An Online Experiment of Privacy Authorization Dialogues for Social Applications

Na Wang

Pennsylvania State University University Park, USA nzw109@ist.psu.edu

Jens Grossklags

Pennsylvania State University University Park, USA jensg@ist.psu.edu

Heng Xu

Pennsylvania State University University Park, USA hxu@ist.psu.edu

ABSTRACT

Several studies have documented the constantly evolving privacy practices of social networking sites and users' misunderstandings about them. Researchers have criticized the interfaces to "configure" privacy preferences as opaque, uninformative, and ineffective. The same problems have also plagued the constant growth of third-party applications and their troubling privacy authorization dialogues. In this paper, we report the results of an experimental study examining the limitations of current privacy authorization dialogues on Facebook as well as four new designs which we developed based on the Fair Information Practice Principles (FIPPs). Through an online experiment with 250 users, we study and document the effectiveness of installation-time configuration and awareness-enhancing interface changes.

Author Keywords

Third-party applications (apps); information privacy; privacy notice and consent; privacy awareness and control; online social networks (OSNs)

ACM Classification Keywords

D.4.6 [Security and Protection]: Access Controls; H.5.2 [Information Interfaces and Presentation]: User Interfaces – Evaluation/methodology; K.4.1 [Public Policy Issues]: Privacy

General Terms

Human Factors; Design; Security.

INTRODUCTION

With the thriving popularity of Online Social Networks (OSNs), an increasingly large number of users share personal information, activities, opinions, photos and videos on OSNs. This trend is giving rise to growing privacy concerns by consumers about the potential misuse of their information by various stakeholders including providers of OSNs, marketers, and other users [1, 16]. Privacy concerns pertain to the acquisition of personal data

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

CSCW '13, February 23–27, 2013, San Antonio, Texas, USA. Copyright 2013 ACM 978-1-4503-1331-5/13/02...\$15.00.

and the potential risks that users may experience as a result of possible privacy breaches [1]. Recently, additional complexities of studying privacy in the context of OSNs have been introduced by the increasing popularity of third-party applications ("apps"). It has been reported by the Wall Street Journal that many popular apps on Facebook have been transmitting users' personal information and their friends' information to various advertising and data tracking firms [28]. Due to the inability to monitor the data use by app providers, users need to account for the inherent uncertainty about the behaviors of many different developers rather one large OSN site (i.e., Facebook).

To address the critical privacy concerns for third-party apps, we conducted this research to investigate whether consumers can more adequately represent their preferences for sharing and releasing personal information with our newly proposed privacy authorization dialogues. Our designs draw upon the internationally recognized Fair Information Practice Principles (FIPPs) and address important interaction problems identified in our previous study [31]. Further, we conduct a series of online experiments to examine the impact of these new interfaces on users' privacy behaviors. We also compare our results with a baseline treatment, i.e., the original authorization dialogue employed by Facebook. This research is not targeted at making value judgments about desirable user practices (e.g., to decide whether an app should be installed or not). Instead, we are interested in understanding the relative observable effect of our proposed redesign elements on the practice of notice and consent on Facebook.

In implementing our online experiment, we aimed for a realistic integration of our design in the typical experience of user-to-app interactions on Facebook. To that end, we recruited real Facebook users who followed our study protocol using their own accounts. In addition, we employed an innovative experimental procedure that mimicked the Facebook's privacy authorization dialogues via Chrome browser extension. The method is similar to a Man-in-the-Middle Attack in the sense that the user expects to communicate exclusively with the OSN but in reality interacts with a modified version of the website.

Our work is significant given the wealth of data that is continuously harvested on OSNs. In addition, the practices we study are broadly alluded to recently released privacy agendas by the FTC [9], the White House Technology

Office [23], and the European Union [8]. However, these developments stand in contrast to the relatively modest investments by the policy and research communities to study their effectiveness and practical implications.

RELATED WORK

Privacy Concerns Pertaining to Third-Party Desktop and Mobile Apps

In the domain of personal desktop computing, third-party apps were already widely in use and available from a variety of sources. For example, CNET's download.com hosts ten-thousands of apps for a variety of operating systems that originate from a wide range of large and small-scale software developers. Unsurprisingly, such repositories may include apps with problematic security and privacy practices. Grossklags and Good analyzed the End User License Agreements (EULAs) of the 50 most popular download.com apps (in a sample from 2006) and found that they contained highly problematic provisions regarding privacy and usage rights. In addition, they found that those statements were opaque, inaccessible and lacked readability [15].

Good *et al.* studied the effectiveness of providing a shortened version of the EULAs for different desktop programs with potentially harmful privacy and security aspects (which were disclosed in the text of the agreement). They found that users (when asked) appreciated the availability of a concise user notice; however, rarely stopped to study them in detail during the installation [13]. In a follow-up study which included a larger user population they recorded a statistically significant reduction of completed installations for the worst programs in their sample. Nevertheless, of the remaining users (who installed consumer-unfriendly programs) a significant share later regretted their decision [14].

Alternative solution approaches to ineffective notification that have been repeatedly suggested include the reliance on a review and reputation process as well as basic quality control. However, external certification has been criticized on the basis of adverse selection, i.e., that mostly those programs seek certification that are suffering from weak reputation and include problematic practices [7].

Recently, research has increasingly focused on privacy leakage and security concerns associated with third-party apps on smart phone systems and closed tablet computing platforms [3, 12, 17]. Those include Google Android, Symbian and Apple iOS for iPad and iPhone who all use some form of application permission and/or review process. On those markets, some apps have been classified as malware by researchers and anti-virus companies. In a recent study of 46 incidents across different mobile platforms, 28 apps were actively trying to exfiltrate user information and 4 searched for user credentials, 24 triggered premium calls or SMS messages and 8 sent unsolicited marketing messages [10].

In a closely related study, Felt et al. conducted a survey of permissions on the Android system for 100 paid and 856 free applications [11]. They found that 93% of free and 82% of paid applications included at least one potentially dangerous permission request. The authors suggested that the associated user dialogues did not allow users to meaningfully discriminate because almost all applications included potentially unwanted practices. The researchers conducted a follow-up Internet survey and usability laboratory study on Android permissions and showed that only a very small share of the study participants were attentive to the tested permissions dialogues, and could answer simple comprehension questions afterwards [12].

Anderson and his colleagues conducted case studies to examine the application markets from personal computers, mobile phones, web browsers, and online social networks. They identified security problems existing in these platforms and also proposed economic solutions [3]. Their study makes the connection to our context of investigation, i.e., authorization dialogues on online social networks, but does not directly address usability concerns.

Facebook Apps' Problematic Privacy Practices

Adding to the previously cited Wall Street Journal (WSJ) report on how apps were exfiltrating identifiable user information and data about users' sharing behavior and interests [28], Krishnamurthy and Wills report how this leaked information was in turn shared with third-party aggregators and advertising partners [21].

However, in addition to these technical findings, there is relatively little user-oriented research on third-party apps on OSNs. In a small-scale qualitative study, Besmer and Lipford examined motivations, intentions, and concerns of users when they engage with applications, as well as their perceptions of data sharing. Their results indicate that Facebook users are not truly understanding and consenting to the risks of apps maliciously harvesting profile information [4]. King and her colleagues conducted a survey study about users' misunderstandings and confusion concerning apps' functionality and information practices [20]. Survey participants self-reported their behavior with respect to the privacy authorization dialogue, but it was not studied experimentally: 44 percent responded that they had read the information, 28 percent answered that they would not read these statements, 25 percent stated that they had read a notice at some earlier time, and 3 percent could not recall whether they had read it or not.

Taking a design perspective, Hull et al. suggest visualization enhancements of the third-party apps' information accessing and publishing practices [18]. In doing so, users might have a better awareness how the app will use their information and thus users might be able to avoid some undesirable information leakage. In a small-scale design study, Tam et al. tested various user interface elements to describe privacy and security consequences

(e.g., icons, paragraphs) [29]. They reported that the variation of the disclosure design had only limited impact on participants' ability to learn about the data practices. Participants further disliked designs that used verbal descriptions in the form of paragraphs (i.e., short descriptions of practices), and preferred icons and images.

THE DEVELOPMENT OF NEW DESIGNS FOR AUTHORIZATION DIALOGUES

Problems in the Current Design

In our previous work, we studied the Facebook authorization dialogue for third-party apps (Figure 1) from different perspectives and identified several significant problems in information transmission and in the options available to the user to configure important aspects of the disclosure consequences (see Wang et al. [31]).

- **Problem 1**: When an app is asking for publishing permissions and data access permissions at the same time, users are confused and may not be able to distinguish these permissions and do not know how the app will use their information.
- **Problem 2**: During the process of adding an app to users' profiles, they do not have any installation-time control to limit or configure the app's access to their information or restrict app's publishing ability. Only after users add the app, they can edit selected categories of data access or publishing options from their privacy settings.
- **Problem 3**: During the process of adding the apps to their profiles, users do not have any control to limit whether other users can see their app activities. Only after they add the app, users can change the visibility of their app activities via adjusting options that are deeply buried in their privacy settings.
- Problem 4: Users may easily give out particularly sensitive private information or share information with third parties from which crucial identifying data can be inferred. For example, information about an individual's place and date of birth can be exploited to predict his or her Social Security Number (SSN) [2].

These observations have helped us to identify suitable design heuristics that we outline below.

Design Heuristics

The information flow between individuals and other entities takes place in the context of ever-present and often conflicting simultaneous information needs. To better understand how to assure privacy and security, we must first understand the flow of personal information among various entities. Xu et al. [32] noted that concerns over information flow may be governed by two larger dimensions: 1) concerns over information release at the front-end where personal data flow in and out of users' accounts (including the data exchange at the moment of

installing apps in our context, i.e., installation time) within a specific platform (e.g., Facebook); and 2) concerns over unwanted access and use of personal information at the back-end where user data are transferred, stored and processed across different platforms. Based on the conceptual distinction above, the practicality and utility of the privacy enhancing interfaces and technologies will eventually depend on how a candidate solution addresses the issues related to these two dimensions.

In the past decades, the U.S. government and privacy scholars have proposed a number of principles to protect users' online privacy [22, 25, 27]. Among these principles, it has been argued that the Fair Information Practices Principles (FIPPs), which originated from a study commissioned by the U.S. Department of Health, Education and Welfare in 1973 [25], have become the de-facto global standard for ethical use of personally identifiable information by large organizations [24]. In particular, four of the total of five FIPPs (excluding enforcement/redress) are directly applicable to the problem context and can potentially empower user control over both frond-end and back-end information flow [32]: notice/awareness that their personal information is being collected, consent/choice with regard to the authorized use of their information, access/participation to view personal information the firm has collected, and security/integrity to prevent these data records from unauthorized access.

In the context of this research, we are aware of the fact that it is hard to implement all of the FIPPs without corresponding policy changes supported by Facebook or other stakeholders. However, we can investigate the impact of a limited implementation of FIPPs at the front-end where the data exchange occurs at the moment of installing apps, i.e., installation time. Consequently, we mainly focus on the following three FIPPs at the front-end: (1) notice/awareness, and (2) choice/consent, and (3) access/participation. Based on these three FIPPs, we propose the following design principles:

- Principle 1 (Notice/Awareness): The authorization dialogue should provide explicit information for users to learn what data would be accessed by the app and how the data would be used.
- Principle 2 (Choice/Consent): The authorization dialogue should provide options for users to control information access or publishing ability before adding the app to the user's Facebook profile (i.e., at installation time).
- **Principle 3 (Access/Participation):** The authorization dialogue should provide options for users to control who can see their app activities.
- **Principle 4 (Notice/Awareness):** The authorization dialogue should provide alert signals for users when the app asks for users' sensitive private information.

Permission Number of apps requesting permission (percentage of apps requesting permission)		Total times a permission is requested by apps		Number of apps requesting permission (percentage of apps requesting permission)		Total times a permission is requested by apps	
basic information	9411	100.00%	502,755,469	friends_education_history	14	0.15%	3,564,500
Email	3234	34.36%	314,855,710	friends_activities	22	0.23%	3,448,300
publish_stream	4702	49.96%	259,917,056	friends_about_me	17	0.18%	3,328,000
user_birthday	902	9.58%	138,257,300	friends_interests	13	0.14%	3,163,500
publish_actions	513	5.45%	125,686,870	user_work_history	73	0.78%	2,961,900
user_location	343	3.64%	55,077,200	friends_relationships	3	0.03%	2,912,000
offline_access	660	7.01%	42,491,210	user_photo_video_tags	98	1.04%	2,779,680
read_stream	528	5.61%	37,863,840	friends_photo_video_tags	32	0.34%	2,423,340
user_photos	491	5.22%	24,940,010	friends_likes	36	0.38%	2,385,960
user_about_me	248	2.64%	23,700,430	user_status	40	0.43%	1,827,500
friends_birthday	206	2.19%	19,237,740	user_checkins	17	0.18%	1,422,000
user_likes	214	2.27%	13,486,760	friends_checkins	6	0.06%	1,350,000
friends_photos	214	2.27%	13,051,340	user_religion_politics	19	0.20%	1,183,100
friends_online_presence	121	1.29%	10,745,500	publish_checkins	11	0.12%	980,700
user_interests	68	0.72%	9,675,600	manage_notifications	8	0.09%	976,000
user_hometown	120	1.28%	9,594,040	user_relationships	34	0.36%	969,600
user_online_presence	110	1.17%	8,298,400	friends_relationship_details	4	0.04%	741,000
friends_location	104	1.11%	8,121,000	read_friendlists	39	0.41%	603,800
xmpp_login	13	0.14%	7,744,000	user_videos	12	0.13%	560,780
user_education_history	51	0.54%	5,920,640	user_relationship_details	19	0.20%	406,200
friends_hometown	21	0.22%	5,862,500	create_event	11	0.12%	336,500
friends_work_history	86	0.91%	5,260,660	user_groups	10	0.11%	294,900
user_activities	53	0.56%	5,204,740	friends_videos	2	0.02%	230,400
manage_pages	60	0.64%	4,725,900	user_website	2	0.02%	130,000

Table 1. Most Frequently Requested Permissions by the Applications.

To further increase the relevance of our work, we also conducted a large-scale measurement study to gain a broad perspective of the data practices of the various app providers and Facebook. The latter study helps us to propose experimental layouts of high relevance to the Facebook user community.

Scope of Permissions

To determine the scope of permissions that should be included into our design of privacy authorization dialogues, we investigated the data practices from the 9,411 most popular third-party apps on Facebook which displayed the typical privacy authorization dialogue to users (as shown in Figure 1). We also collected data to assess the scale of data collection for the apps under consideration.

From the app developer's perspective, there were 63 types of permissions they can request from users. For each of these permissions, we first compiled a list of applications that request each type of permission. We summed up the number of monthly active users for each application on the list to get the total number of users who were affected by a certain data practice (see Table 1). Further details about the measurement methodology can be found in Wang [30].

Table 1 shows how many apps request a particular permission, and given the apps' popularity how many times users on Facebook have shared this information with app providers.



Figure 1. Example of an Authorization Dialogue Page.

As shown in Table 1, users shared their basic information more than 500 million times with apps. The next three most frequently requested permissions are: 1) "email", which allows an app to access a user's primary email address; 2)

"publish_stream", which enables an app to post content, comments, and likes to a user's stream and to the streams of the user's friends; and 3) "user_birthday", which permits an app to access a user's birthday.

Based on the results shown in Table 1, we included users' basic information and an additional twelve frequently requested permissions in our design of the privacy authorization dialogue. We were able to categorize these permissions into three categories: 1) permissions related to users' information releasing behaviour (email, user_photos, user videos, user birthday, user hometown, user location), 2) permissions related to users' friends' information releasing behaviour (friends birthday, friends_hometown, friends_location, friends_photos, and friends_videos), and 3) a permission related to information reposting (publish stream). By selecting these permissions, we kept our privacy authorization dialogues within a reasonable length while adequately representing the most frequently requested data types and categories of permissions. We added the currently less utilized video permissions to account for the trend towards increased multimedia utilization.

Design Considerations about Format and Style

Kelley et al. developed a privacy "nutrition label" that presents to consumers the ways organizations collect, use, and share personal information [19]. Their design aims to:
1) clearly highlight the meaning of different labels so that users can easily understand the different sets of information; 2) use different font highlights to separate sets of information in order to expedite the users' navigation through the list; and 3) have a bold and clear title to inform users with the purpose of the information in each section. Those ideas are derived from previous work on food safety warnings and information about nutritional content which are reviewed in Kelley et al. [19]. We also draw upon recent research by Bravo-Lillo et al. on effective warning mechanisms to protect people from privacy harms [6].

In this research, we aim to include design elements mentioned above. In particular, we present four independent and distinct interfaces covering different aspects derived from our design heuristics. By incrementally adding design elements into our interface, we can test the impact of the design heuristics in a progressive fashion which is highly beneficial given the complexity of typical interactions between users and social apps. In this way, our results are easier to replicate and to rationalize. In the following section, we describe each of our designs in detail.

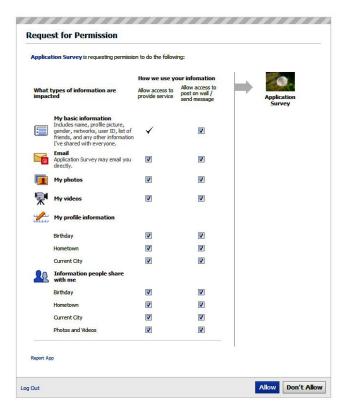


Figure 2. Proposed C Design.

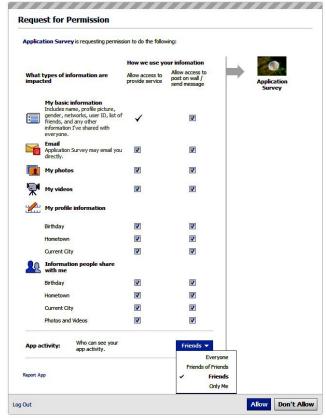


Figure 3. Proposed CAA Design

Four Designs for Authorization Dialogues

Check-box Authorization Dialogue (C): The C interface design of the authorization dialogue (see Figure 2) aims to fulfill the first two design heuristics. Below we describe our major design elements.

- The Granular Layout of Permissions: All types of data (basic information and data reading permissions) required by the app are listed in the first column. The top row displays the information regarding how the app will use the data (including data writing and page management permissions).
- The Tick Marks and Checkboxes: Un-clickable tick marks represent those types of information that will be accessed and used by the app and are non-negotiable. The checked check box means that users will allow the app to access and use certain information. When unchecked, users will not allow the app to access or use the corresponding information.

Taken together this design allows us to investigate the relative impact of granular installation-time configuration (opt-out) options for apps' data practices (compared to a baseline treatment of the current Facebook design applied to our scenario displayed in Figure 1).

Check-box and App Activity Authorization Dialogue (CAA): Our second design of the authorization dialogue, the CAA design, is an enhanced version of the C design, in addition to fulfilling the first two design heuristics, it also aims to address the third one (see Figure 3).

• The "App activity" Drop-down List: It allows the user to decide whether other users (i.e., Friends, or Friends of Friends, or the public) can see users' app activity on Facebook. Users can change this setting by using a drop down menu.

That is, the dialogue now offers the user control options directed towards the app developer as well as other users.

Check-box and Signal Authorization Dialogue (CS): Our third design of the authorization dialogue, the CS design, is another variation of an enhanced version of the C design; in addition to addressing the first two design heuristics, it also considers the fourth one (see Figure 4).

• The "i" Mark and Color Scheme: As users' basic information is always requested by the app, here we use the blue "i" signal to remind users that this information cannot be opted out. We use the red "i" signal to alert users that certain information is particularly sensitive. In highlight email, study, we user birthday, user_hometown, and user_location sensitive as information mainly because privacy advocates advise not to share these information with third parties [2]. Both marks have tooltip information which is accessible to users when they move their mouse pointer over the sign.

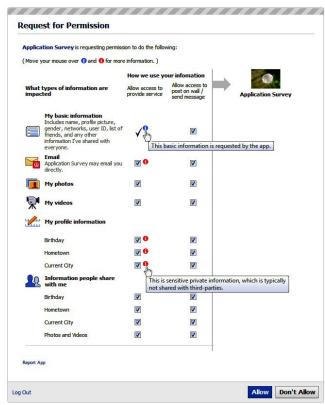


Figure 4. Proposed CS Design.

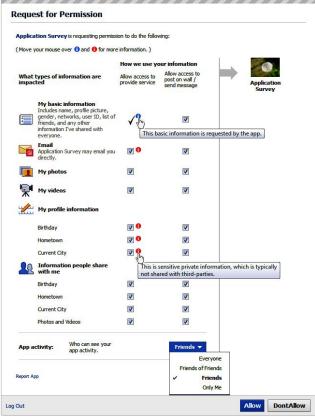


Figure 5. Proposed CSAA Design.

This treatment allows us to study the relative impact of awareness-enhancing interface alerts which signal when an app requests a user's sensitive private information (and allows for a comparison with the checkbox design C and the baseline treatment).

Check-box, Signal, and App Activity Authorization Dialogue (CSAA): The fourth design of the authorization dialogue, the CSAA design, addresses all four design heuristics (see Figure 5). With this interface, we want to examine whether the third and the fourth design heuristics combined together would help users better protect their privacy on Facebook.

EXPERIMENTAL DESIGN

Implementation

To implement our proposed designs, we employed a semifunctional Wizard of Oz approach to visually mimic Facebook's default third-part apps authorization dialogue. The method is similar to a Man-in-the-Middle Attack in the sense that the user expects to communicate exclusively with Facebook.com but in reality interacts with a modified version of the website. This approach was implemented with a Chrome browser extension that integrated into the authorization process by capturing a particular Facebook app's unique ID. Once the app's unique ID was captured, the extension replaced the original authorization dialogue with one of our four proposed designs (for treatment II-V), or left the original interface (the baseline treatment). In all cases (including the baseline treatment) we activated a redirect URL when participants clicked the "Allow" or "Don't allow" button. All these replacements were implemented by modifying the authorization page's HTML Document Object Model (DOM). The browser extension recorded users' interactions with the interface and the time they spent on that page. However, we neither recorded users' identifiable information, nor called the Facebook API to collect information from users' profiles.

Participants

We recruited participants via Amazon Mechanical Turk (MTurk: www.mturk.com). On MTurk, requesters post Human Intelligence Tasks (HITs) by uploading job descriptions onto Amazon's web portal. MTurk maintains each Turker's performance history and additional information, which requesters may use to specify who is eligible to perform a particular HIT. Eligibility may include the Turker's location (country), HIT completion rate (fraction of tasks completed among those signed up for in the past) and approval rate (fraction of tasks accepted by requesters among those completed in the past). The requester must also specify the amount of payment a Turker will receive once the task is completed and the work is accepted by the requester. Once a requester posts a HIT on MTurk, as in our application installation task, eligible Turkers can immediately view it and sign up.

We recruited 276 Turkers with a North American IP address and a previous HIT approval rate of 55% or better. Participants were also required to be Facebook users and needed to be familiar with the Google Chrome browser. To motivate Turkers to complete this study, we paid \$0.80 to each participant after we did a basic evaluation of the validity of task completion. Data from 26 participants was rejected via the MTurk web interface because these participants submitted blatantly incomplete or incoherent work. We also ensured that no individual with a particular MTurk ID would participate in our study twice.

We used a between-subjects design, where the participants were randomly assigned to one of five groups, namely, the *baseline* group I or one of the four treatment groups II-V corresponding to our four new designs of privacy authorization dialogues (see Table 2). As expected, Chisquare tests revealed that subjects assigned to the various treatments did not differ significantly in terms of their age, education, and gender (see Table 3).

Treatment	Interface Presented
I (Baseline)	Current Authorization Interface
II (C)	Check-box Design
IV (CAA)	Check-box and App Activity Design
III (CS)	Check-box and Signal Design
V (CSAA)	Check-box, Signal, and App Activity Design

Table 2. Overview of Treatment Groups.

Among the participants, 56% were females and 44% were males; they had a wide range of education levels (from less than high school to Ph.D.) and covered a wide range of age categories (from 18 to over 50) which is consistent with the diversity of the current Facebook user base. See Table 3 for more details about participants' demographics.

	Baseline	С	CAA	CS	CSAA	Total
	%	%	%	%	%	%
Sex						
Female	56	52	60	46	66	56
Male	44	48	40	54	34	44
Age						
18-24	52	40	44	46	36	43.6
25-29	20	34	26	24	36	28
30-34	14	8	14	10	16	12.4
35-39	10	12	12	12	8	10.8
40-49	0	4	2	6	4	3.2
50 and over	4	2	2	2	0	2
Education						
Less than	2	4	2	6	0	2.8
High school	22	26	20	16	22	21.2
Associate	8	14	14	20	4	12
Current	18	28	30	28	24	25.6
Bachelor's	38	18	24	24	30	26.8
Master's	10	6	8	6	18	9.6
Ph.D.	2	4	2	0	2	2

Table 3. Participants' Demographics.

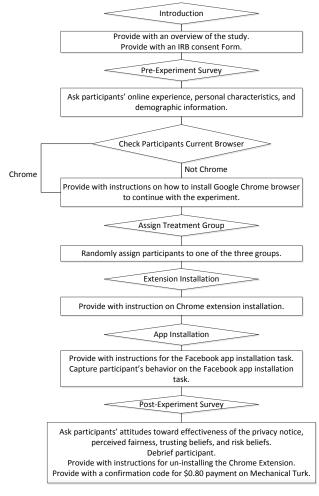


Figure 6. Experimental Protocol

Experimental Procedure and Task

Figure 6 provides an overview of the protocol followed in this experiment. Each participant was guided through this study protocol¹. Upon finishing all tasks shown in Figure 6, participants were asked to complete a questionnaire that was customized to the specific treatment and asked them to evaluate usability, security, and privacy aspects related to the privacy authorization dialogues they interacted with. We also asked participants to evaluate another two alternative designs at the end of the post-experimental questionnaire. See Table 4 for details about the post-experimental survey.

Treatment	Interfaces Evaluated and Compared
I (Baseline)	Baseline, C, and CS
II (C)	C, Baselin, and CS
IV (CAA)	CAA, Baseline, and CSAA
III (CS)	CS, Baseline, and C
V (CSAA)	CSAA, Baseline, and CAA

Table 4. Post-Experimental Survey Customization Details.

RESULTS

Our analysis of the experimental results unfolds in five steps.

1. Task Completion Times:

We estimated that the study would take about 20 minutes to complete and advertised this estimate in the HIT description. The actual average across all treatment groups was 19 minutes and 36 seconds, and Tukey HSK Test revealed that subjects assigned to the five treatment groups did not differ significantly in terms of total time to complete the task.

For participants in the baseline treatment, the average time spent on the authorization dialogue was about 9 seconds, which was significantly faster than what we observed in the other treatments (28 seconds). This effect is significant with p<.001 for treatments II, III and V with p=.025 for treatment IV for the comparisons with the baseline treatment. We attribute this to users' likely familiarity with this dialogue in the baseline condition and the complete absence of any configuration options. We observed no significant difference concerning this metric among the four new designs (i.e., treatment II to V).

2. Overall App Installation Approval Rates:

A total of 50 participants interacted with the original Facebook interface and 42 of them (84%) did "Allow" to add the app to their profiles. The alternative interfaces lowered the participants' readiness to add the apps in all cases. More precisely, 37 out of 50 (74%) for the C design, 39 out of 50 (78%) for the *CAA* design, 30 out of 50 (78%) for the *CS* design, and 30 out of 50 (60%) for the *CSAA* design, allowed the installation of the app.

We are surprised to find that the availability of granular configuration options at installation-time (C and CAA) does not increase the number of installations. We would have expected that the opportunity to opt-out from unwanted practices would make the application more attractive to participants. Instead, there is a non-significant reduction of installations.

In contrast, with the presence of warning signals, we would expect lower installation rates. In fact, this effect is strongly significant for the proposed CS (Chi-square=21.429, p<.001) and CSAA (Chi-square=21.429, p<.001) designs when compared to the original Facebook interface (i.e., the baseline treatment). More importantly, in comparison to the

¹ As participants logged into their Facebook accounts, we asked them to evaluate our Facebook application. The exact instruction was as follows: We have developed a Facebook app that would help us to contact you about the study results and for potential further survey opportunities. To continue with the study, please click the following URL then log into your Facebook account to access the Facebook application's installation page. When you visit this site you can learn more about our app and decide whether you want to continue the installation or cancel the installation. Either option will allow you to continue with the study and will not impact your Mechanical Turk payment.

C and CAA treatments, the impact of the signals is also significant (considering the relevant comparisons the effect is always at least significant at p<0.05). Nevertheless, we are somewhat surprised by this strong finding given the recent research on attention blindness, for example, when considering security indicators for Phishing and other browser warnings [5].

3. The Effectiveness of the Layout of Permissions and Checkboxes:

The original design of the authorization dialogues (i.e., baseline treatment) provides a take-it-or-leave-it option concerning the access and publishing abilities of an app, i.e., users have to accept all of these permissions if they want to use the app. In our proposed interfaces, we separated accessing and publishing permissions into different columns and enabled users to uncheck checkboxes and thereby refuse to give certain permissions. These design components were considered as the most basic ones and were implemented in all of the four proposed designs (i.e., *C*, *CAA*, *CS*, *CSAA*). For analysis purposes, if a user clicked the "Don't Allow" button to step back from the installation of the app, we regard this as if they did not release any information to the app, which is equivalent to the decision to uncheck all the checkboxes in the authorization dialogue.

Table 5 and Table 6 show that when users are interacting with the new designs, they not only tend to release significantly less information in total, but also tend to opt out of publishing permissions to prevent the app from reposting information to their wall compared to the original Facebook interface (all comparison tests of treatments with the baseline are significant at p < 0.001).

Treatment %		
I (I	Baseline)	84
II	(C)	58.27
IV	(CAA)	52.91
Ш	(CS)	46.45
V	(CSAA)	41.27

Table 5. Overall Information Release (in %).

Tre	eatment	%
I ()	Baseline)	16
II	(C)	44.55
IV	(CAA)	51.45
III	(CS)	56.55
V	(CSAA)	61.27

Table 6. Opt-Out from Publishing Permissions (in %).

Also, in the post-installation questionnaire, participants were instructed to rate the effectiveness of the interfaces. We utilized a Likert scale (i.e., how much they agree or disagree with; 1 = strongly disagree, 7 = strongly agree), and presented the participants with the following statements:

- The proposed designs could help users better differentiate between the different purposes of data usage by the app.
- The proposed designs would allow users to better control the app's access and use of a specific type of personal information.

Participants who interacted with the four proposed interfaces rated these two questions significant higher (with at least p<.01) than participants who interacted with the original Facebook authorization dialogue. These factors indicate that the layouts of permissions in our proposed designs are likely to become a more effective way to attract users' attention and to utilize their options to adjust privacy parameters.

4. The Effectiveness of the App-Activity Drop-down List:

With the original design of the authorization dialogue, users cannot modify whether other users can see their App Activity before adding the app to their profile. They can only change it by going through a series of relatively complex steps after installing the app. We believe that if users could modify this setting at installation time, it will be another improvement comparing to the original authorization dialogue. To put it differently, we can investigate the users' responses when we "bring to light" deeply buried privacy configuration options.

If a participant selected "Don't Allow" to refuse to add the app to her Facebook profile, then the user's choice on this App-Activity drop-down list will have no effect in terms of controlling who can see her app activities. Thus, in this section, we only focused on those participants who did "Allow" to add the app to their profiles.

Eight participants out of 39 (20.51%) for the CAA design, and 5 out of 30 (16.67%) for the CSAA design, who added the app to their profiles, decided not to share their app activity with others, and preferred to keep this information private (Selected "Only Me"). Those individuals who tightly restricted their App-Activity also used the opt-out options for the permissions significantly more often (p<.005 for CAA, and p<.0001 for CSAA).

We also found a treatment effect in which participants who interacted with the CAA and CSAA design tend to release significantly less information compared to those who faced the C design (see Table 7, p=.005 for CAA and p=.006 for CSAA). This happened because the app-activity drop-down list enhanced participants' awareness that their interaction with the app might be observed by other users on Facebook, and then triggered them to reduce the information released to the third-party app.

Tre	eatment	%
I (Baseline)		100
II	(C)	79.36
IV	(CAA)	70.86
III	(CS)	76.36
V	(CSAA)	70.61

Table 7. Information Release by Participants' who Installed the App (in %).

5. Further Evidence on the Effectiveness of the "i" Mark Signal:

For this analysis, we continue to focus on participants who chose to add the app to their profiles. From the participants' interaction with the authorization dialogue as well as the post-installation questionnaire, we can infer that the red "i" mark helped users to differentiate between information that is marked as sensitive and other information, and to recognize when a particular type of sensitive information is being collected by an app (i.e., in the CS and CSAA design this affects the four collection permissions for email, birthday, hometown and current city; see Figures 4 and 5). Table 8 shows that participants who interacted with the CS interface tend to release considerably less sensitive information compared with participants facing the C interface. In addition, in the post installation questionnaire, users also indicated that the interfaces (with "i" mark) helped them to better recognize when a particular type of sensitive information is being collected by the app compared with the original authorization dialogue (p < .05).

We also conducted a separate analysis targeted at the blue "i" mark which indicates that basic information is requested by the app (see Figure 7). In particular, we were wondering whether there is a distinct effect on the publishing check box to the right of the blue mark (circled in red in Figure 7). For example, we expected that users might feel that information which is considered basic is subject to fewer opt-outs. This presumes that participants would understand collection and usage of basic information as a bundle that should be treated equally. We find no such distinct effect (see Table 9) that applies to both relevant pairings (i.e., C/CS and CAA/CSAA).

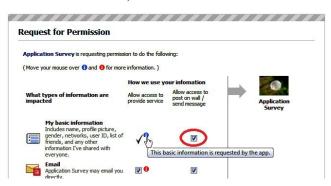


Figure 7. user_basic_information_post Check-Box Circled in Red.

In summary, the red "i" mark motivates participants to install apps less often and to release sensitive information to the experimental app less frequently. The effect of the blue "i" mark (that is attached to the permission designating access to basic information) is primarily informative. It does not appear to have a consistent behavioral impact.

So far, we have reported the overall app installation rate and the effectiveness of different design components. In the next section, we are going to discuss the findings and possible limitations of our current study and possible directions for future work.

Treatment		%	Tr	Treatment	
II	(C)	81.08	II	((
IV	(CAA)	67.31	IV	(CAA	
III	(CS)	70.00	III	(CS	
V	(CSAA)	60.00	V	(CSAA	

Table 8. Sensitive Information Release (in %).

 Treatment
 %

 II
 (C)
 21.62

 IV
 (CAA)
 35.90

 III
 (CS)
 20.00

 V
 (CSAA)
 16.67

Table 9. Opt Outs for Basic Information (in %).

DISCUSSION

While in this study we mainly focus on information disclosure issues pertaining to Facebook third-party apps, similar problems also exist in other platforms, i.e., other social media platforms, smart phone platforms, and desktop platforms. The common features of these third-party app offerings are: they often provide useful services or entertainment to users; but they also typically collect users' information and then transfer it to a server outside of the purview of the platform provider and the user. Users have an extremely difficult task to understand the ramifications and consequences of these practices. Further, beyond vague policy statements there is little actual enforcement by the platform providers to assure adequate treatment of user data.

In our work, we aim to address the vulnerability of users at the front-end of this information exchange through the deployment of improved privacy notice and consent interfaces as favored by most regulatory and self-regulatory proposals. Our study provides nuanced results about the impact of such efforts.

Surprisingly, we find that offering granular configuration options at installation-time does not increase individuals' willingness to install the experimental app. Instead, there is a non-significant reduction of installations. Since more fine-grained control allows users to account for their individual privacy preferences, we would have expected that users negotiate deals that prompt them to reject the app less often. This does not seem to be the case. However, we find that users who eventually decide to install the app make use of the granular choices and opt-out from certain data collection and usage practices. In addition, users who change the default behavior of the app concerning its interaction with other OSN users opt-out even more often.

In light of research that highlights individuals' blindness to security indicators and warnings, for example, in the context of Phishing or Spyware [5], we unexpectedly find that providing additional awareness-enhancing signals in the privacy authorization dialogue significantly lowers the number of installations.

Our study also makes progress along the dimension of methodology. We implemented our new designs as working interfaces and embedded them during controlled experiments into the Facebook environment via a Chrome extension. In this way, participants can actually interact with these interfaces on their own accounts, and we are able to collect users' actual installation operations in a plausible experimental setting. Then, in the subsequent post-experimental survey, we collect additional data for the assessment of the authorization dialogue designs.

Limitations

In our study, we have logged participants' final decisions including whether they "Allow" or "Don't Allow" the app to access and use their information and specifically what kinds of their information can be accessed by the app (which check-boxes are checked). However, we did not track several other types of conceivable interaction data with the authorization dialogue (e.g., mouse movements or eye tracking). We also did not solicit a free text response about the reasons for (not) adding the app to their profiles. Such data might have enabled us to provide further intuition about the effectiveness of certain design elements, e.g., the "i" mark, and whether they attracted users' attention and enhanced participants' privacy awareness.

We mentioned in our instructions on Mechanical Turk that the app to be installed is an "application survey" (and it indeed triggered the post-experimental survey). We also mentioned to the participants that the app would enable us to contact them later about potential future survey opportunities. We selected this framing because it naturally fit our goal to include a series of survey questions in the experimental process. It would be interesting to conduct additional treatment conditions with different framings (e.g., that control for the reputation of the app developer or the type of the application).

The study use a monitoring infrastructure based on a Google Chrome extension. According to StatCounter [http://gs.statcounter.com/] Google Chrome has now achieved a market share of about 35.7% (similar to Microsoft Internet Explorer's share and significantly higher than Firefox). While we cannot categorically exclude that Chrome users are more tech-savvy or differ in some other regard, we feel that its deep market penetration will moderate such concerns.

We used Mechanical Turk as our recruitment platform. To increase the relevance of our study, we carefully limited the user population to Northern American IP addresses. As Smith et al. [26] pointed out, different countries have approached privacy issues differently in their social norms and regulatory structures. Thus, in our study, we restrict eligibility to Turkers with a North American IP address because the technological and regulatory privacy environments in countries in North America are comparatively similar [26]. Thus a future research opportunity could be to conduct a comparative study by recruiting participants from other regions (e.g., from the E.U. or Asia).

CONCLUSION

Based on qualitative and theoretical considerations and the results from a substantial measurements study on Facebook, we proposed four new designs of the authorization dialogues for third-party apps and conducted a rigorous online experiment to investigate whether users can more adequately represent their preferences for sharing and releasing personal information with these improved designs. We uncovered significant treatment effects that may contribute to improvements of the effectiveness of authorization dialogues for third-party applications and beyond.

In the future, we intend to use our experimental setup to explore a number of related questions (e.g., the relevance of opt-in versus opt-out) to provide the CSCW research community with comparable results across a spectrum of design choices. We also would like to explore opportunities for future collaboration with Facebook and application developers to conduct large-scale field experiments in the context of naturally occurring user practices.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their comments on a previous version of this paper. Heng Xu and Na Wang gratefully acknowledge the financial support of the National Science Foundation under grant CNS-0953749. Any opinions, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Jens Grossklags gratefully acknowledges the financial support of a Google Faculty Research Award.

REFERENCES

- [1] Acquisti, A. and Gross, R. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In *Proc. PETS 2006*, ACM Press (2006), 36-58.
- [2] Acquisti, A. and Gross, R., Predicting Social Security numbers from public data. *Proceedings of the National Academy of Sciences* 106,27 (2009), 10975-10980.
- [3] Anderson, J., Bonneau, J., and Stajano, F. Inglorious Installers: Security in the Application Marketplace. In *Proc. WEIS 2010* (2010).
- [4] Besmer, A. and Lipford, H. Users' (Mis)conceptions of Social Applications. In *Proc. GI 2010*, ACM Press (2010), 63-70.
- [5] Böhme, R. and Grossklags, J. The Security Cost of Cheap User Interaction. In *Proc. NSPW 2011*, ACM Press (2011), 67-82.
- [6] Bravo-Lillo, C., Cranor, L., Downs, J., and Komanduri, S., Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security & Privacy* 9,2 (2011), 18-26.
- [7] Edelman, B., Adverse Selection in Online "Trust" Certifications and Search Results. Electronic

- Commerce Research and Applications 10,1 (2011), 17-25.
- [8] European Commission, Proposal for a Regulation of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of such Data (General Data Protection Regulation), COM (2012) 11 final, 25 January, 2012.
- [9] Federal Trade Commission, Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers. http://www.ftc.gov/os/2012/03/120326privacyreport.pdf
- [10] Felt, A., Finifter, M., Chin, E., Hanna, S., and Wagner, D. A Survey of Mobile Malware in the Wild. In *Proc.* SPSM 2011, ACM Press (2011), 3-14.
- [11] Felt, A., Greenwood, K., and Wagner, D. The Effectiveness of Application Permissions. In *Proc. WebApps 2011* (2011), 75-86.
- [12] Felt, A.P., Ha, E., Egelman, S., Haney, A., Chin, E., and Wagner, D. Android Permissions: User Attention, Comprehension, and Behavior. In *Proc. SOUPS*, ACM Press (2012).
- [13] Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D., and Konstan, J. Stopping Spyware at the Gate: A User Study of Privacy, Notice and Spyware. In *Proc. SOUPS* 2005, ACM Press (2005), 43-52.
- [14] Good, N.S., Grossklags, J., Mulligan, D.K., and Konstan, J.A. Noticing Notice: A Large-Scale Experiment on the Timing of Software License Agreements. In *Proc. CHI* 2007, ACM Press (2007), 607-616.
- [15] Grossklags, J. and Good, N. Empirical Studies on Software Notices to Inform Policy Makers and Usability Designers. In *Proc. FC 2007 and USEC* 2007, Springer Press (2007), 341-355.
- [16] Hoadley, C.M., Xu, H., Lee, J.J., and Rosson, M.B., Privacy as Information Access and Illusory Control: The Case of the Facebook News Feed Privacy Outcry. *Electronic Commerce Research and Applications* 9,1 (2010), 50-60.
- [17] Howell, J. and Schechter, S. What You See is What They Get: Protecting Users from Unwanted Use of Microphones, Camera, and Other Sensors. In *Proc. Web 2.0 Security and Privacy* (2010).
- [18] Hull, G., Lipford, H., and Latulipe, C., Contextual gaps: Privacy issues on Facebook. *Ethics and Information Technology* 13,4 (2010), 289-302.

- [19] Kelley, P., Bresee, J., Cranor, L., and Reeder, R. A Nutrition Label for Privacy. In *Proc. SOUPS* 2009, ACM Press (2009).
- [20] King, J., Lampinen, A., and Smolen, A. Privacy: Is There an App for That? In *Proc. SOUPS 2011*, ACM Press (2011).
- [21] Krishnamurthy, B. and Wills, C. On the Leakage of Personally Identifiable Information Via Online Social Networks. In *Proc. WOSN 2009*, ACM Press (2009), 7-12.
- [22] Lederer, S., Hong, I., Dey, K., and Landay, A., Personal Privacy Through Understanding and Action: Five Pitfalls for Designers. *Personal and Ubiquitous Computing* 8,6 (2004), 440-454.
- [23] Obama, B., Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy, 2012.
- [24] Reidenberg, J.R., Setting Standards for Fair Information Practice in the US Private Sector. *Iowa L. Rev.* 80,(1994), 497.
- [25] Secretary, Records, Computers and the Rights of Citizens - Report of the Secretary's Advisory Committee on Automated Personal Data Systems, 1973.
- [26] Smith, H.J., Dinev, T., and Xu, H., Information Privacy Research: An Interdisciplinary Review. MIS Ouarterly 35,4 (2011), 989-1015.
- [27] Spiekermann, S. and Cranor, L.F., Engineering Privacy. *IEEE Transactions on Software Engineering* 35,1 (2009), 67-82.
- [28] Steel, E. and Fowler, G., Facebook in Privacy Breach. http://online.wsj.com/article/SB100014240527023047 72804575558484075236968.html.
- [29] Tam, J., Reeder, R.W., and Schechter, S., I'm Allowing What? Disclosing the Authority Applications Demand of Users as a Condition of Installation, MSR-TR-2010-54, 2010.
- [30] Wang, N. Third-Party Applications' Data Practices on Facebook. In *Proc. CHI 2012*, ACM Press (2012), 1399-1404.
- [31] Wang, N., Xu, H., and Grossklags, J. Third-Party Apps on Facebook: Privacy and the Illusion of Control. In *Proc. CHIMIT 2011*, ACM Press (2011).
- [32] Xu, H., Teo, H.H., Tan, B.C.Y., and Agarwal, R., Effects of Individual Self-Protection, Industry Self-Regulation, and Government Regulation on Privacy Concerns: A Study of Location-Based Services. *Information Systems Research* (2012), doi: 10.1287/isre.1120.0416