The Impact of Social Navigation on Privacy Policy Configuration

Andrew Besmer
UNC Charlotte
9201 University City Blvd
Charlotte, NC 28223
arbesmer@uncc.edu

Jason Watson
UNC Charlotte
9201 University City Blvd
Charlotte, NC 28223
jwatso8@uncc.edu

Heather Richter Lipford
UNC Charlotte
9201 University City Blvd
Charlotte, NC 28223
heather.lipford@uncc.edu

ABSTRACT

Social navigation is a promising approach to help users make better privacy and security decisions using community knowledge and expertise. Social navigation has recently been applied to several privacy and security systems such as peer-to-peer file sharing, cookie management, and firewalls. However, little empirical evaluation of social navigation cues has been performed in security or privacy systems to understand the real impact such knowledge has on user behavior and the resulting policies. In this paper, we explore the application of social navigation to access control policy configuration using an empirical between subjects study. Our results indicate that community information does impact user behavior, but only when the visual representation of the cue is sufficiently strong.

Categories and Subject Descriptors

H.5.3 [**Information Interfaces and Presentation**]: Group and Orgazational Interfaces—*Evaluation/methodology*

General Terms

Human Factors, Security, Experimentation

Keywords

Social Navigation, Social Networking, Privacy, Policy Configuration

1. INTRODUCTION

More and more details of our lives are moving into the digital world. Hundreds of millions of people already maintain a social network site profile, post photos online, share music and files, and many other online social activities. On many of these sites, users are charged with managing their personal information by controlling what gets shared with whom. They must determine appropriate and desired privacy policies for a wide variety of data and contexts. A number of policy mechanisms have been proposed to protect personal information in a variety of such settings [2, 4, 25]. Yet, many solutions rely on explicit user input without providing much assistance with policy creation and decision making. Policy interfaces can be time consuming and difficult to use [18, 22], and privacy and security decisions are often complex and highly contextual. Without adequate privacy support, users are not able to maintain control over their personal information. Privacy intrusions, both personal and publicized, may lead to reduced participation and benefits of online information sharing, hurting both users and the businesses built upon their information.

Social navigation may aid users in making better decisions by informing them of the previous decisions made by themselves or others. Social navigation is defined as the use of social information to aid a user's decision [7]. In the real world, social navigation is commonly used in everyday interactions. For example, a person might decide to visit a store based on the number of cars parked outside. We use cues like this to make an interpretation of the attractiveness of the store. More cars may indicate better prices or a wider selection, while fewer cars may indicate higher prices and more exclusivity.

DiGioia and Dourish have argued that social navigation can also be used in a security context [6]. The goal is that by relying on the collective decisions of a community of users, people will make more informed and appropriate security decisions. DiGioia and Dourish then demonstrate this approach by visualizing which folders are shared or accessed in Kazaa, a peer-to-peer file sharing application. The visual cues provide users with knowledge of conventional use and activities of others, and may prevent users from inadvertently sharing more than intended. Goecks and colleagues have also explored social navigation in the domain of cookie management for web browsers [10] and firewall policy configuration [9]. In these systems, users may lack the technical knowledge to make good privacy and security decisions. The knowledge of other users' actions can help close this gap. Thus, social navigation may help users make better decisions if they can understand the social cues, and if those social cues are actually correct [11].

Despite these examples, little empirical evaluation of social navigation has been performed in security or privacy systems to understand the real impact such knowledge has on user behavior and the resulting policies. Social navigation merely provides cues that users do not have to consider, or could overlook or ignore. However, if users are very reliant on such cues, they could instead be led into making inap-

Symposium on Usable Privacy and Security (SOUPS) 2010, July 14–16, 2010, Redmond, WA USA
.

propriate decisions due to herding or malicious users gaming the system. We need to understand the impact of such cues on security and privacy policies in order to determine when and how they can benefit users' decision making.

In this paper we seek to evaluate whether a social navigational cue impacts user decisions. We evaluated social navigation in the domain of access control policies for social applications on Social Network Sites (SNS). We have previously proposed the creation of an access control policy governing the profile information that a 3rd party application is allowed to access on sites such as Facebook [3]. We added a social navigation cue to the policy interface, which provides information about what percentage of other users shared individual pieces of information with a particular application. Using this interface, we conducted an experiment where users set their policies for a number of different applications, and where we varied the social cues provided. Our results indicate that in this domain, a social navigational cue can impact user decisions, but only when the cue is sufficiently strong.

2. BACKGROUND

The policy decisions involving online personal information often regard who that information can be shared with: whether photos are public or private, who has access to a music list, or which friends can subscribe to status updates. These decisions are often considered privacy decisions, where people determine the boundaries between what is private and what is public based on the social situation [17]. Indeed, information sharing is highly governed by the social norms of a given context [16] and privacy problems will arise when information is shared beyond the social expectations of the context. Managing the privacy of online personal information can be challenging, as users must decide a priori how to create policies that reflect the perceived future contexts for that information. Users tend to underestimate the audience of their information [1, 13, 22], or misunderstand the information flows and implications of privacy settings, resulting in information being shared out of the intended social context [14].

Researchers have examined various security mechanisms for protecting personal information, such as on social network sites [3, 8, 21]. Others have investigated interfaces for representing such policies [18, 20, 24]. But little work has examined how to help users make better policy decisions. These decisions are made and influenced by the social context of the user. Users often have similar goals and want to share information that is appropriate to a situation, while protecting sensitive information that is not relevant. In exploring the use of social navigation, we are further utilizing that social context to provide guidance for the user's policy decisions.

Social navigation has been studied and used in a variety of collaborative domains such as helping users select news stories [19], recipes [23], and research articles [15]. A more complete review can be found in [12]. In the domain of security and privacy, social navigation has been explored in a peer to peer file sharing application, cookie management, and personal firewalls [6, 10, 9] yet with minimal evaluation of the impact. For example, in a small deployment study, Goecks et al. found that users of their Acumen system did view and utilize the community information in making cookie decisions [10]. However, it is unclear how much users' decisions were really influenced by the social navigation information.

Thus, despite the potential of social navigation, we do not yet understand the real impact and benefits such cues may have on users' privacy policies. This is important not only to understand the potential benefits, but also because social navigation is not without its own unique problems and challenges. One issue with using social information is the potential for herding or informational cascades [9]. This situation arises because decisions made based on the cues contribute to community knowledge. More and more users then rely on that cue, resulting in a cycle where a cue no longer reflects actual community knowledge. The prevalence of such informational cascades, and their potential solutions, will depend upon the influence the community data has on a user's decision.

We aim to add to the understanding of social navigation by empirically testing the impact that community information has on the decisions made by users in the domain of access control for social applications. In order to expand functionality, many of the leading social network sites have created a platform to allow 3rd party developers to create applications that utilize and enhance users' profiles. These platforms provide the ability for applications to consume users' profile data such as names, birthdays, interests and more. Popular applications allow users to share books they have read and movies watched, play games, and share virtual gifts. We have previously proposed and evaluated an access control mechanism to allow fine-grained control over the information that applications can access from a user's profile [3]. We added social navigation to our access control interface as a mechanism to motivate more users to modify their policies, however we never evaluated the impact of this mechanism. In this paper, we have modified this interface in order to explore social navigation in detail. We believe this domain is attractive for experimentally testing social navigation because the user can be asked to make multiple policy decisions, over a variety of data and applications, in a short amount of time.

3. PROTOTYPE

Our prototype interface for our current study is an extension of a previously developed Facebook application for managing fine-grained access control for social applications [3]. We updated our previous prototype by modifying the social navigation cues, adding features that allowed us to record interactions and modified the design to work with Amazon Turk. An example of the updated prototype is presented in Figure 1.

Users would view the interface on the first time visiting an application. They are presented with a form to authorize the application to access profile information. The user is presented with a set of data items that the application is requesting, both required data fields and optional ones. The user can choose to not share any optional data items using the checkboxes. In addition, in order to convey to the user that the application is also able to access information about a user's friends, the information from a random friend is also added to the interface. By clicking *Continue to...*, users indicate they authorize access, and by clicking *Cancel* they do not authorize access and will not interact with the application.

The prototype interface is itself implemented as a Face-

Figure 1: Social Navigational Prototype

book application because it allows us to access real user profile information to populate the data fields, which adds realism to our study. The prototype was designed to look as real as possible so the user can make privacy decisions in a real world context. We also designed the prototype to be aesthetically similar to the current Facebook authorization screen. While the interface is created to appear as real as possible, it is in fact a simulation. Users are told the prototype is an application container, and they are asked to test the container by adding a set of chosen applications to their profile using that container. The prototype presents a set of hard-coded scenarios, where no applications are actually added and no information is shared with any application besides our own prototype. In a previous study, our participants did believe they were actually performing the authorization task [3] and adding the applications to their profiles. In this study, feedback from participants suggested that they believed these tasks were real as well. Thus, we are able to manipulate scenarios to conduct an experiment in as realistic a setting as possible.

For this study, we chose seven applications, like the one in Figure 1, that reflect the types of applications commonly found on Facebook. We chose real applications that were not popular, but were similar to popular applications. We wanted to minimize the chance that participants had already interacted with an application, which would impact their policy decisions. The interface provides a link to the real application page if users want to learn more about its functionality. We also chose the types of data that each application would request, asking for a variety of profile information that is available to applications. We decided that each application would require that a user's name, networks, and friends be shared. While not necessary for the social navigation study, these items are commonly needed by applications in order to run properly, so we required them for realism. We also continued to show the information about a friend on the interface to again reflect actual platform implementations and make the interface as realistic as possible.

Each data item occupies a row, with a checkbox to determine whether the data is shared or not. By default, the box is checked, again to reflect the default policy on today's application platforms. When the data item is restricted the row is struck through. In addition, the corresponding row for the friend's information is struck through. Figure 1 illustrates how not sharing books results in the strikethrough for both the user and the user's friend's data. Beside each check box is a short simple description of the data item, for example, Books or Music. Adjacent to the description is actual data from the user's own profile. Finally, the social navigation cue is displayed on the far right.

The value of the cue represents the percentage of people who have allowed this application to access the data item present on that row. The bar appears green when the percentage of people sharing is high and red when the percentage is low. This cue was our independent variable in our experiment. We chose the values of the cues in advance using a random number generator. We implemented two versions of this cue. In the weak cue version, only the red/green bar is shown. In the strong cue version, the interface highlights the entire row in red for a negative social cue. The red highlight disappears if a user decides not to share that data item. We decided not to add a stronger positive cue for several reasons. First, we thought the interface would look cluttered if too many red and green highlights were present. Second, the red highlight indicates that the user's decision to share is not in line with the community's decisions. This means that the equivalent positive cue would only show up if the user decided not to share a data item that was commonly shared. Thus, the cue would not have influenced the initial decision to not share information and might be confusing.

4. STUDY DESIGN

Our study was designed around answering two main questions regarding the impact of social navigation use. $Q1$: Will a positive social navigational cue result in a willingness to increase sharing? $Q2$: Will a negative social navigational cue result in a decrease in the willingness to share? In our interface we refer to a high percentage as a positive cue, as it indicates a general willingness to share in the community. A low percentage we refer to as a negative, as it indicates that few other users are willing to share.

4.1 Methodology

We designed our study to determine if the introduction of social navigation would result in differences in user policies. To accomplish this, we conducted an online between subjects user study using a combination of Amazon Mechanical Turk and the Facebook Application Platform. Amazon Turk is system which lets anyone post a small task or job (called a HIT) to be completed in exchange for money. We used Amazon Turk as a way to recruit a wide range of Facebook users to complete our study in exchange for a small monetary inducement. We offered 50 cents for 5 minutes of time.

A Turker's assignment was to "use our container to add applications as you see fit." If a turker accepted the HIT, they were directed to use their Facebook account to visit our Facebook application, shown in Figure 1 and described above, to complete the user study. Once they completed the study they were offered a code which could be entered into Mechanical Turk to receive payment.

Participants were told our application was a container for other applications. We asked participants to use our application container to review seven random applications we had picked for them. Their job was to make decisions about sharing data items with those applications. We informed them that others' decisions had been made available to them. Finally, we informed participants that at the completion of the

study we would assist them in removing the applications authorized by them if they wished.

As participants first encountered our study application, they were asked to read and electronically accept an informed consent agreement. After consent was obtained, participants were randomly assigned to one of the five treatments, described below. As the participant progressed through the study, they then could uncheck boxes and hit *Continue to* or *cancel* to go on to the next scenario. We recorded a number of metrics for later analysis including timing, which applications were authorized, the values of each of the checkboxes, and the length of the data from their profiles. After they completed the seven scenarios they were given a post study survey which gathered demographics about their age, ethnicity, and location. In addition, they were asked to self rate the impact the cue had on their decision making using a Likert scale. Finally, we asked participants to quickly explain what they liked and disliked about the interface.

We assigned participants to one of five groups. The first group was the control group which received no social navigational cues. Data from this group is used to establish a baseline to compare against the other 4 groups. Groups 2 (G2) and 3 (G3) were given what we refer to as a weak cue. A weak cue did not highlight the entire row as depicted in Figure 1, but just displayed the red/green percentage bar. Groups 4 (G4) and 5 (G5) contained a strongly negative cue which highlighted the entire row for the low percentages. The red background on the row would disappear only if the data item was restricted from being shared with the application. The positive cue for G4 and G5 remained as it did in G2 and G3.

All five groups received the same set of applications, requesting the same set of data items, and in the same order. The applications and data items were determined by us. We chose applications similar to the ones we saw being installed on many profiles but that were new enough that we did not think many participants would have previously come into contact with them.

The value of the social navigation cue was predetermined. We "seeded" the cue by determining in advance whether G2 would see a positive or negative cue for each item. When picking whether the cue would be positive or negative we did not concern ourselves with the type of data. Instead, we tried to pick different combinations of positives and negatives so participants would believe the cue represented actual community feedback and they would see a variety of both positive and negative cues [11]. Using that seed we then assigned values to the remaining groups.

Figure 2 illustrates the assignment of cues to the different groups. G2 and G4 received the same cue value and varied only by the strength of the cue. G3 and G5 received the exact opposite treatment as G2 and G4. For example, if G2 and G4 had a negative cue for some data item, say books, then G3 and G5 had a positive one. Figure 2 illustrates this assignment for the first application, a books application. Notice movies is low, or negative, for G2 and G4, but positive for G3 and G5. The opposite is true for books. This pattern was true for all but two of the applications which received negative cues or positive cues for all groups. After determining the positive or negative cues for each data item for each group, we used a random number generator to assign values of 1-20% to negative cues and 80-100% to positive cues. While in reality, values may be seen between 20%-80%, we chose extreme values to influence the participant as much as possible and simplify analysis.

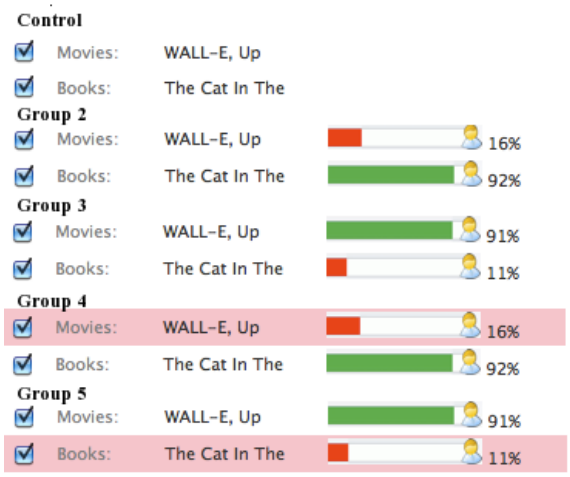

Figure 2: Example of Treatment Assignment

4.2 Analysis

We used the recorded session data to compute sharing scores for positive and negative cues for each participant. First, we calculate the authorization value *Auth* for each application i using a simple binary function based on the user clicking *Continue to* or *cancel*:

$$Auth_i = \begin{cases} 1 & i \text{ authorized} \\ 0 & i \text{ unauthorized} \end{cases}$$

In addition, we assigned values for the data items j represented by the checkboxes for each application $Box_{i,j}$:

$$Box_{i,j} = \begin{cases} 1 & i,j \text{ true} \\ 0 & i,j \text{ false} \end{cases}$$

Now we create two sets of scores, $set1, set2$ which are calculated as follows:

$$set1 = \sum_{i=0}^{n} Auth_i * \sum_{j=0}^{n} Box_{i,j} \text{ for negative seeded } j's$$

$$set2 = \sum_{i=0}^{n} Auth_i * \sum_{j=0}^{n} Box_{i,j} \text{ for positive seeded } j's$$

This represents how many items were shared for items in $set1$ and $set2$ for each user. This gives us the ability to compare the positive cues and negative cues against the control group for each data item across weak and strong cues. We use $set1$ for negative data items in G2 and G4 and positive items in G3 and G5. The control, which had no cue, establishes our baseline to test against. Similarly, $set2$ contains the positive data items for G2 and G4 and negative items for G3 and G5, all of which are comparable to the control. For example, in Figure 2, Movies would be in $set1$ and Books in $set2$.

Simply combining the scores into one sharing score per participant would not have helped us, as we are interested in whether the bar being positive increases sharing while

the bar being negative decreases sharing, and users saw a combination of both positive and negative cues. We will see later that this separation, while necessary, comes at the cost of running multiple statistical tests.

5. STUDY RESULTS

Using Amazon Turk for recruitment, we collected data from 408 Facebook users in late 2009 and early 2010. One potential problem with using Turk as a recruitment platform is the possibility that participants will only be motivated to complete the task as quickly as possible in order to obtain the offered inducement. In order to account for this, we removed outliers based on timing. We first converted the total time spent configuring the application to a Z-Score. We then trimmed the excess 5% of times on both tails leaving us with 390 participants. We felt that 5% would be adequate in removing participants who spent either too little time to have taken the study seriously or so much time that they did not represent the 'average' user.

Of the 390 remaining, 286 participants were aged 18-30, 67 aged 31-40, and 36 aged 41 and older. 35% indicated that they had completed a 4 year college degree and 39% indicated they had completed at least some college or earned a 2 year degree. The remaining participants either had extensive college degrees or no college education. Participants were asked to self identify their ethnicities to which 188 indicated they were Caucasians and 133 identified themselves as Asian. The majority of the remaining participants classified themselves as other.

5.1 Default Decisions and No Decisions

Out of the 390 participants, 15 of them chose not to add any of the applications. In addition, roughly 59% or 231 participants added every application using the default setting of share. In other words, they authorized every application and did not uncheck any boxes. This number, while alarmingly high, was not unexpected. In our previous small in lab study [3] where we did not evaluate social navigation, we noticed many participants allowing the default policy. We do not believe this represents participants who did not take the task or study seriously. Instead, we believe these participants were either unwilling or unmotivated to make privacy decisions.

Table 1: Overall types of decisions made

Group	No Apps	All Apps + Default	Custom
Control/G1	2	46	28
G2	2	46	28
G3	3	49	28
G4	0	51	27
G5	8	39	33
All Groups	15	231	144

The remaining participants made custom decisions either to allow some applications or not to share some data items. A group by group breakdown is shown in Table 1. Overall, each group has similar rates of participants accepting the defaults, not adding applications, and setting custom policies. The only notable exception is group 5 where more applications were not added and fewer participants accepted the default policy. However, approximately the same number of participants created custom policies. For the remainder of our analysis, we only focus on the custom groups as we are not interested in those who made no decisions and shared everything or those who decided not to authorize any of the scenarios in our study. Only the custom group would have been influenced by social navigation information.

5.2 Overall Effects

We previously described the computation of the two scores that will be used in our analysis, $set1$ and $set2$. They contain the paired scores for positive and negative data items in relation to the control for G1, G2, and G4 as well as G1, G3 and G5. Since we need to separate the positive and negative treatments as well as the groups and data items we must run two Kruskal-Wallis nonparametric tests for K Independent Samples. To adjust for the running of two statistical test we used a Bonferonni adjustment to set our new p value to $p = .025$. Results indicated that there were significant differences between both sets of scores for the data items that were cued negatively, $\chi^2(2, N = 83) = 8.088, p = .018$ for $set1$ of G1, G2, G4 and $\chi^2(2, N = 89) = 14.282, p = .001$ for $set2$ of G1, G3 and G5.

In order to determine which types of negative cues resulted in a difference, we chose to use the Tamhane post hoc criterion as it does not assume equal variance and our data is non-normal. We identified significant differences between the control group and the strong negative cues in G4, $p = .018$. In addition, we discovered significant differences between the control and the strongly negative cues for G5, $p = .010$, and between the weak negative cue in G3 and the strong negative cue in G5, $p = .001$. No significant differences existed between the control group and any of the weak negative cues. These results indicate that there was little, if any, difference when participants were presented with a weak negative cue. There was only a difference in the policies configured when participants were given a strong negative cue.

In addition, when presented with a positive cue, participants' decisions were not statistically different from those they would have made without social navigation $\chi^2(2, N = 83) = 4.781, p = .092$ for G1, G2, G4 and $\chi^2(2, N = 89) = 4.186, p = .0.123$ for G1, G3, and G5. Note however that we did not use a stronger positive cue.

We were also interested in seeing if the introduction of the social navigation cue would cause participants to take more time in configuring policies. We used a one-way ANOVA to test for differences among the groups in timing. We found that timing did not significantly differ across the five groups, $F(4, 143) = .454, p = .770$. Each of the groups spent about two minutes configuring the seven applications.

Table 2 illustrates the overall policies configured for those participants in each group in addition to the overall timing. The number in each cell reflects the percentage of users who shared the data with the application. In all cases, when a strongly negative cue was presented, the percentage of participants willing to share decreased in comparison to those with a weak cue or the control. We calculated the difference between the percentage of those willing to share the data item and found on average a 20% drop when a strongly negative cue was provided. Using the harmonic mean of the sample sizes, we work this out to be roughly 6 out of 29 people who changed their decision to not share that data item.

For the weak cues in G2 and G3 we see that in certain cases there seem to be major decreases in the percentage

Table 2: Percentage of participants sharing

Application	Attributes Requested	Control/G1 n=28	Weak G2 n=28	Weak G3 n=28	Strong G4 n=27	Strong G5 n=33
	Total Time	$m = 113.63$ $\sigma = 53.93$	$m = 105.43$ $\sigma = 44.28$	$m = 111.46$ $\sigma = 41.84$	$m = 118.11$ $\sigma = 34.41$	$m = 118.73$ $\sigma = 40.77$
Books	Movies	71%	82%	93%	63%	67%
	Books	71%	93%	82%	78%	58%
I Remember	Birthday	61%	50%	68%	59%	48%
	Hometown	61%	64%	68%	56%	42%
	Interests	68%	75%	68%	63%	42%
The Heist	Interests	54%	54%	68%	33%	36%
	Quotes	50%	57%	68%	33%	39%
Magic Fortune	Birthday	46%	32%	54%	30%	58%
	Work History	36%	21%	50%	26%	36%
	Sex	64%	46%	79%	33%	64%
	Interests	54%	54%	57%	33%	30%
Stickers	Religion	61%	50%	61%	48%	30%
	Political	68%	39%	61%	44%	42%
	Music	68%	75%	68%	44%	45%
	Movies	68%	64%	64%	48%	55%
Name Analysis	Hometown	57%	57%	57%	48%	58%
	Interests	61%	71%	71%	56%	61%
My Poll	Hometown	54%	50%	64%	15%	52%
	Religion	61%	50%	68%	26%	36%
	Political	64%	57%	54%	37%	24%
	Music	61%	71%	64%	41%	45%

* Negative treatments are indicated by red or gray on printed text.
* Positive treatments are indicated by green or light gray in printed text.

of those willing to share. While this is not significant at the overall group level, it may be that there are individual differences for types of data items. Unfortunately, it is not feasible to test every single attribute for every single group in relation to the control and find anything significant. The decrease is also certainly not consistent. For example, the second application, *I Remember*, had higher rates of sharing under a negative treatment than the control and lower in some cases for the positive treatment. The stronger cue was far more consistent in leading to the significance we found.

Two of the applications in Table 2, *The Heist* and *Name Analysis*, received identical cues across all treatments varying only by the strength of the negative indicator. These provide further evidence that strong negative cues decrease sharing while positive cues seem to have little effect, if any. For example, *The Heist* has similar levels of sharing for both attributes in the control, as well as G2 and G3. However, once the cue is represented more strongly, the percentage of those willing to share is reduced.

Another application, *Name Analysis*, does not have a negative indicator on any data attribute. As a result, the control was presented with no social navigation while groups 2 through 5 saw the exact same screen. Across each of the groups we see similar rates of sharing the two data items, again providing evidence that the positive cue had no effect.

5.3 Willingness to authorize applications

In certain situations the participant may feel that it is easier not to authorize an application rather than configure a custom policy and uncheck a number of boxes. Table 3 shows the number of cases in which a participant decided not to authorize an application instead of setting a custom policy. A Kruskall-Wallis test shows no significant differences in the number of applications added between groups, $\chi^2(4, N = 144) = 2.787, p = .0594$.

Table 3: Decision not to authorize app instead of custom policy

	Control n=28	G2 n=28	G3 n=28	G4 n=27	G5 n=33
Books	2	1	1	1	5
I Remember	6	4	4	3	5
The Heist	9	7	6	6	6
Magic Fortune	9	7	6	5	6
Stickers	6	5	4	5	7
Name Analysis	6	3	4	3	4
My Poll	9	5	5	9	8

Ironically, the biggest difference in mean rank was for the control group. This is also apparent in the table as the control group appears to have used not authorizing an application as a strategy more often than those with social navigation. As we did not have statistical significance this could have been due to sampling error, but it was surprising never-the-less as it went against our expectations. Thus, it appears that the difference in sharing for the strong negative cues, was due to more participants unchecking boxes.

5.4 Strategies for handling data

A close examination of Table 2 shows a lot of variation between data items, even the same ones in different scenarios. For example, we displayed *Hometown* with a strong negative cue for two applications: *I Remember* and *Name Analysis*. The respective percentage of respondents who chose to share was 42% and 15%. Additionally, users seemed less willing to share data items with certain applications, such as *Magic Fortune*. Across most applications in our study, items such as *Movies*, *Music* and *Interests* showed higher rates of sharing than data items like *Religion* and *Hometown*. Thus, the cue seems to have had different levels of impact on different data items.

These results show that the decision to share is likely a complex set of factors, each impacting the other, and is confounded by the interpretation of the application and sensitivity of the data item. Although it is possible to see some overall patterns in the sharing percentages, there is no obvious structure that can be attached to categorize or group the data items. Grouping these items would provide insight into other factors that influence our participants' decision to share. As a method to potentially detect possible structure between the data items, we performed principal axis factoring. Principal axis factoring is used to reduce multiple variables into factors which can then be used to model data. In our case, loading results provided three factors that contained data items that seemed arbitrary and unrelated and did not provide any useful patterns.

5.5 Participant Perceptions

After completing the study, participants were asked to provide demographics and answer a few short questions. We asked participants to use a 7 point Likert scale to rate the degree to which they agreed with the following statement. "The green/red bar helped me make a decision." A score of 1 indicated they strongly disagreed and a score of 7 meant strongly agree. Across Groups 2-5 we saw similar responses indicating that participants neither agreed nor disagreed with the effect of the cues on their decision making ($mode = 4$). Table 4 contains the means for each group. The control (G1) received no social navigational cue and therefore has no score.

Table 4: Self report of cue's effect on decision making

	Control n=28	G2 n=28	G3 n=28	G4 n=27	G5 n=33
m		4.30	4.11	4.20	4.0
σ		1.77	1.83	1.61	1.61
$mode$		4	4	4	4

Many participants reported that the interface allowed them to feel more in control, that they liked the transparency provided by the interface, and thought the task was easy to do. However the focus of our analysis here is to draw out issues around social navigation. Some wondered why users chose not to share certain attributes creating a need for more information, P139: "*It left me more unsure as to why people excluded things and if I should have done the same as well.*" Information about the application was available if the user clicked on the application's name but we have no measurements on how many times this occurred.

Even with social navigation, social context clearly played a role in participants' decision making. P319: "*[I was also] thinking about what others may think about particular applications - would they like it and add it too, or dislike it and think condescendingly toward me (good quality vs. low quality applications). In this way, adding applications was certainly biased based on several factors.*" Thus, decisions made by participants included a variety of information sources including social norms (from cues), their own evaluation, and the context of the applications.

Additionally, we found that some users doubted the cue, because others would not take the time to provide good feedback. For example, P360: "*The (green/red) bar is not that useful since my friends almost always add applications that spam since they do not research what apps they add... so what they do does not matter to me.*" This can have detrimental effects for social navigation [9, 11].

It was also interesting to see how participants interpreted what the text "others have shared" above the social navigation cue meant. Previously we had actually labeled the cue with "friends of yours shared" and got comments about how it was impossible for any of their friends to have used the application. When we generalized it to "others have shared" we anticipated they would rationalize that users of Facebook had shared these data attributes. Instead, many of the comments referred to what their friends had decided to share with the application, indicating users' mental models of the cue actually defaulted back to their closer circle of friends rather than a larger unknown community.

6. DISCUSSION

6.1 Behavioral Impact

Social navigation has been proposed as a mechanism to help users make informed security and privacy decisions. Our study empirically determined whether the use of social navigation modified users' decisions in one particular interface. The primary result of our study is that social navigation does have some impact on users in the domain of access control settings for social applications, but only with a strong cue. Using the weaker visual cue had no overall effect on modifying our participants' behavior. It may be that the weak cue did help a few users with a few data items of importance to them. There were users in the weak cue groups G2 and G3 who commented on the cues in the free-form survey question, so they were noticed and considered at least by some. But such a weak social navigation cue seems unlikely to have much impact, if any at all, on this policy configuration task.

The strong cue had significant impact, resulting in an average of 20% less sharing. While the community information did influence users to some degree, their decisions still seemed to be heavily based on other aspects, such as their privacy preferences for their data items and sharing them within their social contexts. One limitation is that our experimental set up may have made the social cue less believable. Our cue was manipulated to be artificially high or low, and some cues may not have made much sense. Thus, some users may have questioned the accuracy of the cue, or the motivation of the community of users the information was based upon. For social navigation to have any influence,

users must be able to understand and trust the integrity of the cue, which may require providing additional information about how the community information was calculated or why users made certain decisions. An interesting question is whether the value of the cue makes a difference, or merely the presence of the red highlight. For our experiment all negative cues were between 1 and 20%. If, for example, values under 50% resulted in the same visualization, would it have the same impact? This needs further study to understand how such a social navigation cue would function in a real application.

Not surprisingly, our result implies that just with all kinds of security warnings, the presentation of the social navigation cue matters [5]. There are many things competing for a user's attention, and in this policy decision, clearly many aspects of privacy that users are considering, such as the sensitivity of the information and the purpose of the application. Users were making rapid decisions based on the most salient information available. The strong cue was much harder to ignore, drawing the eye to those rows of data and always visible until the user chose to not share the information.

We had hoped that the social navigation cue, particularly the strong one, would act as motivation for more users to not accept an application or to customize their policies. In other words, we were expecting differences between groups in Table 2. Yet there were no differences. Across all groups, including the control, a large number of users accepted the default, and open, policies. Thus, the social navigation cue only impacted the smaller number of people who would already be modifying these policies. One method to increase the use of community information by the users inclined to simply accept the defaults is to provide an additional button to authorize the application with community default policies. Utilizing the community decisions would then require no additional effort. We would need to study this additional condition to determine how many users would decide to use such an option and under what conditions.

The presence of multiple negative cues could have also led users customizing their policies to become skeptical of more applications and refuse to authorize more as a result. Yet, instead of impacting their decision to authorize an application, the strong social cue led participants to uncheck more data items. This could be either a positive or negative effect. Users would still be inclined to interact with applications, gaining the benefits they offer. However, users may also continue to add and share at least some information with applications with less privacy than they really desire.

One limitation of our study is that we did not examine a strong positive social cue. In this domain, the current default of application platforms is to share information, which we followed in order to remain realistic. Thus, a decision to not share is the only one requiring any user interaction, and a cue to help with that decision would be more useful. However, this and other interfaces could be configured with different default decisions, where a positive cue may have been useful and had a similar impact.

6.2 User Perceptions

Our participants were neutral in their assessment of the usefulness of the cue. Few strongly disagreed that it was useful, implying that many had noticed it and at least considered the information. Yet, few also strongly agreed, again implying that the cue was only one factor among many. And there was no difference in this perception between the groups with strong and weak cues, even though the strong cue did influence users. Users may not have been aware that they were actually being influenced, either because the influence was small for each participant, or the influence of the cue was subtle, perhaps persuading users to weigh more heavily the reasons not to share over the reasons to share.

The decisions of other users could have been interpreted either subjectively or objectively by our participants [9]. Subjective data is based on users' personal preferences, and when users see social navigation cues they understand that those cues are based on other people's subjective preferences, which they may or may not agree with. In certain security situations, the decisions are more objective. Most users agree on the criteria, generally wanting to do the common and safe actions, such as to not accept cookies from dangerous sites. In this situation, users must judge the expertise of the community to make such determinations in deciding whether to use the social navigation guidance. The domain of our experiment seems to have a mixture of both. In general, users are willing to share information that an application needs to work, but not as willing to share information that is unnecessary. So, users find it acceptable to share their birthday with a horoscope application, but not their hometown [3]. So objectively many users may have the same desire to only share context appropriate information with trustworthy applications. But, the willingness to share personal information is also subjective, and users may have either very sensitive or completely false information entered on their profile, also influencing their privacy decisions. In the comments we received, users seemed to interpret the community information as more subjective in nature. Thus, in interpreting the meaning of the cue, users are likely going to be considering whether the privacy preferences of the community match their own.

6.3 Open Questions

While our study did investigate the concrete impacts of social navigation, there are several interesting questions raised that we could not examine in this experiment. The goal of social navigation is to help users make better decisions. However, our study cannot make any determination as to whether the policies that were modified due to the social navigation cues were really any "better." Our negative and positive cues were balanced between groups, and were not designed to reflect a "good" policy. In fact, determining what a good policy is in this domain is difficult. While from a security perspective, any reduction in personal data sharing with social applications is a good thing because of the potential misuse of information by 3rd party application developers, this does not necessarily reflect the desires of the users. Other stakeholders, such as Facebook and application developers may prefer users share more information to increase interaction. An acceptable measure may be that the policy reflects the actual privacy preferences of each individual, reducing unintentional disclosures while allowing beneficial sharing and interaction. In previous research, users indicated that they desired to share information that was needed and appropriate for an application to work [3]. However, we did not query users as to their expectations of the needs of each application to be able to make this determination for this study.

Another aspect of social navigation that needs additional

investigation is the impact of the community on the trust and use of the social cue. Because our application was clearly an experimental system, we chose to vaguely label the cue as what "others" have allowed. Interestingly, there were participants who indicated in their comments that they interpreted the cue, at least partially, as what their friends shared. We would expect that an individual would trust their friends' information more than unknown "others", but that there would be fewer data points to aggregate for a social navigation cue for any one application. Yet, one participant indicated that he knew his friends just added applications without much thought, so he actually did not value the cue. So we need to examine these perceptions of different communities to further understand the impact that social navigation may play in different security and privacy tasks.

7. CONCLUSION

Users are being confronted with more and more policy decisions governing the sharing of their personal information online. While social navigation has been advocated to help users make security and privacy decisions, and even deployed in prototype systems, there is still much to understand about the true impact that community information would have on the policy decisions of users. Our paper begins that investigation with an experiment that determines the impact that a social navigation cue had on application access control policies. Our results demonstrate that such cues can impact user decisions, but only if the cue is sufficiently strong. These results have implications for the design and use of social navigation in privacy policy interfaces:

- Social navigation can be used to impact user behavior and their resulting privacy policies, although that impact may be small.
- The presentation of a social navigation cue matters, and needs to be sufficiently strong to draw user attention if it is to have any impact at all.
- In this and similar domains, social navigation is not useful for motivating more users to consider their privacy and modify their policies. The cue may only have an impact on those who are already making policy decisions.
- Developers need to carefully consider how to construct such a social navigation cue. If all users' sharing behaviors were simply aggregated together in our interface, the cue would always remain very high due to the large numbers of users accepting defaults. Instead, similar to the notion of "mavens" used in the Acumen cookie management system [10], only users who regularly make policy decisions should be included in the community information.
- Developers need to consider how users will interpret the behaviors of the chosen community, and possibly provide some way to learn about those behaviors if they would be unknown.

Our experiment also revealed a number of remaining questions, even for this particular policy interface. In order to design and deploy a successful social navigation cue, we need to determine the impact of the community choice and the value of the social cue to determine how to construct and display such information. How can we determine what is a "good" policy and whether the community knowledge is benefitting users? We would also need to understand the real world impact. Would users behave similarly when accessing applications of their choice instead of those chosen by an experimenter? Just like any warning, users could become habituated to the social cue, reducing the influence over time. And would users who only accept the default settings be willing to base those defaults on the community's decisions? And finally, how do users react to social navigation in other domains? We continue to examine these questions in order to help designers understand how to provide useful guidance for users in security and privacy systems.

8. REFERENCES

[1] A. Acquisti and R. Gross. Imagined communities awareness, information sharing, and privacy on the facebook. In *Privacy Enhancing Technology*, Cambridge, United Kingdom, June 2006.

[2] E. Bertino, F. Paci, R. Ferrini, and N. Shang. Privacy-preserving digital identity management for cloud computing. In *IEEE Data Engineering*, pages 21–27, 2009.

[3] A. Besmer, H. R. Lipford, M. Shehab, and G. Cheek. Social applications: exploring a more secure framework. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–10, Mountain View, California, 2009. ACM.

[4] B. Carminati, E. Ferrari, and A. Perego. Enforcing access control in web-based social networks. *ACM Transactions on Information and System Security*.

[5] L. F. Cranor. What do they "indicate?": evaluating security and privacy indicators. *interactions*, 13(3):45–47, 2006.

[6] P. DiGioia and P. Dourish. Social navigation as a model for usable security. In *Proceedings of the 2005 symposium on Usable privacy and security - SOUPS '05*, pages 101–108, Pittsburgh, Pennsylvania, 2005.

[7] P. Dourish and M. Chalmers. Running out of space: models of information navigation. In *HCI '94*, Glasgow, UK, Aug. 1994.

[8] A. Felt and D. Evans. Privacy protection for social networking APIs. In *Proceedings of Web 2.0 Security and Privacy 2008*, 2008.

[9] J. Goecks, W. K. Edwards, and E. D. Mynatt. Challenges in supporting end-user privacy and security management with social navigation. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, Mountain View, California, 2009. ACM.

[10] J. Goecks and E. Mynatt. Supporting privacy management via community experience and expertise. In *Communities and Technologies 2005*, pages 397–417. 2005.

[11] A. Herzog and N. Shahmehri. User help techniques for usable security. In *Proceedings of the 2007 symposium on Computer human interaction for the management of information technology*, page 11, Cambridge, Massachusetts, 2007. ACM.

[12] K. Hook, D. Benyon, A. J. Munro, D. Diaper, and C. Sanger, editors. *Designing information spaces: the*

social navigation approach. Springer-Verlag, 2003.

[13] C. Lampe, N. Ellison, and C. Steinfield. A face(book) in the crowd: social searching vs. social browsing. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, pages 167–170, Banff, Alberta, Canada, 2006. ACM.

[14] H. R. Lipford, G. Hull, C. Latulipe, A. Besmer, and J. Watson. Visible flows: Contextual integrity and the design of privacy mechanisms on social network sites. In *Computational Science and Engineering, IEEE International Conference on*, volume 4, pages 985–989, Los Alamitos, CA, USA, 2009. IEEE Computer Society.

[15] S. M. McNee, I. Albert, D. Cosley, P. Gopalkrishnan, S. K. Lam, A. M. Rashid, J. A. Konstan, and J. Riedl. On the recommending of citations for research papers. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 116–125, New Orleans, Louisiana, USA, 2002. ACM.

[16] H. Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1), 2004.

[17] L. Palen and P. Dourish. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 129–136, Ft. Lauderdale, Florida, USA, 2003. ACM.

[18] R. W. Reeder, L. Bauer, L. F. Cranor, M. K. Reiter, K. Bacon, K. How, and H. Strong. Expandable grids for visualizing and authoring computer security policies. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1473–1482, Florence, Italy, 2008. ACM.

[19] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl. GroupLens: an open architecture for collaborative filtering of netnews. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 175–186, Chapel Hill, North Carolina, United States, 1994. ACM.

[20] H. Richter, A. Besmer, and J. Watson. Understanding privacy settings in facebook with an audience view. In *UPSEC '08*, San Francisco, CA USA, Apr. 2008. USENIX.

[21] M. Shehab, A. C. Squicciarini, and G. Ahn. Beyond User-to-User access control for online social networks. In *Proceedings of the 10th International Conference on Information and Communications Security*, pages 174–189, Birmingham, UK, 2008. Springer-Verlag.

[22] K. Strater and H. R. Lipford. Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group Annual Conference on HCI 2008: People and Computers XXII: Culture, Creativity, Interaction - Volume 1*, pages 111–119, Liverpool, United Kingdom, 2008. British Computer Society.

[23] M. Svensson, K. Hook, J. Laaksolahti, and A. Waern. Social navigation of food recipes. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 341–348, Seattle, Washington, United States, 2001. ACM.

[24] J. Watson, M. Whitney, and H. R. Lipford. Configuring audience-oriented privacy policies. In *Proceedings of the 2nd ACM workshop on Assurable and usable security configuration*, pages 71–78, Chicago, Illinois, USA, 2009. ACM.

[25] C. M. A. Yeung, L. Kagal, N. Gibbins, and N. Shadbolt. Providing access control to online photo albums based on tags and linked data. In *Proceedings of the AAAI Spring Symposium on Social Semantic Web: Where Web 2.0 Meets Web 3.0*, Stanford, CA, Mar. 2009.