

Overall, participants stated that they wanted to see the notice when downloading or using an app (78%), wanted to remember the information in the notice a day later (60%), and cared about the information shown in the notice (66%). Half of the participants found the content of the notice surprising (50%). The results of the questions are shown in Figure 6.

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

6. FOLLOW-UP WEB SURVEY ON APP STORE NOTICES

The first web survey and field experiment indicated that when the notice was shown in the app store participants had low rates of recall. We used an app store design that matched what had been proposed by the multi-stakeholder group that developed the notice code of conduct [29]. This design did not require any changes to the app store itself as the privacy notice could be inserted as a screenshot. However, the notice was not displayed prominently. In

the follow-up web survey, we evaluated whether the notice in the app store was less effective due to the smaller size and distractions (such as other elements describing the app). We find that a larger notice, and a notice that users must click on perform better than a screenshot in the app store, but not as well as the during-use timing.

6.1 Follow-up Web Survey Participants

Our follow-up web survey used the same method as the first web survey; the app store install process was simulated through a series of clickable screenshots displayed in a iFrame. Three of the conditions (Not Shown, App Store, and During Use) were the same as the previous web survey and field experiment. We added two new conditions designed to show the notice more prominently. The two new conditions were:

App Store Popup. The privacy notice was shown to the user as a pop-up after the permission dialog popup. The app store was greyed-out, and the privacy notice dominated the screen (see Figure 8).

App Store Big. The notice is in the same location as the screenshots in the app store (see Figure 7), but the image is as wide as the store, and replaces other screenshots.

Figure 7: App store with the big privacy notice shown in place of screenshots.

Web survey participants were paid $1.01 and were recruited via Amazon MTurk. The median age of the 326 participants was 31 years (range 19–69). Forty-six percent of participants were female; 3 participants opted not to state their gender. Almost half (49%) had a bachelors degree or graduate degree. Forty-four out of 50 U.S. states were represented. Most of our participants owned and used a smartphone (94%). There were no significant differences between timing conditions and the following demographics: age (ANOVA F=1.09 p=.36), gender ($\chi^2(8)$=8.11, p=.423), and smartphone type owned ($\chi^2(12)$=18.7, p=.096) respectively). Participants completed the survey in a median of 8.71 minutes (range 2.68-27.8).

Figure 8: Privacy notice as a popup displayed after the Android permission screen.

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

Table 5: Number of participants per condition in follow-up web survey, and correct recall of notice by condition. Values significantly different from "not shown" are marked with * (Mann-Whitney U with Bonferroni correction).

The timing conditions were randomly assigned, and there were between 52 and 84 participants in each condition. Table 5 shows the number of participants in each condition.

6.2 Follow-up Web Survey: Notice Recall

The follow-up web survey found that the app store notice was recalled at better rates when it was displayed more prominently in the app store than when it was just one of many screenshots in the app store. However, when the notice was displayed during app use, participants still remembered it more accurately than any of the app store conditions, as shown in Table 5.

The timing condition had a significant effect on RecallCorrect (KW $\chi^2(4)$=81.2, p=0.001). Post-hoc tests (Mann-Whitney U with Bonferroni correction) showed statistically significant differences between 'not shown' and and all of the conditions except 'app store', as seen in Table 6. The two new app store conditions were not significantly different from each other, but were significantly better than the previous app store condition, indicating that size and prominence improves recall. However, despite the improvements with the new app store conditions, participants who saw the notice

in any of the app store conditions were still less likely to remember the notice than those who saw it during app use, and this difference was statistically significant although the effect was smaller. This indicates that despite our efforts to improve the app store notice, during app use notices still had better rates of recall. Furthermore, it suggests that timing has a significant effect on retention which is only marginally influenced by how the notice is displayed.

condition	not shown		app store		a.s. big		a.s. popup	
	r	p	r	p	r	p	r	p
app store	0	1						
a.s. big	0.3	.005*	0.2	.073				
a.s. popup	0.4	.001*	0.3	.001*	0.0	.761		
during use	0.6	.001*	0.6	.001*	0.4	.001*	0.2	.007*

Table 6: Follow-up web survey: r (effect size) and p-values of pairwise comparisons on recallCorrect using Mann-Whitney U with Bonferroni correction. Significant results marked with * (α=.05).

7. LIMITATIONS

While our studies successfully show that when a privacy notice is shown significantly impacts users' ability to recall the notice content, our work has a number of limitations.

We did not study the impacts of habituation on users' ability to recall the notice. Although we used a privacy notice modeled after a standardized notice, there is little indication that many app developers have adopted this notice yet. It is possible that if this notice is widely adopted across apps, smartphone users may begin to ignore them, no matter when they are shown.

While we tried to recruit participants without introducing too much bias in the sample, some bias is inevitable, and we do not claim to have a representative sample. While we attempted to design our questions about the notice to reduce bias, by including negative and positively worded questions, responses might still be biased.

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

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

We only used one app to isolate the timing effects. The app itself was rather innocuous – the nature of the app content and interaction might not raise red flags for many users. While we designed it this way to isolate the impact of the privacy notice, users may be likely to scrutinize an obviously intrusive app more carefully. Furthermore, this study was done on only one notice.

We do not assume that recall will necessarily change behavior, as we recognize that many elements go into a smartphone users' decisions about privacy. Furthermore, we specifically studied recall of a privacy notice, and did not examine users' ability to or desire to control data sharing, which could be included in future work.

8. DISCUSSION

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

A notice shown in app use may be more salient to users, leading to the better recall we found. The fact that the notice interrupted the app usage may have helped the user pay attention to it. Further work is needed to examine habituation to notices shown during app use and determine how frequently the notices should be displayed.

When the notice is shown in the app store as one screenshot of many, it competes with other information on the screen (such as app title, developer, the install button), while the notice shown during app usage was a modal dialog that occupied the entire screen. Our second web survey attempted to understand if this is why the app store was ineffective, by testing options to display the notice more prominently in the app store. The more prominent conditions in the app-store had better rates of recall than not showing the notice, but were still recalled significantly less than when the app was shown during app use. Since the app store options we tested in the follow-up survey are not currently available to app developers, we propose that app store designers consider offering new options for app-store privacy notices that allow the notices to be shown in larger sizes with fewer distractions.

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

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

This work also indirectly opens discussions about methods for research on smartphone privacy notices. Although web surveys are often criticized for a lack of ecological validity, previous work on smartphone privacy notices has relied on using web surveys on Amazon MTurk. Our findings that the field experiment and web survey yielded similar effects between comparable conditions is interesting, if preliminary. The findings from this study may suggest

that web surveys can be a useful tool in examining smartphone privacy notices.

This research was funded in part by NSF grants CNS-1012763, CNS-1330596, and DGE-0903659 and a John and Claire Bertucci Fellowship. Many thanks to the assistance from phone-lab.org.

9. REFERENCES

[1] A. Acquisti. Nudging Privacy: The Behavioral Economics of Personal Information. *Security & Privacy, IEEE*, 7(6):82–85, 2009.

[2] A. Acquisti and J. Grossklags. Privacy and Rationality in Individual Decision Making. *Security & Privacy, IEEE*, 3(1):26–33, 2005.

[3] A. Adams and M. Sasse. Users Are Not the Enemy. *Communications of the ACM*, 42(12):40–46, 1999.

[4] I. Adjerid, A. Acquisti, L. Brandimarte, and G. Loewenstein. Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. In *Proc. of SOUPS*. ACM, 2013.

[5] H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. Cranor, and Y. Agarwal. Your Location has been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proc. of CHI*. ACM, 2015.

[6] J. J. Argo and K. J. Main. Meta-analyses of the Effectiveness of Warning Labels. *Journal of Public Policy and Marketing*, 23(2):193–208, 2004.

[7] R. Balebako, J. Jung, W. Lu, L. Cranor, and C. Nguyen. "A Lot of Little Brothers:" Measuring User Confidence in Smartphone Security and Privacy. In *Proc. of SOUPS*, 2013.

[8] R. Balebako, A. Marsh, J. Lin, J. Hong, and L. F. Cranor. The Privacy and Security Behaviors of Smartphone App Developers. *Workshop on Usable Security*, 2014.

[9] R. Balebako, R. Shay, and L. F. Cranor. Is Your Inseam a Biometric? Evaluating the Understandability of Mobile Privacy Notice Categories. *Cylab Tech. Report*, 2013.

[10] R. Böhme and J. Grossklags. The Security Cost of Cheap User Interaction. In *Workshop on New Security Paradigms*, pages 67–82. ACM, 2011.

[11] R. Böhme and S. Köpsell. Trained to Accept?: A Field Experiment on Consent Dialogs. In *Proc. of CHI*. ACM, 2010.

[12] L. F. Cranor. A Framework for Reasoning About the Human in the Loop. *UPSEC*, 8:1–15, 2008.

[13] S. Egelman, A. Felt, and D. Wagner. Choice Architecture and Smartphone Privacy: There's A Price For That. In *Proc. of WEIS*, 2012.

[14] S. Egelman, J. Tsai, L. Cranor, and A. Acquisti. Timing is Everything?: The Effects of Timing and Placement of Online Privacy Indicators. In *Proc. of CHI*. ACM, 2009.

[15] A. Felt, S. Egelman, M. Finifter, D. Akhawe, and D. Wagner. How to Ask For Permission. *HOTSEC 2012*, 2012.

[16] A. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android Permissions: User Attention, Comprehension, and Behavior. *Proc. of SOUPS*, 2012.

[17] H. Fu, Y. Yang, N. Shingte, J. Lindqvist, and M. Gruteser. A Field Study of Run-Time Location Access Disclosures on Android Smartphones. In *Workshop on Usable Security (USEC)*, 2014.

[18] C. Gates, J. Chen, N. Li, and R. Proctor. Effective Risk Communication for Android Apps. *IEEE Trans Depend. Sec. Comp.*, 11(3):252–265, May 2014.

[19] C. Gates, N. Li, H. Peng, B. Sarma, Y. Qi, R. Potharaju, C. Nita-Rotaru, and I. Molloy. Generating Summary Risk Scores for Mobile Applications. *IEEE Trans. Depend. Sec. Comp.*, 11(3):238–251, 2014.

[20] N. S. Good, J. Grossklags, D. K. Mulligan, and J. A. Konstan. Noticing Notice: A Large-scale Experiment on the Timing of Software License Agreements. In *Proc. of CHI*. ACM, 2007.

[21] J. L. Hall. NTIA Multistakeholder Process Delivers Increased App Transparency. https://cdt.org/blog/ntia-multistakeholder-process-delivers-increased-app-transparency/.

[22] M. Harbach, M. Hettig, S. Weber, and M. Smith. Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions. In *Proc of CHI '14*, pages 2647–2656. ACM, 2014.

[23] M. M. Jan Boyles, Aaron Smith. Privacy and Data Management on Mobile Devices. *Pew Internet and American Life Project*, August 2012.

[24] P. Kelley, L. F. Cranor, and N. Sadeh. Privacy as Part of the App Decision-Making Process. In *Proc. of CHI*. ACM, 2013.

[25] I. Liccardi, J. Pato, and D. J. Weitzner. Improving User Choice Through Better Mobile Apps Transparency and Permissions Analysis. *Journal of Privacy and Confidentiality*, 5(2):1, 2014.

[26] J. Lin. Understanding and Capturing People's Mobile App Privacy Preferences. Technical Report Ph.D Thesis CMU-CS-13-127.

[27] J. Lin, S. Amini, J. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy through Crowdsourcing. *Proc. of UbiComp 2012*, 2012.

[28] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 15(4):336–355, 2004.

[29] National Telecommunication and Information Administration. Privacy Multistakeholder Process: Mobile Application Transparency. `http://www.ntia.doc.gov/category/privacy/u`, Jul. 2013.

[30] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A Design Space for Effective Privacy Notices. *Proc. of SOUPS*, 2015.

[31] J. Tan, K. Nguyen, M. Theodorides, H. Negrón-Arroyo, C. Thompson, S. Egelman, and D. Wagner. The Effect of Developer-Specified Explanations for Permission Requests on Smartphone User Behavior. In *Proc. of CHI*, pages 91–100. ACM, 2014.

[32] M. S. Wogalter, D. DeJoy, and K. R. Laughery. *Warnings and Risk Communication*. CRC Press, 2005.

APPENDIX

A. SURVEY QUESTIONS

This appendix shows the final web survey that participants were asked to respond to after playing the app and experiencing a distractor or delay. This corresponds to steps 4 and 5 described in the Section 3.

A.1 App Review: Please tell us what you thought of this app.

Only field study participants completed this page.

1) How would you rate this app? (5 stars means a great app)

2) Please write a review of the app. (Imagine this is for the app store)*

3) I read the privacy policies of smartphone apps and websites.*
( ) Always
( ) Sometimes
( ) Rarely
( ) Never

4) How did you hear about this survey?*

5) Did someone who did the study before you did tell you anything about the study before you did it? If so, what did they tell you?

6) How did you hear about this survey?*

A.2 Game Review Questions

All participants (field experiment and web surveys) completed this page. Question order was randomized. We have marked the correct answers here with an "X".

7) What was the title of the app?*
( ) US History Questions
( ) History of US Inventions Quiz
( ) Inventions in US History Quiz
(X) US Inventors History Quiz

8) Do you remember seeing the privacy notice?*
( ) No, not at all
( ) Vaguely
( ) I remember most of it
( ) Yes, I remember it well

9) Which of the following people were you asked about in this app?*
( ) Louis Armstrong
(X) Elijah McCoy
( ) Willie Brown
( ) Benjamin Banneker

10) What information was collected by the app?*
( ) Financial Information
( ) Which other apps are installed on my phone
(X) Browser History
( ) User Files
( ) I don't remember
( ) Nothing

11) With whom does the app share data?*
( ) Government entities
( ) Social Networks
(X) Ad networks
( ) Consumer Data Reseller
( ) I don't remember
( ) No one

12) Which of the following people were you asked about in this app?*
( ) Barack Obama
(X) Valerie L. Thomas
( ) Ralph Ellison
( ) Frederick McKinley Jones

13) What color was the background of the app?*
( ) Green
( ) Red
( ) Blue
( ) White
(X) Black

A.3 Purpose of the Study:

*All participants (field experiment and web surveys) completed this page. Participants were shown the privacy notice again, along with the following text, "This is the privacy notification for the app. Please note, we did *not* collect this information, but please imagine your reaction if this really occurred on your phone." Participants could not return to the previous page to correct their recall answers.*

14) *The next four questions were shown in random order. Possible answers were: Strongly disagree, Disagree, Neutral, Agree, Strongly agree, and Not Applicable/Don't Remember.*

Please select whether you agree or disagree with the following statements about when you saw the notification:

- The privacy notification was shown at a time when I could pay attention to it.
- The privacy notification was shown at a time when I could make decisions about whether to allow the data collection.
- The privacy notification disrupted my use of the app.
- The privacy notification occurred at an unexpected time.

15) *The next six questions were shown in random order. Possible answers were: Strongly disagree, Disagree, Neutral, Agree, and Strongly agree.*

Please select whether you disagree or agree with the following statements about the information in the privacy notice:

- I would want notifications like this when I download or use an app.
- It is important for me to remember what the notification says while I'm using the app over time.
- This notification could be improved so I understand it better.
- The privacy notification gave me information I care about.
- I was surprised by what I learned from the privacy notification.
- I expected the app to collect my browser history and share it with ad networks.

16) Is there anything you would like to know that wasn't clear from the notification?

17) Is there anything else you would like to tell us about the privacy notification?