

Facebook and Privacy: It's Complicated

Maritza Johnson
Columbia University
maritzaj@cs.columbia.edu

Serge Egelman
UC Berkeley
egelman@cs.berkeley.edu

Steven M. Bellovin
Columbia University
smb@cs.columbia.edu

ABSTRACT

We measure users' attitudes toward interpersonal privacy concerns on Facebook and measure users' strategies for reconciling their concerns with their desire to share content online. To do this, we recruited 260 Facebook users to install a Facebook application that surveyed their privacy concerns, their friend network compositions, the sensitivity of posted content, and their privacy-preserving strategies. By asking participants targeted questions about people randomly selected from their friend network and posts shared on their profiles, we were able to quantify the extent to which users trust their "friends" and the likelihood that their content was being viewed by unintended audiences. We found that while strangers are the most concerning audience, almost 95% of our participants had taken steps to mitigate those concerns. At the same time, we observed that 16.5% of participants had at least one post that they were uncomfortable sharing with a specific friend—someone who likely already had the ability to view it—and that 37% raised more general concerns with sharing their content with friends. We conclude that the current privacy controls allow users to effectively manage the outsider threat, but that they are unsuitable for mitigating concerns over the insider threat—members of the friend network who dynamically become inappropriate audiences based on the context of a post.

Categories and Subject Descriptors

K.6 [Management of Computing and Information Systems]: System Management, Security and Protection; H.1.2 [Models and Principles]: User/Machine Systems—*Human Factors*

General Terms

Security, Human Factors, Design

Keywords

Social networking, privacy, access control

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2012, July 11–13, 2012, Washington, DC USA

1. INTRODUCTION

People spend an unprecedented amount of time interacting with social network sites (SNS) and uploading large quantities of personal information [22, 19]. The dramatic growth in SNS use has created a myriad of privacy concerns. In this paper, we focus on the interpersonal privacy concerns that arise between SNS users and how they manage their concerns by expressing preferences for who should be allowed to access posted content.

Access control management is known to be a difficult problem for end-users in other domains [25, 24]. Not surprisingly, the task of correctly configuring privacy controls, a particular type of access control management task, is out of reach for many SNS users [21]. Knowing that SNS privacy settings are difficult to manage correctly, our research furthers the goal of designing a more usable mechanism, beginning with the question, *How likely are Facebook users to share content with unintended audiences, and what mitigation strategies do they use?*

We constructed an interactive Facebook application to survey 260 Facebook users about specific pieces of content that they had posted to their profiles, as well as their levels of comfort sharing content with randomly selected people from their friend networks. We observed that many participants (94.6% of 260) deny access to their profile content—posts or photos—to people outside their friend network (e.g., strangers). Participants who were concerned with sharing specific posts with strangers were significantly more likely to block strangers from accessing their profiles. Thus, our results indicate that users effectively mitigate their concerns over sharing content with strangers. At the same time, we found that users increasingly experience sharing concerns that involve members of their friend networks and that these concerns are not mitigated by the existing access control settings. Specifically, we found that 37% raised concerns over sharing specific posts with subsets of their friend networks or with allowing specific friends to view their profile information. Thus, while global privacy settings have helped users cope with the threat of strangers viewing content, our results indicate that they do not adequately address the "insider threat." We observe that out of necessity, many users have developed several ad hoc approaches to preserving interpersonal privacy.

We begin by discussing prior research on SNS usage, and prior studies of SNS users' privacy concerns and mitigation strategies (Section 2). Then, we present our methodology (Section 3), including a description of the Facebook application we implemented to execute the study. Next, we present

data on the privacy concerns users experience and the techniques they employ to mitigate their concerns (Section 4). We conclude by discussing how our data demonstrate a shift in privacy concerns from situations that involve outsiders to situations that involve people within the friend network (Section 5). Based on our results, we suggest a new focus for future research and highlight aspects of our approach that need to be adjusted to ensure meaningful progress toward the goal of usable privacy controls for SNS users.

2. BACKGROUND AND RELATED WORK

We examine how SNS users manage their interpersonal privacy while sharing and interacting with other users. We subscribe to Altman’s definition of privacy and equate interpersonal privacy to “an interpersonal boundary regulation process used by people to regulate their interactions with others” [2]. Palen and Dourish’s discussion of digital privacy is also relevant, specifically the point that privacy is dynamic and requires users to satisfy constraints that vary across contexts [23]. We focus our discussion of related work on Facebook privacy and the available privacy controls, the difficulties users face in their attempts to manage interpersonal privacy, qualitative studies of SNS users’ privacy concerns, and users’ strategies for mitigating their concerns.

The aggregate of media reports and the results of prior work create a perplexing view of SNS users and privacy. Despite the multitude of privacy concerns that accompany SNS use [9], the number of people who are active SNS users continues to grow and users feel there are real benefits to interacting with others via an SNS and are motivated to share personal information online [11]. Privacy concerns related to the use of Facebook have grown with the addition of new features and an expanding user base. Originally, Facebook membership was limited to university students, and the default privacy settings were configured to allow ‘network members’ access to user content. In 2005, Gross and Acquisti found that only 0.06% of a university network—three people—had changed the default settings [9].

Since 2006, Facebook has introduced many new sharing opportunities including photo albums, status updates, notes, etc., giving people more ways to share more personal information, and creating a fertile ground for researchers [5] (see [4] for an overview of the evolution of Facebook’s features and privacy controls). In light of the increase in the amount of content shared and the increase in the number of users, recent research results indicating that users’ actual privacy settings do not match their sharing intentions are particularly troubling [21, 18]. These results confirm those of earlier work. In 2006, Acquisti et al. found that a significant minority of users were aware of the privacy settings available [1]. A later study by Egelman et al. indicated that users have difficulty configuring Facebook privacy settings to satisfy task requirements in a laboratory setting [7]. Some of the difficulties that participants experienced were related to a failure to understand the limitations of the privacy settings.

2.1 SNS Users’ Privacy Concerns

To identify categories of SNS privacy concerns, Krasnova et al. held focus groups with university students in Berlin about their concerns with Facebook use [14]. The most frequent theme was concern over unwanted audiences viewing shared content, where the list of audiences mentioned included future employers, supervisors, family members, peers,

and subordinates. Participants also frequently mentioned “organizational threats” related to the collection and use of their data by the SNS provider and third parties. Concerns about social threats were another common theme for concerns including people purposefully posting content to harm the individual, and general concern over a lack of control over the actions of other users.

Tufekci investigated the relationship between users’ privacy concerns and their level of disclosure on an SNS, and found no relationship [30]. Even users who expressed many privacy concerns divulged large amounts of personal information on their profiles. However, the study only asked about the relatively static fields of a profile like age, sex, gender, religion, political affiliation, interests, and favorite books, rather than concerns over dynamic content (e.g., status updates, comments, etc.).

In a three year longitudinal study of university students, Lampe et al. found that users’ imagined audiences for their profiles were changing over time [15]. For example, in 2008, significantly more users expected family members had viewed their profiles compared to 2006 (an increase from 49% to 70%). Similarly, more students thought a total stranger might have viewed their profile (24% in 2008, 14% in 2006). Some of the changes in attitude can be attributed to the evolution of Facebook’s sharing features and default privacy settings.

The shift from Facebook as a social network for universities to a social network for everyone forced users to adapt to a new model of sharing: suddenly users’ friend networks included coworkers, family members, and friends from other life stages, in addition to classmates. Interested in understanding the tensions that arise from a heterogeneous friend network, Lampinen et al. conducted 20 semi-structured interviews about participants’ friend networks and their methods for managing group co-presence. They reported that many users fear that a boss or acquaintance might see something embarrassing that was not intended for them, and that users attempt to avoid these situations through self-censorship and using context to carefully selecting a suitable communication medium.

Skells and Grudin also studied the dynamics of group co-presence, but focused on SNS usage in the workplace, and found that users have trouble coping with the co-presence of coworkers and other contacts in an SNS friend network [26]. Many participants noted the burden associated with constantly maintaining an awareness that the two groups are present in their audience. Participants also noted the need to limit access to select content based on relationship.

These studies provide a strong foundation for the observation that protecting content from unwanted audiences is more than simply a matter of preventing strangers from accessing profiles. However, we are unaware of any large-scale studies that have attempted to quantify the extent to which users are sharing content inappropriately with members of their friend networks through the use of users’ previously posted content or questions about specific friends. We also build upon previous work by recruiting a more generalizable sample, rather than members of a particular institution.

2.2 Strategies for Mitigating Privacy Concerns

In terms of users’ strategies for mitigating their privacy concerns, SNS users regulate their interactions with others using many techniques and not all are based on the official

privacy controls. Young and Quan-Haase identified boundary regulation mechanisms that include deleting tags, and using direct messages to limit audiences [31]. Stutzman and Kramer-Duffield found that users who employed supplemental privacy preserving behaviors, like curating the posts on their wall and collaboratively adjusting SNS behavior among friends, were more likely to have a “friends only” profile [28].

Several papers have reported that users cope with conflicting social spheres by maintaining separate profiles, limiting access to subsets of the friend network, carefully selecting a communication medium, or using separate SNSs for different audiences [29, 26]. PEW Internet reports that in 2011, 63% of Facebook users had removed someone from their friend network [20], an increase compared to the 56% of users who reported to have “unfriended” someone in 2009. The same survey found deleting and untagging posts to be common among all user demographics.

Some users resort to changing their offline behavior to mitigate their privacy concerns. In a study of sharing photos in an SNS, Besmer and Lipford found that users adjusted their offline and online behavior to mitigate their privacy concerns: participants reported avoiding having their pictures taken in the first place, untagging photos, and asking friends to remove photos rather than adjusting privacy settings [3].

2.3 Summary

Usable privacy controls are critical to SNS users’ boundary regulation process. While it may be possible for some users to achieve their desired privacy without the help of technical mechanisms, it is unlikely that this is the case for all users considering the wide range of privacy concerns and the overhead involved with using ad hoc techniques. Usable privacy controls are needed, but first a thorough understanding of users’ privacy concerns is necessary such that the design can optimize the number of concerns addressed and the number of users who benefit.

Prior work leaves an important question unanswered— which privacy concerns are rampant enough that they ought to be designed for in the controls, and which mitigating behaviors are prevalent enough to motivate the design of new privacy controls? Prior work has demonstrated the wide range of users’ privacy concerns and that users manage their concerns through a number of techniques other than the use of the access controls. This suggests that the existing access controls can be improved. However, it is unreasonable to expect that an access control mechanism will prevent all users from ever sharing content inappropriately. Therefore, we need metrics to determine how often problems currently occur and what would be an acceptable failure rate [6]. In our study, we attempted to answer the former.

3. METHOD

We collected data on Facebook users’ interpersonal privacy concerns with regard to specific subgroups of their friend networks. We also collected data on users’ strategies for mitigating privacy concerns. We chose to focus on Facebook based on the functionality of the API and because of the large user base. We instrumented the survey as a Facebook application; this enabled us to pose questions using real profile data. We were specifically interested in how users manage their friend networks, their use of Facebook’s privacy features, and whether users had privacy concerns surrounding their posted content (e.g., photos and comments).

3.1 Survey Content

The survey had three sections. In the first section, we asked participants general questions about their Facebook usage so that we could compute correlations with real and perceived privacy risks. In the second section, we asked participants to report their level of concern with general scenarios describing situations with common unwanted audiences. Finally, in the third section, we used the API to ask questions about individual Facebook friends and shared posts.

3.1.1 General Usage

We asked about participants’ Facebook habits to measure the activities users engage in most often, the amount of time spent on each activity, the relationship between the user and the people in their friend networks, which privacy features are used, and whether other means of controlling access to information are employed.

3.1.2 Concerns with Unwanted Audiences

Previous work asked users to report the perceived likelihood of specific audiences viewing their profiles (e.g., employers, law enforcement, thieves, political parties, or sexual predators) [31]. We reused many of the scenarios that were used by Young and Quan-Haase, but instead of asking participants to guess the likelihood that the scenario was already occurring, we asked participants to rate their level of concern—unconcerned to concerned on a 5-point Likert scale—that “each scenario could happen by using Facebook.” We asked these questions to examine participants’ levels of concern in the general sense before asking similar questions about specific posts randomly selected from their profiles.

3.1.3 Incorporating Profile Data

In the final section of the survey, we used the Facebook API to select content from participants’ profiles and ask questions that incorporated that content. The questions were designed to ascertain the composition of participants’ friend networks and the perceived level of sensitivity of content that they and others had posted to their profiles. We designed this section to help us identify specific instances in which participants’ posts were inappropriately being viewed by members of their friend networks.

For instance, if a participant indicated that she would not want coworkers to see a specific photo, and her friends network included coworkers, this may indicate a situation where fine-grained control is needed to manage access to content within her friend network, particularly if access to the photo was not restricted to anything more granular than friends. In this manner, we attempted to quantify the frequency with which privacy violations may be occurring. That is, previous work has focused on qualitatively describing privacy concerns, whereas we were interested in measuring the likelihood with which these concerns come to fruition.

For each participant, we asked questions about nine randomly selected friends to gain an understanding of how Facebook users know the members of their friend networks, as well as to measure how much they trust their friends with access to their profile information. For each of these friends, we asked the following questions:

1. What is your relationship to *FRIEND-NAME*?
2. How do you feel about *FRIEND-NAME* viewing all the information you have uploaded to Facebook?

	General	Friends	Posts
Member of your immediate family	✓	✓	✓
Member of your extended family	✓	✓	✓
Coworker	✓	✓	✓
Someone you know from high school, college, or grad school	✓	✓	✓
Friend of a friend	✓	✓	✓
Someone you have not met in person	✓	✓	✓
Someone you socialize with in person		✓	✓
Not sure		✓	
Stranger			✓

Table 1: We asked about common Facebook audiences throughout the survey. ‘General’ shows the groups used for a question about general friend network composition. ‘Friends’ shows the groups presented for the classification of individual friends. ‘Posts’ shows the groups used in questions about individual posts.

	Sample	Facebook
Age		
18-24	16%	25%
25-34	48%	25%
35-54	31%	30%
55+	5%	11%
Gender		
female	75%	55%
male	25%	43%

Table 2: Comparison of our sample’s demographics to the demographics reported by iStrategyLabs.com. The table shows only the age groups that are present in our sample, and iStrategyLab’s numbers for gender total 98% (2% of users were recorded as unknown).

To determine whether participants were being diligent in describing their friends, we also asked these questions about a fictitious friend whose profile picture we took from a free stock photo archive. Thus, we asked these questions for ten friends, nine of whom were actually members of their friend networks, the tenth was a randomly assigned male or female fictitious person.¹

To understand the type of content a user might be uncomfortable sharing and why, ten posts were randomly selected from the participant’s profile. We posed eight questions to measure their level of comfort with sharing each post. The audiences used for this set of questions are listed in the last column of Table 1.

3.2 Participants

We recruited participants via ResearchMatch, a website that pairs researchers with potential participants.² The recruiting email did not mention privacy, it requested “Facebook users to take a twenty minute survey on their Facebook usage habits.” As compensation, participants were entered

¹The fake profile picture appeared as the fifth of ten. We observed that 83.1% of participants correctly answered that they did not know this person (95% CI [78.0, 87.2]), though found no correlations with demographics or Facebook usage and thus did not analyze this further.

²www.researchmatch.org/about/




Figure 1: Estimated time spent per week.

in a drawing for one of five \$100 gift cards. We received completed surveys from 260 respondents, ages ranged from 18-62 ($\mu = 33.8$, $\sigma = 10.6$).

We conducted the study remotely, and communicated with participants by email and the study application.³ The recruiting material specified that participants would be required to install a Facebook application in order to participate. Prior to encountering the installation dialog, all potential participants reviewed a consent form describing the research and our data collection and storage policies. We minimized risk to participants by designing the application to require only the permissions necessary to execute the study, minimizing the amount of data collected, and retaining data for the shortest time possible (e.g. we retained answers to the survey questions but only temporarily stored the Facebook identifiers of friends and posts to ensure uniqueness). We also provided uninstallation instructions upon completion of the study.

Facebook does not publish detailed demographic data and so we rely on the statistics reported by iStrategyLabs [10]. Based on their most recent demographics report, released in June 2010, our sample closely resembles the larger Facebook population of users. However, we underrepresent the 18-24 group, and overrepresent the 25-34 age group (Table 2). Our underrepresentation of younger users is in contrast to the related work focused on youths or undergraduates (e.g. [1, 4, 7, 21, 28]). We restricted participation to users in the United States, and our sample represents users from several states including New York (25% of 260), Alabama (13%), Minnesota (10%), and California (9%). Approximately 50% of our sample had completed at least some college.

Most of our participants reported using Facebook several times a day (68.8% of 260), while very few participants said they log in less than once per week (5% of 260). Most of the participants have had their Facebook account for more than two years (77.3% of 260). We also asked about the amount of time spent on specific activities: reading the newsfeed, creating new posts, or browsing friends’ profiles (see Figure 1). In general, participants spend more time consuming content than they do creating content (see Figure 2).

4. RESULTS

We collected data from February to April 2011, to examine Facebook users’ privacy concerns and privacy-preserving strategies. Overall, we observed that the most concern-

³Columbia University IRB protocol #AAAI1077.




Figure 2: Responses to, “How frequently do you use Facebook to ...?” Answers correspond to: look up information about a friend, upload photos, view friends’ photos, browse profiles of people you don’t know, and find new friends.

ing threat to participants’ privacy comes from fears about strangers viewing their profiles. However, we observed that the vast majority of participants successfully navigate the privacy settings interface in order to mitigate this concern. Participants raised concerns about certain subgroups of their friend networks being able to view inappropriate content, but were less prepared to deal with these threats: problems arise from the intermingling of “friends” who are associated with differing contexts (e.g., work, family, etc.). Participants currently use a variety of ad hoc approaches that are unlikely to completely address their concerns.

In this section we first present our results in terms of participants’ strategies to prevent strangers from viewing their content, concerns about disclosures to strangers, and specific examples of content that participants would not want strangers to view. Next, we examine what we call “the insider threat,” which involves inappropriately sharing content with members of the friend network. In this context, we present the compositions of participants’ friend networks, their concerns with regard to sharing content with specific subgroups of their friend networks, and the strategies they employ to address their concerns.

4.1 Stranger Danger

4.1.1 Strategies to Block Strangers

We used the survey application to check the amount of profile information that was viewable to all Facebook users (i.e., the number of users not using any access control settings). The application checked the visibility of each participant’s wall (e.g., status updates and comments), the people in her friend network, and her photos. We found that across the sample:

- 14.2% had a public wall (e.g., status updates and comments).
- 6.5% had a public photo album.
- 53.8% had the list of people in their friend network public.

Almost half of our participants (45.4% of 260) had no information accessible to strangers. The application was not able to measure whether access was further restricted beyond complete strangers (e.g., whether certain subsets of friends were also prohibited from viewing certain content or if ‘Friends of Friends’ were prohibited from viewing content). Additionally, since Facebook’s default privacy settings have changed over time, we cannot definitively say whether participants’ had actively blocked strangers from accessing this content or if it could be partially attributed to changes in default settings. However, these numbers do indicate that the vast majority of participants (94.6% of 260) have either photos or posts blocked from strangers; 84.6% had both photos and posts blocked from strangers.

4.1.2 Concerns with Broad Scenarios

We measured general privacy concerns using ten scenarios about unwanted audiences and asked participants to indicate on a 5-point Likert scale their level of concern that each could happen as a result of using Facebook (the markers were “unconcerned,” “neutral,” and “concerned”). We measured participants’ level of concern that each *could* happen, as opposed to prior work that used similar scenarios to measure users’ belief that the scenarios were already taking place [31]. The set of concerns that involved profile access by strangers (i.e., people who are not members of the friend network) are depicted in Table 3. The table also presents the percentage of participants who reported being concerned with each scenario,⁴ as well as the median ranking from the Likert scale. Twenty-eight participants (10.8% of 260) reported being unconcerned with any of the scenarios involving strangers (85.7% of those participants had a private profile—neither photos nor walls were accessible to strangers). We observed no statistically significant correlations between participants’ concerns for the scenarios and whether they were mitigating them through the use of private profiles. We hypothesize that this may be because only four participants (1.5% of 260) were not concerned by any of the scenarios nor had private profiles.

4.1.3 Specific Concerns

The aforementioned scenarios, while plausible from various media accounts, are unlikely to affect most users. To create a more realistic view of how often Facebook users share content inappropriately, we showed participants ten random pieces of content that they had previously posted to their profiles. These included comments, photos, and status updates. For each piece of content, we asked participants to rate how concerned they would be if a stranger were to view it. As before, answers were reported on a 5-point Likert scale that ranged from “concerned” to “unconcerned,” with “indifferent” as the neutral option.

We observed that only 48 participants (18.5% of 260) were unconcerned with sharing all ten posts with a complete stranger. On average, each participant was concerned with 51.6% of the ten posts. Upon performing a Pearson correlation between participants’ mean levels of concern averaged over the ten posts and whether or not their posts were private, we observed a statistically significant negative correlation ($r = -0.141$, $p < 0.023$). Thus, participants whose posts were private had lower average rates of concern than participants whose posts were accessible to strangers.

⁴We define *concern* as reporting a 4 or 5.

Scenario	Concerned	M
1. Thieves using Facebook to track, monitor, locate, and identify you as a potential victim.	68.8%	4
2. Your employer seeing an inappropriate photo or comment on your profile.	62.7%	4
3. Your employer using your profile to assess your suitability for the company.	55.0%	4
4. Sexual predators using Facebook to track, monitor, locate, and identify you as a potential victim.	51.9%	4
5. Your employer using Facebook to monitor your conduct while you're at work.	46.2%	3
6. Your employer using Facebook to monitor your conduct while you're away from work.	44.6%	3
7. A stranger will see an inappropriate photo or comment on your profile.	40.8%	3
8. Political parties using Facebook to target you through the use of ads and data mining.	30.4%	3
9. Your university using Facebook to identify you as a university code violator.	20.0%	1
10. Law enforcement using Facebook to track drug use and other illegal activities.	17.3%	1

Table 3: We asked participants to rate their level of concern that each scenario could occur on a 5-point Likert scale from “unconcerned” to “concerned,” with “indifferent” as the neutral option. The second column reports the percentage of the sample that was concerned about each scenario, while the third represents the median rating for each question.

One interpretation of this is that participants who knew that strangers could not access their profiles were therefore less concerned about the likelihood of it happening. The corollary to this is that participants who did not use privacy controls were more concerned about the threat of strangers.

4.2 The Insider Threat

4.2.1 Friend Network Composition

Our sample’s average friend network size was 357 friends (median = 291, range = [22, 3280], $\sigma = 319.5$). This is much larger than the statistic Facebook reports: the average user has 130 friends [8]. It is likely that this discrepancy is indicative of a very long tail; Facebook reports the average network size for all users—including users who add a handful of friends and never access their accounts again—whereas we limited our sample to only active users. Our numbers are consistent with other academic studies of active Facebook users. For instance, Kelley et al. reported a median of 222 friends [13], Young and Quan-Haase reported an average of 401 friends, and Stutzman and Kramer-Duffield reported an average of over 400 friends [28].

To get a general sense of friend network composition, we asked participants, “Which of these groups are you friends with on Facebook?” and instructed them to select all options that applied to their friend networks. The choices for this question are presented in the first column of Table 1. As shown in Figure 3, a minority of participants selected people that they had not met in person, whereas the other five groups were each selected by over 80% of our participants.

For a broad overview, we asked each participant, “What percentage of your friend network do you trust with access to your profile and shared information?” The choices were presented in increments of ten, from 0-100%. On average,




Figure 3: Responses to, “Which of these groups are you friends with on Facebook?”

participants claimed to trust 75.4% ($\sigma = 26.3$, median = 90%) of their friend networks. Of the 55 people who answered 50% or less, only 35 (63.6% of 55) of them claimed to have modified their privacy settings so that some of their friends have limited access to their profiles.

We validated participants’ perceptions about their friend network composition by asking them to categorize a random sampling of their friends. The survey randomly selected nine people from each participant’s friend network and asked questions about each selected friend. For example, if *Alice Smith* was selected, the participant was shown Alice’s profile picture and was asked, “What is your relationship to *Alice Smith*?” We asked each participant to select one category from those listed in the ‘Friends’ column of Table 1.

A plurality of the friends were reported to be known from school (42.6% of 2,340). The remaining groups were also chosen at least once, though immediate family members and *not sure* were chosen least frequently.⁵ Based on the categorization of the nine friends, we can estimate the average composition of participants’ friend networks based on the frequency that each group was selected (see the ‘Frequency’ column of Table 4). We hypothesized that while labeling individual friends, participants might discover their friend networks contained more groups than they remembered. However, we found that this was not the case; none of the nine randomly selected friends was a member of a group that a participant had not already selected in the first part of the survey. Thus, participants were by and large aware of the composition of their friend networks.

For each of the nine friends we asked participants to categorize, we also asked how they felt about that friend viewing all the content they had uploaded to Facebook. Participants responded using a 5-point Likert scale that ranged from “uncomfortable” to “comfortable,” with “indifferent” as the neutral option. We provided participants with a text box to optionally explain their responses when they selected uncomfortable or slightly uncomfortable. We define “unopposed to sharing” throughout the rest of this paper to mean participants who answered with either “comfortable,” “slightly comfortable,” or “indifferent;” we consider those who answered with either “uncomfortable” or “slightly un-

⁵This says more about the *not sure* category, as immediate family members were likely selected infrequently because each participant was likely to have only a limited number of immediate family members in real life, relatively speaking.

	Frequency	Participants	Comfort
School	42.6%	88%	97.0%
Socialize with	15.4%	57.3%	98.9%
Friend of a friend	12.4%	62.7%	97.0%
Coworker	11.1%	45%	96.9%
Extended family	9.4%	48.5%	95.4%
Have not met	5.3%	20%	95.2%
Immediate family	2.1%	14.2%	98.0%
Not sure	1.7%	13%	75.0%

Table 4: Summary of responses about individual friends. ‘Frequency’ shows the number of times each group was selected across the sampling of 2,340 friends (i.e., nine friends for each of 260 participants). ‘Participants’ shows the percentage of participants represented in the frequency column. ‘Comfort’ shows the percentage of selected friends that the participant is unopposed to sharing with in that group.

comfortable” as being opposed to sharing. The majority of participants were unopposed to sharing with all nine of the selected friends (79.2% of 260). Participants indicated that they would be opposed to sharing at least some of their profile with 3.3% of the 2,340 selected friends (this number corresponds to 54 unique participants). When a participant indicated they were opposed to sharing with a specific friend, we asked a follow-up question to prompt an explanation. We present and discuss this data in Section 5.

We asked participants to rate their levels of concern that two additional scenarios may happen, similar to those presented in Section 4.1.2. While the first ten scenarios were centered around strangers—people unlikely to appear in participants’ friend networks—the additional scenarios focused on concerns with sharing inappropriate content with known recipients: family members and coworkers. In Section 4.1.2, we asked about “employers,” whereas here we discuss coworkers. We intended for the distinction between coworkers and employers to be that the latter are in management positions (i.e., have the ability to hire and fire), and therefore not apart of participants’ social circles. We cannot say with certainty that this distinction was apparent to all participants. However, McNemar’s test between participants’ concern levels between when a coworker and an employer see “an inappropriate photo or comment on your profile” yielded statistically significant differences ($\chi^2 = 9.50, p < 0.002$). This indicates that participants viewed these two groups differently.

We observed that participants claimed to be significantly more concerned with the prospect of coworkers viewing content than family members ($\chi^2 = 5.80, p < 0.016$). We performed Phi correlations to examine whether participants’ concerns over sharing content with these two groups were correlated with having members of these groups included in their friend networks, but found no significant correlations when examining both coworkers and family members. If there is a correlation, it is too small for us to observe among our relatively large sample. In either case, since roughly half of the participants indicated they would be concerned by these scenarios (Table 5), they choose to mitigate them in ways beyond preventing family members and coworkers from being included in friend networks.

Scenario	Concerned	M
1. A coworker seeing an inappropriate photo or comment on your profile.	55.0%	4
2. A family member will see an inappropriate photo or comment on my profile.	46.5%	3

Table 5: We asked participants to rate their level of concern that each scenario could occur on a 5-point Likert scale from “unconcerned” to “concerned,” with “indifferent” as the neutral option. The second column reports the percentage of the sample that was concerned about each scenario, while the third represents the median rating for each question.

	Posts	Participants
Immediate family	99.0%	91.2%
Socialize with	98.9%	96.5%
Extended family	98.9%	91.2%
School	98.8%	95.0%
Friend of a friend	97.2%	90.0%
Coworker	97.0%	83.9%
Have not met	91.6%	72.3%
Stranger	84.4%	55.4%

Table 6: Responses to questions about individual posts. Columns depict percentage of posts (of 2,600) participants were unopposed to sharing with the given groups, as well as the percentage of participants (of 260) who were unopposed to sharing all ten posts with the given groups.

4.2.2 Concerns over Specific Content

Access control decisions are typically phrased in terms of who can access a particular resource. For this reason, we also asked questions about sharing preferences based on specific posts. As we explained in Section 4.1.3, we randomly selected ten posts from each participant’s profile and asked questions to measure how they perceived the sensitivity of the content. For each post, we asked eight questions of the form, “How would you feel if *group* saw this?” The groups are listed in the last column of Table 1. Again, participants responded using a 5-point Likert scale from “uncomfortable” to “comfortable,” with “indifferent” as the neutral option. We asked participants for an optional explanation when they selected uncomfortable or slightly uncomfortable.

The groups with whom participants were least comfortable were strangers, people they had not met in person, coworkers, and those who were a friend of a friend, respectively. Table 6 depicts participants’ levels of comfort sharing their ten posts with the various groups. The first column of the table depicts the total percentage of posts (of 2,600) that participants reported that they were unopposed to sharing with each group. However, these numbers by themselves do not give an accurate representation of the number of participants who have posted sensitive content. Thus, the second column shows this information on a per-participant basis. Specifically, this shows the percentage of participants (of 260) who were unopposed to sharing all ten randomly selected posts with each of the eight groups.

We compared the number of participants who were opposed versus unopposed to sharing with each of the eight groups using McNemar’s test across each pair of groups.

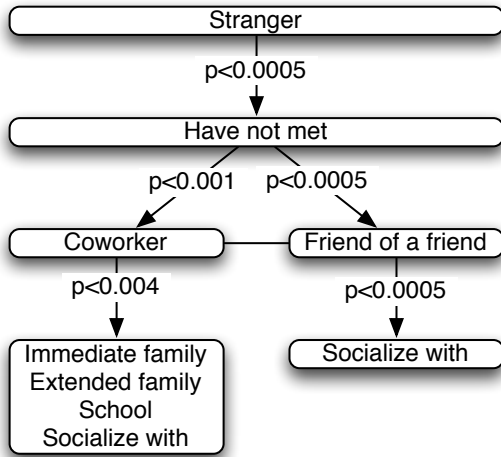


Figure 4: The hierarchy of participants’ comfort for sharing all ten randomly selected posts.

We then applied the Holm-Bonferroni correction to account for multiple tests. Based on the results of these tests, we were able to partition the eight groups based on significant differences with regard to participants’ comfort sharing all ten posts. This hierarchy can be seen in Figure 4. As expected, participants were significantly less comfortable sharing with complete strangers than any other group, whereas they were most comfortable sharing with people with whom they socialize. Because we found no significant differences among our sample between sharing with immediate family, extended family, people from school, and people with whom participants socialize, as well as no significant differences between a coworker and a friend of a friend, we consolidated our hierarchy into four discrete levels.

We took this analysis a step further in order to detect the frequency with which participants’ concerns were being realized: we examined whether participants who were concerned about posts being viewed by specific groups also included members of those groups in their friend networks. For instance, if a participant indicated that it would be inappropriate to share some of the ten aforementioned posts with coworkers, we also examined whether that participant had also categorized one of the nine aforementioned friends as being a coworker. We performed this analysis individually for coworkers, immediate family, and extended family. We found that:

- 14 of 117 (12.0%) participants whose nine friends included one *coworker* were uncomfortable sharing at least one of the ten content items with coworkers.
- 1 of 37 (2.7%) participants whose nine friends included one *immediate family member* were uncomfortable sharing at least one of the ten content items with immediate family members.
- 9 of 125 (7.2%) participants whose nine friends included one *extended family member* were uncomfortable sharing at least one of the ten content items with extended family members.

Overall, when we examined all of the groups in Table 1 together, we observed that 43 participants (16.5% of 260) had positively identified at least one content item that they were uncomfortable sharing with members of their friend networks. Thus, we have quantitatively showed that the friend network consists of varying subgroups with whom participants are not always comfortable sharing all of their content. Unlike strangers, who participants were able to block through the use of privacy settings, members of the friend network are able to view a profile and its contents when it is set to be private. Thus, we also examined whether users were employing fine-grain privacy settings or any other strategies to restrict access to certain groups or individuals within their friend networks.

4.3 Strategies for Mitigating Concerns

We know from the results of qualitative surveys and anecdotal experience that SNS users employ a multitude of techniques for managing their interpersonal privacy in addition to using the actual privacy settings: custom lists, culling their friend network, deleting posts, untagging posts, asking friends to delete posts, and maintaining more than one account. We can build a better understanding of SNS users’ privacy management needs by collecting data on the instances in which these privacy preserving behaviors are employed. In our survey, we asked participants whether they had employed each technique, and if so, to provide an explanation.

4.3.1 Custom Friend Lists

Custom friend lists allow users to subdivide their friend networks and can be used to configure fine-grained privacy settings—to allow or deny additional access. To examine how lists are used, we accessed each participant’s set of custom friend lists using the Facebook API, and for each list asked, “Why did you create the custom list *LIST-NAME*, and what do you use it for?” Our sample yielded 555 custom friend lists: most users had created at least one list (52.3%, 95% CI [46.3, 58.3]). From the explanations provided, three themes emerged: 100 lists were for privacy (18% of 555 lists), 372 were for use with other features (67%), and 83 were for created for reasons that participants could no longer remember (15%).

Approximately a quarter of the participants had created a custom list for privacy reasons (23.8% of 260). We further categorized the privacy lists based on whether the list was created to include or exclude specific friends. Exclusive lists were created to prevent access by that group (75% of the 100 privacy lists). Inclusive lists were created to give additional access to that group (17% of the 100 privacy lists). In some cases, the intended use of the list was unclear (8% of the 100 privacy lists). Most of the lists were created to separate groups of friends or to differentiate contacts by closeness. A few users created custom lists for family members (24 lists) and coworkers (14 lists).

The majority of the custom lists were created for use with other Facebook features (67% of 555 lists, created by 136 participants). The descriptions provided were generic and did not explain how the lists were actually used (66.1% of 372 lists). Although these lists might not be used for privacy reasons, the names and descriptions of the lists provide insight to friend network composition and the user’s many social identities like shared interests, activities, and location.

Several participants created a custom list to group friends that play the same Facebook game, filter their newsfeed, and manage their chat list (14.7%, 6.7%, and 6.5% of the 372 lists, respectively).

4.3.2 Curating the Friend Network

In an SNS like Facebook, where a friend relationship is reciprocal and friends are granted additional access to content, curating the friend network can be a privacy preserving behavior. The options for curating the friend network are to deny a friend request or delete (unfriend) a person. Nearly every participant had turned down a friend request in the past (96.2% of 260), and so we asked them to select their reasons. The most common answer was ‘didn’t know the person’ (selected 211 times), followed by ‘knew the person but did not want them to have access to my profile’ (129 times).

Most participants had unfriended at least one person (69.6% of 260). A commonly selected reason was ‘because we were no longer friends in real life.’ Forty-five participants had deleted a friend because they were unsure whether they knew them. Of the forty-six explanations that were provided, twenty-five participants noted they wished to stop seeing updates from the person (“because they kept posting negative or critical things without relent”, “their political views were exactly opposite of mine and I did not agree with any of their posts”).

4.3.3 Control via Deletion

SNS users can also manage access to the data associated with their profiles by deleting or ‘untagging’ posts. We posed a set of questions related to untagging and deleting posts to measure how often they are used in a privacy preserving manner. Each question was posed in the format: *have you ever...? If yes, why?* We then asked participants to check all the reasons that applied from the options: you didn’t want anyone to see it, you didn’t want a specific person to see it, you didn’t like it, or other. We also provided a text box for the participant to describe the circumstances.

We asked participants, “Have you ever untagged yourself in a photo that was posted by a friend?” Over half (58.5% of 260) the participants answered yes, and the participants provided sixty-two descriptions. Most of the descriptions related to reputation or image management: the picture looked bad (e.g., “I was making a very unattractive face”), they did not want people to see them partying/drinking/etc., or the photo was not actually them (e.g., “I’ve been tagged in spam before”).

More than half the participants responded affirmatively when we asked if they had ever deleted a photo they had uploaded to Facebook (60.8% of 260). The explanations of the circumstances varied. In some cases a photo was deleted to satisfy the request of a friend who was also in the photo (e.g., “a relative requested I remove it”). In other cases, it was to preserve privacy (e.g., “I have taken most of the pictures of my kids off because when I think about it’s weird to me that random people I don’t know well are looking at my kids”).

About one-fifth of the participants (22.3% of 260) said that they had asked a Facebook friend to delete a photo that they were in. The most popular reason was because they ‘didn’t like it’ (e.g., “Me getting drunk at a party, not appropriate.”).

We asked participants if they had ever posted a status update or comment and later deleted it. More than half answered yes (65.4% of 260). The most popular reason was that they ‘didn’t like it’ (e.g., “I decided it was stupid.”), followed by ‘didn’t want anyone to see it.’ Some of those who selected ‘other’ explained “I’ve written posts that later seem too personal,” and “I changed the way I felt.”

4.3.4 Control via Per-Post Privacy Settings

In addition to providing global access control settings, Facebook also allows users to customize access control settings on a per-post basis through a drop-down menu at the time the post is made. We asked participants, “Have you ever changed the privacy settings for a single status update?” Then asked the 92 participants who answered yes to describe an instance where they had used the feature. From the explanations it was clear that only 47 participants had actually used the feature (18.1% of 260), the other 45 conflated the global privacy settings with the per-post feature.

We coded the situations described by the 47 participants who correctly used the feature based on whether they desired to exclude specific people, include specific people, or did not specify. Thirty-one answers were for the purpose of exclusivity, for example:

- “I have blocked family members and conservative friends from status updates they might view as inappropriate.”
- “When I was talking about my roommates, I didn’t want one of them to see it.”
- “It was hidden from a person to announce a surprise party about them.”

Fourteen answers were situations that specifically included some subset of friends (e.g., “posting a personal link regarding a vacation, I made it viewable only to a specific group”). The remaining descriptions were either unspecified or too ambiguous to categorize.

4.3.5 Multiple Accounts

Although the terms of service mandate that each person have at most one account, twenty-two participants (8.5% of 260) had two or more accounts. Among these, sixteen cited reasons related to managing social spheres, either for dividing friends and family, friends and game friends, or their professional and social lives.

5. IMPLICATIONS

Our results contribute to an understanding of Facebook users’ privacy concerns and the strategies they employ to mitigate their concerns. We found most users are concerned about strangers viewing their profiles, and that many users are also concerned with the *insider threat*—inappropriately sharing content with members of the friend network. Many users have private profiles (i.e., profiles that are only visible to friends), either by default or manually adjusting the global privacy settings, which means that strangers are unable to view posted content (e.g., status updates, photos, and comments). However, these settings do not adequately address the insider threat, and therefore these concerns likely go unmitigated. In this section we discuss the implications of protecting against the insider threat, as well as the limitations and generalizability of our work. We conclude with future work.

5.1 Users are Concerned with Strangers and Many Effectively Mitigate Their Concern

We found evidence that many users are concerned with the possibility of an outsider accessing their shared content. In fact, the broad scenarios that elicited the highest concern were those that involved audiences that were not represented in most participants' friend networks (see Table 3 and Figure 3), we also found that participants were least comfortable sharing individual posts with a 'stranger' or someone they had 'not met in person' (see Table 4).⁶

We found that 89.2% of our participants were concerned with the outsider threat (232 of 260 indicated concern for at least one of the scenarios involving strangers). This figure represents the participants who selected 'concerned' for at least one of the general scenarios in Table 3. Of this set of concerned participants, we observed that 84.5% had a private profile. Which means that overall 15.5% of our participants were highly concerned with the outsider threat but they were not managing it through the available privacy controls. Put another way, 86.2% of participants (224 of 260; 95% CI [81.4, 91.1]) were either not very concerned with the stranger threat or they were concerned and were able to address their concerns.

Facebook users are increasingly opting for a 'friends only' profile. In 2010, Stutzman and Kramer-Duffield reported more than half of their participants had a 'friends only' profile [28]. This is a significant increase in adoption compared to Gross and Acquisti's 2005 observation that 0.06% of their sample had a 'friends only' profile. Thus, the existing privacy settings interface and the default settings may reasonably address users' privacy concerns regarding strangers.

5.2 The Insider Threat Prevails Unmitigated

We identified interpersonal sharing concerns in 37% of our sample (96 of 260; 95% CI [31.0, 43.1]), this figure represents the participants who either expressed concerns about sharing with a specific friend or expressed concerns about sharing a specific post. We have reason to believe that this figure is a lower bound since the sample represents such a small portion of most users' friend networks and shared data. Even so, we can positively say that this group of participants have interpersonal sharing concerns that have been realized. Furthermore, our data on the number of users who supplement the privacy controls with ad hoc strategies suggests that the privacy controls are unsuited for dealing with this threat. It is important to understand the characteristics of users' actual privacy concerns to determine which should be addressed with privacy controls and which would be better handled through other means like self-censorship or removal.

Based on our findings, we recommend additional research to understand the taxonomy of privacy concerns experienced by users. We found that users' conceptualization of the insider threat varies depending on the context of a situation. For example, some participants described their concerns by their relationship with a person while others described them

⁶The option 'not met in person' was used in two sections of the survey: reporting friend network composition and in the section on individual posts. We intended it to describe a person known online but not in the real world, however, users understood the intended meaning in the first section but did not in the posts section. Based on the explanations offered, many equated 'not met in person' with 'stranger' in the posts section.

based on the content of posts. In situations where participants were uncomfortable with specific friends, some participants described their discomfort as a general distrust of the person, while others described specific types of content they would not want the person to see. Also, in some cases the participants described their concern in terms of their intended audience, while others described their concern by the people that should be excluded. Participants almost never described their concerns in terms of broad groups, though this could be an artifact of our question style. In Section 4.3.2, we provide sample participant responses to illustrate users' reasons for deleting content and using per-post privacy settings.

Facebook users could use custom friend lists to address the insider threat. However, based on our data and prior work, it seems this feature is largely a failure. Nearly all of our participants who had created a custom list also utilized additional privacy preserving behaviors, and many of the custom lists we recorded were not used for managing privacy concerns. Furthermore, according to prior work, although users are able to organize their friends into lists, the lists are effectively useless because the user-created lists fail to accurately capture their desired audiences for shared content [12, 13]. As mentioned above, the threat model has changed such that now the problem consists of edge cases that are highly contextual. One reason custom lists do not address the problem is that they are created a priori, before the user thinks about the situations in which they will be used [13]. We note this problem is most likely present in other SNSs that rely on audience management, such as Google Plus.

5.3 Generalizability

We believe that our sample more accurately reflects the current demographics of Facebook users than most prior work on this topic. Despite significant changes in the demographics of Facebook users, most of the empirical research on Facebook users' privacy concerns has been limited to undergraduates or teenagers [1, 4, 7, 14, 17, 21, 28, 30]. Notable exceptions include two studies on group co-presence [16, 26]. In the early days of Facebook, when membership was limited to university students, studying undergraduates made sense. However since 2009, Facebook has become popular with other demographics as well [27]. College-aged individuals now represent a minority of the Facebook population [10].

5.4 Limitations

The numbers we present in this paper are likely inexact due to several confounding factors. First, we only asked participants about a limited number of their friends and posts. Thus, participants' concerns for these likely represent lower bounds. Additionally, we only asked participants about their levels of concern, rather than their perceived likelihood of negative outcomes or whether they regretted sharing specific content. Thus, while participants may have concerns, it is unclear under what circumstances these concerns may rise to the level of altering behavior.

It is likely that our method inherently introduced bias: at least two people refused to participate because of the use of a Facebook application. We cannot estimate how many others chose not to participate for similar reasons. Thus, it would at first seem that our sample is biased toward users who are unconcerned with privacy. While our sample might not

include the users most concerned with privacy, the number of privacy concerns and mitigation strategies recorded indicate that our participants had very clear privacy concerns. It is possible that without this bias, our sample would reflect even stronger privacy concerns.

5.5 Future Work

The existing privacy controls fail to empower users to adequately manage their concerns because they mostly focus on strangers. Our study identifies several privacy preserving behaviors that users rely on to manage their interpersonal privacy. We suggest that familiarity with the strategies used and when they are employed will help researchers identify specific weaknesses to address in future privacy controls. We do not mean to imply that privacy controls must replace all privacy preserving behaviors or that the privacy controls are the only way to manage interpersonal privacy. Rather, we believe SNS users would benefit from improved controls that match their needs.

By measuring Facebook users' interpersonal privacy concerns we learned that the insider threat is a significant concern for many users, but not a concern for all. As a result, we collected detailed data about the insider threat from only a subset of our sample. Our data suggest that users' conceptualization of the insider threat is individualized and dependent on context, but they do not reveal generalizable themes. A follow-up study might take a similar approach to measuring the threat and choose to focus on participants who are concerned with the insider threat.

As we mentioned previously, we believe our estimate for the number of people who are concerned with the insider threat is a lower bound. A follow-up study could test this by using a similar methodology and asking questions about a larger sample of the participants' friend networks or a larger sample of the participants' shared content. In most cases, our sampling of the friend networks represented less than 3% of a user's friends. For the purpose of understanding the insider threat, it may be useful to devise a way to select friends of interest or shared items that are most likely to be problematic.

We found that the threat model has evolved, which suggests that our approach to the problem needs to evolve as well. In the past, we measured the usability of access control interfaces in strict terms: How many unintended parties can view the content? How many intended audiences cannot view the content? These metrics may have served us well in the past, but applying it as we move forward would be a mistake. One complication is that with the insider threat, the threat is dynamic, unlike the more static stranger threat; the appropriateness of the audience is highly contextual. Moreover, because of the complexity of this problem, we will never have an interface that provides perfect coverage in all situations, for all users. Instead, we should focus on designing fine-grained controls that work for *most* users, *most* of the time. Ideally a mechanism that achieves this goal would also effectively communicate its limitations and promote alternative privacy preserving behaviors that mitigate users' residual concerns.

6. CONCLUSION

Our survey contributes to human-centered design efforts to correct the inadequacies of existing SNS privacy settings. Prior work has shown that users' privacy settings do not

match their desired level of privacy and that more usable access control mechanisms are needed. We quantify the problems that users have when they attempt to keep certain profile information private on Facebook. Specifically, we found that 86.2% of participants were either unconcerned with the threat of strangers viewing their content, or they were able to mitigate those concerns through the use of global privacy settings. Thus, we believe that strangers are no longer the greatest threat. Instead, our data indicate that threats from within users' friend networks are more concerning because they are much less likely to be mitigated through the use of privacy settings. We observed that 37% of our 260 participants indicated concern with allowing specific friends to view their profiles or with showing certain posted content to certain groups of people among their friend lists. The existing privacy settings do not address these types of threat, which is why we observed users performing ad hoc mitigations, such as self-censorship or culling the friend network. While no access control system in this context will ever address all privacy concerns, our results indicate that existing systems must be improved to better address emergent threats.

7. ACKNOWLEDGMENTS

This work was partially supported by NSF Grant CNS-0905403 and Intel, through the ISTC for Secure Computing.

8. REFERENCES

- [1] A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy Enhancing Technologies*, pages 36–58. Springer Berlin, 2006.
- [2] I. Altman. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Brooks/Cole Pub. Co., 1975.
- [3] A. Besmer and H. Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *CHI 2010*, pages 1563–1572. ACM, 2010.
- [4] d. boyd and E. Hargittai. Facebook privacy settings: Who cares? *First Monday*, 15(8), August 2010.
- [5] d. m. boyd and N. B. Ellison. Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13:210–230, 2007.
- [6] S. Egelman and M. Johnson. How good is good enough? the sisyphian struggle for optimal privacy settings. In *Proceedings of the CSCW Reconciling Privacy with Social Media Workshop*, 2012.
- [7] S. Egelman, A. Oates, and S. Krishnamurthi. Oops, I did it again: Mitigating repeated access control errors on Facebook. In *CHI 2011*, pages 2295–2304, 2011.
- [8] Facebook. <http://newsroom.fb.com/content/default.aspx?NewsAreaId=22>, 2012.
- [9] R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *Proceedings of Workshop on Privacy in the Electronic Society*, pages 71–80. ACM, 2005.
- [10] istrategylabs. Facebook demographics and statistics report. <http://www.istrategylabs.com/2010/06/facebook-demographics-and-statistics-report-june-2010-privacy-concerns-dont-stop-growth/>, June 2010.

- [11] A. N. Joinson. Looking at, looking up or keeping up with people?: motives and use of Facebook. In *CHI 2008*, pages 1027–1036. ACM, 2008.
- [12] S. Jones and E. O’Neill. Feasibility of structural network clustering for group-based privacy control in social networks. In *SOUPS 2010*, pages 9:1–9:13. ACM, 2010.
- [13] P. G. Kelley, R. Brewer, Y. Mayer, L. F. Cranor, and N. Sadeh. An investigation into Facebook friend grouping. In *INTERACT ’11*, pages 216–233. Springer-Verlag, 2011.
- [14] H. Krasnova, O. Günther, S. Spiekermann, and K. Koroleva. Privacy concerns and identity in online social networks. *Identity in the Information Society*, 2:39–63, 2009.
- [15] C. Lampe, N. B. Ellison, and C. Steinfield. Changes in use and perception of Facebook. In *CSCW ’08*, pages 721–730, New York, NY, USA, 2008. ACM.
- [16] A. Lampinen, S. Tamminen, and A. Oulasvirta. All my people right here, right now: management of group co-presence on a social networking site. In *GROUP 2009*, pages 281–290. ACM, 2009.
- [17] K. Lewis, J. Kaufman, and N. Christakis. The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication*, 14(1):79–100, 2008.
- [18] Y. Liu, K. Gummadi, B. Krishnamurthy, and A. Mislove. Analyzing Facebook privacy settings: User expectations vs. reality. In *Proc. of Internet Measurement Conference (IMC)*. ACM, 2011.
- [19] S. Lohr. How privacy vanishes online, March 17 2010. <http://www.nytimes.com/2010/03/17/technology/17privacy.html>.
- [20] M. Madden. Privacy management on social media sites. <http://pewinternet.org/Reports/2012/Privacy-management-on-social-media.aspx>, February 2012.
- [21] M. Madejski, M. Johnson, and S. M. Bellovin. A study of privacy settings errors in an online social network. In *Proceedings of the 4th IEEE International Workshop on Security and Social Networking, SESOC ’12*, 2012.
- [22] Nielsen. Social media report, 2011. <http://blog.nielsen.com/nielsenwire/social/>.
- [23] L. Palen and P. Dourish. Unpacking “privacy” for a networked world. In *CHI ’03*, pages 129–136. ACM New York, NY, USA, April 5–10 2003.
- [24] R. W. Reeder and R. A. Maxion. User interface dependability through goal-error prevention. In *Proc. of the International Conference on Dependable Systems and Networks 2005*, pages 60–69. IEEE, 2005.
- [25] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, Sept. 1975.
- [26] M. Skeels and J. Grudin. When social networks cross boundaries: a case study of workplace use of Facebook and LinkedIn. In *GROUP 2009*, pages 95–104. ACM, 2009.
- [27] J. Smith. Fastest growing demographic on Facebook: Women over 55, February 2 2009. <http://www.insidefacebook.com/2009/02/02/fastest-growing-demographic-on-facebook-women-over-55/>.
- [28] F. Stutzman and J. Kramer-Duffield. Friends only: examining a privacy-enhancing behavior in Facebook. In *CHI 2010*, pages 1553–1562. ACM, 2010.
- [29] F. D. Stutzman and W. Hartzog. Boundary Regulation in Social Media. *SSRN eLibrary*, 2009.
- [30] Z. Tufekci. Can you see me now? audience and disclosure regulation in online social network sites. *Bulletin of Science, Technology & Society*, 28(1):20–36, 2008.
- [31] A. L. Young and A. Quan-Haase. Information revelation and Internet privacy concerns on social network sites: a case study of Facebook. In *Proceedings of the 4th International Conference on Communities and Technologies, C&T ’09*, pages 265–274. ACM, 2009.

APPENDIX

A. QUESTIONS ABOUT FACEBOOK USAGE AND PRIVACY PRESERVING BEHAVIORS

1. How long have you had a Facebook account?
 - Less than 1 year
 - Between 1 and 2 years
 - More than 2 years
2. How often do you use Facebook?
 - Several times a day
 - Once a day
 - Once every few days
 - Once a week
 - Once a month
 - Less than once a month
3. About how much time do you spend on Facebook reading your news feed each week?
 - Less than 1 hour
 - Between 1 and 2 hours
 - Between 2 and 4 hours
 - 4 hours or more
4. About how much time do you spend on Facebook posting information and updating your profile each week?
 - Less than 1 hour
 - Between 1 and 2 hours
 - Between 2 and 4 hours
 - 4 hours or more
5. About how much time do you spend on Facebook browsing your friend's profiles or photos each week?
 - Less than 1 hour
 - Between 1 and 2 hours
 - Between 2 and 4 hours
 - 4 hours or more
6. How frequently do you use Facebook for the following:
 - To look up information about a friend
 - To communicate with a friend
 - To upload photos
 - To view photos your friends uploaded
 - To share a link to a news story
 - To browse the profiles of people that are friends with your Facebook friends
 - To browse the profiles of people that you do not know
 - To find new friends
7. Are you Facebook friends with:
 - Members of your immediate family (parents/siblings)
 - Members of your extended family
 - Coworkers
 - People you know from high school/college/grad school
 - People you met through friends
 - People you have not met in person
8. How many Facebook friends do you have? If you're unsure, an estimate is fine.
 - 0-99
 - 100-299
 - 300-599
 - 600 or more
9. Do you have more than one Facebook account?
 - yes
 - no
10. How many accounts do you have?
 - 2
 - 3
 - 4
11. Have you ever used the option to change the privacy settings of a single status update?
 - yes
 - no
 - Please describe an instance where you changed the settings of a single update.
12. Have you ever turned down a friend request?
 - yes
 - no
13. Why did you turn down the request?
 - You did not know the person
 - You knew the person but did not want them to see your profile
 - You suspected the profile was fake
14. Have you ever unfriended someone?
 - yes
 - no
15. Why did you unfriend them?
 - You were no longer friends in real life
 - You did not want to share your Facebook profile with them any longer
 - You were unsure whether you knew them
16. Have you ever sent a friend request to someone you did not know in person?
 - yes
 - no
17. Why did you friend them?
 - They were a friend of one of your Facebook friends
 - You had common interests
 - You were interested in meeting them offline
18. Have you changed your privacy settings such that some of your Facebook friends have limited access to your profile?
 - yes
 - no
19. How do you know the friends that have limited access to your profile?
 - Members of your immediate family (parents/siblings)
 - Members of your extended family
 - Coworkers

- People you know from high school/college/grad school with
 - People who you met through friends
20. Have you ever untagged yourself in a photo that was posted by a friend?
 - yes
 - no
 21. Why did you untag yourself in your friend's photo?
 - You didn't want anyone to see it
 - You didn't want a specific person to see it
 - You didn't like it
 22. Have you ever deleted a photo you uploaded to Facebook?
 - yes
 - no
 23. Why did you delete the photo you uploaded?
 - You didn't want anyone to see it
 - You didn't want a specific person to see it
 - You didn't like it
 24. Have you ever asked a Facebook friend to delete a photo they uploaded of you?
 - yes
 - no
 25. Why did you ask your Facebook friend to delete the photo?
 - You didn't want anyone to see it
 - You didn't want a specific person to see it
 - You didn't like it
 26. Have you ever posted a status update or comment and deleted it later?
 - yes
 - no
 27. Why did you choose to delete the post?
 - You didn't want anyone to see it
 - You didn't want a specific person to see it
 - You didn't like it
 28. Have you ever deleted a comment that was posted by a Facebook friend?
 - yes
 - no
 29. Why did you delete the comment?
 - You didn't want anyone to see it
 - You didn't want a specific person to see it
 - You didn't like it
 - You suspected the comment was spam

B. GENERAL SCENARIOS

The study application presented the following questions one at a time as Likert items on a 5-point scale with the anchor points unconcerned, indifferent, and concerned.

1. A stranger will see an inappropriate photo or comment on my profile.

2. A family member will see an inappropriate photo or comment on my profile.
3. A coworker will see an inappropriate photo or comment on my profile.
4. An employer will see an inappropriate photo or comment on my profile.
5. Your employer using Facebook to monitor your conduct while you're at work.
6. Your employer using Facebook to monitor your conduct while you're away from work.
7. Thieves using Facebook to track, monitor, locate, and identify you as a potential victim.
8. Your employer using the information on your Facebook profile to assess your suitability for the company.
9. Law enforcement using Facebook to track illegal activities (illegal drug use, underage drinking, etc.).
10. Your university using Facebook postings, personal information, and images to identify you as a university code violator.
11. Sexual predators using Facebook to track, monitor, locate, and identify you as a potential victim.
12. Political parties using Facebook to target you through the use of advertisements and data mining.

C. QUESTIONS ABOUT PROFILE INFORMATION

The study application presented the following questions one at a time and listed the possible answers as percentages from 0% to 100% in increments of ten. The application asked the second question for each network the participant was associated with.

1. What percentage of your Facebook friends do you trust with access to your profile and shared information?
2. What percentage of the NETWORK-NAME network members do you trust with access to your profile data?

C.1 Custom Friend Lists

The study application presented the following question for each of the participant's custom friend lists.

1. Why did you create the friend list LIST-NAME and what do you use it for? Friend lists are a Facebook feature for grouping friends. If you cannot remember, enter that as your response.

C.2 Individual Facebook Friends

The study application asked the following questions in regard to nine of the participants' Facebook friends. The second question was presented as a 5-point Likert item with the anchor points uncomfortable, indifferent, and comfortable.

1. What is your relationship to FRIEND-NAME? Select the answer that fits best.
 - A member of your immediate family (parent/sibling)
 - A member of your extended family
 - A coworker
 - Someone you know from high school/college/grad school

- A friend of a friend
 - Someone you have not met in person
 - Someone you socialize with in person
 - Not sure
2. How do you feel about FRIEND-NAME viewing all the information you have uploaded to Facebook?

C.3 Individual Posts

The study application asked the following questions in regard to ten of the participants' posts. The questions were presented one at a time as Likert items on a 5-point scale with the anchor points uncomfortable, indifferent, and comfortable.

1. How would you feel if a stranger saw this?
2. How would you feel if a member of your immediate family saw this?
3. How would you feel if a member of your extended family saw this?
4. How would you feel if a coworker saw this?
5. How would you feel if someone you know from high school/college/grad school saw this?
6. How would you feel if a friend of a friend saw this?
7. How would you feel if someone you have not met in person saw this?
8. How would you feel if someone you socialize with in person saw this?