Privacy Manipulation and Acclimation in a Location Sharing Application

**Shomir Wilson, Justin Cranshaw, Norman Sadeh, Alessandro Acquisti, Lorrie Faith Cranor,
Jay Springfield, Sae Young Jeong, and Arun Balasubramanian**

Carnegie Mellon University

{shomir, jcransh, sadeh, lorrie, jspringf, saeyoung}@cs.cmu.edu, acquisti@andrew.cmu.edu, arunb@cmu.edu

ABSTRACT

Location sharing is a popular feature of online social networks, but challenges remain in the effective presentation of privacy choices to users, whose location sharing preferences are complex and diverse. One proposed approach for capturing these nuances builds on the observation that key attributes of users' location sharing preferences can be represented by a small number of privacy profiles, which can provide a basis for configuring individual preferences. However, the impact of this approach on how users view their privacy is relatively unknown. We present a study evaluating the impact of this approach on users' location sharing preferences and their satisfaction with the decisions made by their resulting settings. The results suggest that this approach can influence users to share significantly more without a substantial difference in comfort. This further suggests that the provision of profiles for privacy settings must be carefully considered, as they can substantially alter sharing behavior.

Author Keywords

Location sharing; user study; mobile computing; privacy

ACM Classification Keywords

H.5.m Information Systems: Information Interfaces and Presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Online social networks such as Facebook and Google+ expose a plethora of settings to users, enabling them in principle to configure these environments to their likings. In particular, the number of privacy settings that these platforms offer has steadily grown over the past few years. This trend partly reflects the increasing functionality provided by these platforms as well as recognition that people's privacy settings are complex and diverse. Location sharing privacy on social networks exemplifies the extent of

this complexity [6]: research on this topic has shown that a user's willingness to share their location with others varies based on who is requesting their location [1], where the user is at the time of the request [5,21], the day of the week [5], the time of day [5], as well as other factors such as geographical proximity between the two parties [22]. These nuances, combined with the diversity of privacy preferences expressed by different users, create a challenging tension between privacy and user burden. The significance of this challenge—that is, the need for users of social networks to control the privacy of their location, in spite of the complexities of their preferences—motivates our research.

Expecting users to spend the time to fully configure their privacy settings has been shown to be unrealistic, and even willing users often find it difficult to accurately express their preferences [28]. A promising alternative involves identifying sets of aggregated settings, or privacy *profiles*, which collectively capture the privacy preferences of diverse subpopulations of users, with users being given the option to select one or more profiles that best capture their privacy preferences. For example, this approach has been featured commercially in Internet Explorer to support the management of compact P3P privacy policies. Recently a similar approach, using machine learning techniques to identify clusters of like-minded users, has been proposed to generate location sharing privacy profiles [25,27]. However, the prior work did not explore the impact of such profiles on users' sharing behaviors and preferences. Also unanswered was how this impact unfolded over time, and how it interacted with trends reported in prior studies, such as the selective relaxation of privacy preferences [28,30].

This paper reports on a three-week study of location sharing privacy preferences to explore the effects of offering privacy profiles to users. Participants used an actual location sharing app named *Locaccino* on their smartphones for the duration of the study. They were divided into two conditions, differing only by the procedures that they used to specify their initial location sharing preferences. Participants in the treatment group used a wizard that required them to choose among sets of location sharing profiles derived from prior work [25,27], and their choices were translated into rules for sharing their location. Participants in the control group used a wizard that required

them to create rules directly, without exposure to the profiles. During the study, participants provided feedback on hypothetical requests for their actual locations throughout the day from people they had identified as their friends, indicating their satisfaction with how their settings would have handled those requests. This combination of the deployment of a functional location sharing app and feedback based on actual locations and specific friends allowed us to gauge participants' satisfaction more realistically than arrangements in prior studies (e.g. [5,21]).

We hypothesized that, although the location sharing application provided users with ample opportunities to refine their initial preferences, the impact of the initial wizards would remain visible after several weeks of use. The results of our study seem to confirm this hypothesis, suggesting that the profiles can have a substantial impact on users' sharing behaviors and on their privacy preferences. A second, related hypothesis was that, while people would tweak their privacy settings over time, those changes would remain limited. This second hypothesis proved generally correct in that the absolute number of edits made by users remained somewhat limited. However, we did observe that users in both conditions seemed to gradually and selectively relax some of their privacy settings over the three weeks of the study. More interestingly, we observed that users who were exposed to the privacy profiles seemed more inclined to share than those who were not, and that their edits resulted in a larger increase in sharing than those made by their counterparts. This particular effect was a surprise to us: while we did expect the profile-based wizard to result in more sharing directly after initialization, the results suggest that the difference does not diminish over time. The profiles made users feel more inclined to share their location, leading them to substantially relax their privacy preferences over the course of the study.

We view the contributions of this research as falling into two areas. First, *the complexity and diversity of people's privacy preferences creates a major tension between privacy and usability*. Privacy profiles have been proposed as a way of alleviating this tension, and we report on a first study looking at both the short and medium-term impact of such profiles on user behavior and privacy preferences. Second, *our study suggests that privacy profiles can have a substantial impact on a user's behavior as well as their privacy preferences*. In this particular study, we demonstrate a set of profiles that led users to share more and to be satisfied with their sharing. It is likely that profiles could also be designed to have the opposite effect. As such, this research shows how critical it is for people who design these systems to carefully think about the ramifications of their design decisions. These decisions impact more than just the usability of interfaces; they lead to substantial differences in behaviors and attitudes, which in turn can impact both adoption and user privacy.

BACKGROUND

Related Work

Although location sharing has been studied by the research community for well over a decade, over the last five years it has seen considerable commercial attention thanks to the proliferation of smartphones and other locative technologies. Individual location sharing applications can vary greatly in their features and design, but they can be broadly categorized based on whether they track the locations of users continuously (e.g. Latitude [15], Find My Friends [12], and Locaccino [24]) or whether the location sharing happens discretely such as via check-ins (e.g. Foursquare [13] and Facebook [10]).

Check-in based systems have gained traction in recent years in part because they simplify the complex privacy concerns inherent with continuous sharing. However, there are a number of advantages to continuous location sharing over the check-in model. Check-in based systems require constant engagement from the sharing users, which is not always practical (e.g., while driving). This can lead users to grow frustrated or fatigued with the constant need to check-in, or they can simply forget to do so. This may explain some of the difficulty in maintaining user activity in check-in applications. As of November 2012, Foursquare reported having 25 million users, but only 8 million of them were monthly active users [19]. Also, there are several use cases that continuous location sharing supports that check-in based sharing does not, including coordinating with friends, family and colleagues, and automatic notifications of nearby points of interest. However, researchers have shown that privacy remains a significant barrier to the large-scale adoption of such systems [3,9,26,28].

A number of efforts have explored the complex privacy-related issues that arise in location sharing and ubiquitous computing [2,8]. Barkhuus and Dey studied users' concerns for location privacy in the use of location-based services and location-tracking applications [4]. While their participants did indicate privacy concerns for these applications, their concerns could be mitigated with proper controls. Consolvo et al. conducted a three-phased study of 16 participants, examining their location disclosure preferences in sharing with social relations [9]. Their investigation used experience sampling to explore participants responses to hypothetical discloser requests, both in lab and *in situ*. They found participants' responses were most influenced by who was requesting their location, why the requester wanted the participant's location, and what level of detail the requester needed. Sadeh et al. showed that sharing is often influenced by temporal and location factors as well as the granularity at which one's location would be disclosed [5,21,28]. Tang et al. found that the privacy concerns differ depending on whether the sharing was purpose-driven (i.e., utilitarian) or whether it was socially driven [29]. Users of socially driven applications in their study preferred sharing semantic place names with their friends, thus blurring their exact location

to mitigate privacy concerns, and increasing the utility of location-sharing as a means for self-presentation.

Researchers have also looked at what motivates users to share their location. Barkhuus et al. studied an always-on location-based status updating system called *Connecto*, in order to better understand the landscape of social interactions that such systems foster [3]. Lindqvist et al. surveyed users of Foursquare to better understand what drives users to check in [23]. They found various factors influence participation, including gaming aspects of the system, keeping up with friends and other social relations, the discovery of new places, and the journaling of daily activities. Despite the fact that privacy is less of a concern in check-in systems, Lindqvist et al. found that users were concerned about managing their privacy on Foursquare, and privacy concerns affect how they use the system.

Techniques have been proposed to obfuscate location data once it has been collected [20], including k-anonymity [16] and semantic obfuscation [14]. Brush, Krumm, and Scott logged the GPS traces of 32 participants (from 12 households) in a two-month long study investigating participants' levels of comfort with various location obfuscation techniques [7]. Their study showed that obfuscation does affect a person's willingness to share their personal location traces with outside entities, such as government organizations or academic researchers, but they also found a lack of understanding among participants of the broader privacy risks inherent in the data, indicating that obfuscation alone will not solve the problem.

In this work, we attempt to mitigate privacy concerns on online social networks by improving users' abilities to control their sharing. Our approach operates in the context of a continuous location sharing system that allows users to express their sharing preferences via access control rules. Prior research has shown that users sometimes have difficulty in specifying and maintaining their preferences in such access control systems [11,17,28], and that user-oriented machine learning can ease the process [18]. Accordingly, we have developed a small set of default profiles based on an analysis of users' expressed privacy preferences (explained in the subsection below), and our study examines what effects these profiles have on user sharing, comfort, and behavior.

Prior Work

Our work was motivated by and follows that of Mugan et al. [25], which is itself based in part on prior work by Ravichandran et al. [27], who took a data-driven approach to studying default profiles for privacy management in location sharing. These studies analyzed data from a three-week long location-tracking field study [5] that collected the expressed preferences of 27 participants for sharing their location with various social groups at various locations and times. By converting the raw geographic and temporal data into canonical concepts for time, day, and location (such as "weekday" or "at work"), this work showed how k-means clustering could discern a small set of policies, or *profiles*, that encoded common sharing preferences observed across many of the study participants. Their work considered four groups of hypothetical recipients of location information ("Friends and Family", "Facebook", "University", and "Advertisers") and evaluated how well different combinations of default profiles for sharing with these groups captured users' privacy preferences.

This earlier work showed that, as the number of default profiles increased, accuracy (i.e., the correspondence between each user's designated profile and their expressed preferences) increased as well, though it reached a plateau for relatively small numbers of such profiles. For instance, when looking at privacy preferences for sharing one's location with other members of the university community or with advertisers, having three default profiles was substantially better than having two, but adding a fourth profile contributed very little value. Because the analysis was post-hoc, the utility and the usability of the profiles were not tested *in situ*, and the impact of these profiles on users' preferences was not investigated.

Here, we extend Mugan et al.'s work by evaluating the effectiveness of a slightly modified set of privacy profiles in a real-world deployment of a location sharing application. We evaluate both how useful these profiles are at initially capturing users' preferences and how these initial policies evolve as users refine them over time. Although our analysis of users' preferences relies on simulated requests, it is an improvement over the prior work in two crucial ways. First, participants in the present study installed and used a functional location sharing app on their phones. Second, participants in the present study received simulated requests from specific, named individuals whom they had identified as belonging to social groups in relation to them. Participants were asked to evaluate the hypothetical results of these requests (i.e., location shared or withheld) based on their sharing settings at the time of the request. Both of these improvements situated participants more realistically with regard to the social aspects of location sharing that we wish to study.

METHODS

This section describes our three-week experiment to determine the effects of the privacy profiles from the previous work. In brief, we recruited participants to use a mobile location sharing application on their smartphones, and we monitored their privacy settings, their satisfaction with those settings, and their willingness to share their location with their social connections. We placed users in two conditions: a treatment condition which used a profile-inspired wizard to simplify the initialization of settings, and a control condition which initialized those settings directly.

Software

Participants were required to use *Locaccino*, a mobile location sharing application we have developed as a platform for studies. Locaccino allows users to selectively share their locations with their Facebook friends and to

Figure 1. The map screen from the Locaccino app, which subjects were required to run on their Android phones.

make requests to see the locations of fellow Locaccino users[2]. The service consists of a mobile app (available for Android, iOS, and Symbian) and a website with extended functionality. The app determines a user's location via GPS or nearby Wi-Fi networks, and it sends that information to the Locaccino server. A user may make a request to view another user's location using the mobile app or the website. Figure 1 shows a screen shot from the mobile app.

In addition to serving as a platform for studies, Locaccino has been publicly available for the past three years [24], allowing us to steadily refine its usability and performance in response to users' feedback. To encourage adoption, Locaccino also provides a service to view the current locations of Carnegie Mellon University campus shuttles, and these shuttles have received up to 1,000 location requests per day. Together, Locaccino's mobile apps have been downloaded by approximately 35,000 users.

Locaccino users are able to grant their friends access to their location by creating *sharing rules*. Initially, a new user's location is shared with no one, and if a user creates no rules his location remains completely private. Each rule grants access to a set of one or more friends under specific circumstances. Those circumstances, as specified by the

---

[2] Throughout this paper, a *requestor* is a Locaccino user requesting another's location, and a *requestee* is a Locaccino user whose location has been requested by another user.

user, can be either or both of the requestee's location (i.e., whether the requestee is within a certain geographic area) and a time range (by time of day or day of week). If all specified conditions of a rule are satisfied (i.e., a requestor's identity, the requestee's location, and the current time all satisfy the requestee's rules), then the requestor is allowed to see the requestee's location. Otherwise, the request is denied. Users also have the option of toggling an "Invisible Mode" for exceptional circumstances: when it is activated, the user's rules are overridden and their location is disclosed to no one. These options allow users to build simple or articulate patterns of sharing as appropriate.

Additionally, Locaccino users have access to tools that illustrate the effects of their sharing preferences. When logged in on the website, a "Who can see me?" display is available to the user. Intuitively, it shows two lists of friends: those whose requests for the user's location would be granted, given the user's rules, the present time, and their current location, and those whose requests would be denied. A "Who saw me?" display is also available, which shows the user a list of all previous requests for their location and the outcome of each request. Prior work by Tsai, et al. showed that tools like these increase users' comfort with sharing [30].

Protocol

Study participants were sought from the university community using fliers and email lists. To facilitate use of Locaccino, each participant was required to have a Facebook account and to own an Android phone with a data plan. The set of 37 subjects that enrolled in the study consisted of 26 undergraduates, ten graduate students, and one university staff member. Most subjects were frequent users of online social networks, with 84% reporting in an entrance survey that they used Facebook at least daily. However, most subjects had not used location sharing services before or were infrequent users: only 19% used such a service at least once a week. No participants were or had been regular users of Locaccino, though six participants (two in the treatment condition and four in the control condition, described below) had tried earlier, substantially different versions of the service no more recently than two years prior to the study.

Subjects were required to attend an introductory session during which they viewed a presentation on the participation requirements, answered a study entrance survey, and completed several setup tasks. They initialized their location sharing settings using one of the two wizards, to which they were randomly assigned; these comprised the study conditions, and they are described in detail in the following section. Subjects then viewed a list of their Facebook friends and were instructed to place a small number of them (2-4) in each of three categories: close friends and family, friends affiliated with the university, and friends unaffiliated with the university. They made these selections using an interface (Figure 2) on the Locaccino website, and their selections were stored for use
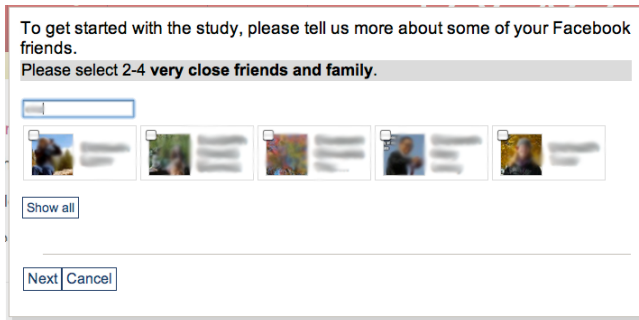
Figure 2. One of the three friend categorization screens.

later in the study. To encourage subjects to use Locaccino and to create a natural sharing experience, they were then asked to send invitations (facilitated by Facebook) to use Locaccino to all of the people that they had categorized. Finally, they installed the Locaccino app on their Android smartphones and verified that it functioned properly.

Subjects ran Locaccino on their phones for three weeks, and every evening they completed an auditing task to gauge their satisfaction with the results of their privacy choices. To do this, they visited a page on the Locaccino website where they were presented with a list of hypothetical "what-if" requests for their location at various times throughout the day. These requests were simulated by the system by taking a sample of real location observations of the user and pairing them with hypothetical requestors, sampled randomly from the user's friends so that each of the three labeling groups in the introductory session were equally represented on average. Figure 3 shows the interface for the auditing task. For each request, the user was shown the name of the requestor, the time of the request, the user's location at the time, and the outcome (e.g., location request allowed or denied) according to the user's sharing rules. The location shown for each what-if request was the subject's actual location at the given time.

At most 20 what-if requests were generated each time the user navigated to the auditing interface. The location observations were sampled so that the what-if request were no closer than 15 minutes apart, and they were uniformly sampled from the available location observations of the user since the last generated what-if request stored in the system. This allowed us to get an even sample of feedback points about users' locations while simulating the notion that location requests are continually being made. Actual requests for the user's location from their friends were also shown, though these were much lower in quantity and thus (as anticipated) insufficient for analysis.

For each auditing request, a participant was instructed to assign one of three labels to the outcome: *satisfied* (i.e., not bothered, as explained during the introductory session), *unsatisfied*, or *not there* (indicating that Locaccino had incorrectly identified their location at the given time). Subjects were presented with approximately 20 audits of hypothetical requests per day, with some variation when they visited the site more than once or missed a day. To ensure that they followed the nightly regimen, subjects' response rates were monitored, and the experimenters promptly contacted delinquent participants.

Subjects' interactions with Locaccino were recorded; in particular, their friend labeling, sharing rules, and auditing responses were logged for analysis. Subjects also completed an exit survey following the three-week period, and a subset participated in semi-structured interviews to discuss their sharing rules and their reactions to using Locaccino.
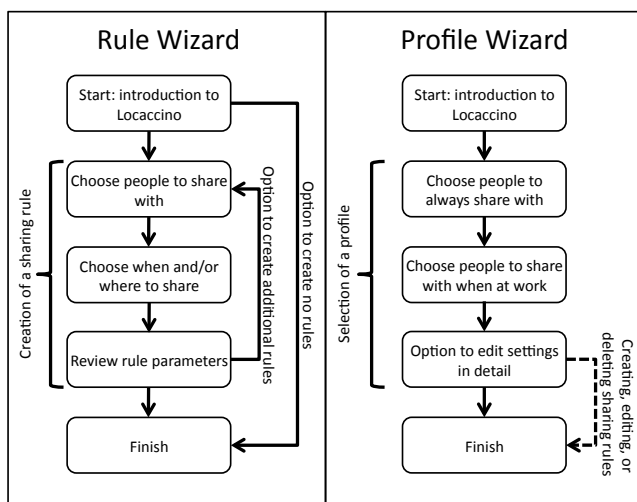
Conditions

Subjects were randomly placed into two conditions, with *ex ante* allocation to balance the gender ratios. Those in the treatment condition created their initial set of location sharing rules using a wizard that captured the functionality of the privacy profiles described in the previous section. Since the profiles for each sharing group shared common features, they were restructured as a sequence of

Figure 3. The auditing interface, with instructions and audit requests.

Figure 4. A comparison of the structure of the Profile Wizard and the Rule Wizard.

discriminative questions. Subjects in the control condition initialized their rules using a rule-centric wizard. There were no other differences between the conditions, and all subjects were encouraged to change their rules if they found them unsatisfactory during the study. After using their assigned wizards once, participants in both conditions had access to identical interfaces to alter their sharing rules.

Figure 4 outlines the structure and content of the two wizards. The rule-centric wizard guides the user through the creation of each sharing rule in steps: first specifying whom a rule should concern, then adding time restrictions as desired, and finally adding location restrictions as desired. Those three steps provide Locaccino with sufficient information to create a single sharing rule for the user, and the user is able to create arbitrarily many rules, or none at all. This wizard exposes the user to the full functionality of Locaccino's privacy settings with a nominal level of guidance to make appropriate sharing decisions.

The profile-centric wizard guides the user through the selection of the default profiles identified in our earlier work. The wizard presents two tasks to the user, in sequence: first, to identify friends with whom the user would like to share their location all the time, and second, to identify friends with whom the user would like to share their location only when at work[4]. The choice of whom to share with is made explicit in these questions, while location and time restrictions are reduced and eliminated (respectively) in line with the prior findings. Through these simplifications, the wizard is able to create rules for the user that correspond with their choices; these rules may be as few as zero and as many as two. Some rule nuances available in the rule-centric wizard are not shown in the

---

[4] The "at work" location was found to be important in Mugan et al.'s prior work on location sharing [25].

Statistic	Rule condition	Profile condition
Total	5,230	5,412
Median	260	313
Mean	291	338

Table 1. Aggregate statistics on auditing responses.

profile-centric wizard, and a link to Locaccino's rule editing interface is provided for users who change their minds about their choices or wish to see more options.

RESULTS

Of the 37 subjects who enrolled, 33 completed the study requirements. These consisted of 17 subjects in the rule-centric wizard condition and 16 subjects in the profile-centric wizard condition (henceforth referred to as the *rule condition* and the *profile condition*, respectively). The four subjects who did not complete the requirements consisted of one in each condition who dropped out partway through the study, one in the profile condition who produced aberrant auditing responses, and one in the rule condition who completed all requirements except for the exit survey (whose available data is included in analysis below).

Auditing Responses

Table 1 shows aggregate totals for auditing responses from subjects over the entire study period, as well as per-condition means and medians. For brevity, we use the following terms in the remainder of the paper:

- An auditing *request* refers to a hypothetical request for the subject's location;

- An auditing *outcome* refers to how the subject's settings handled a request (i.e., either *allowed* their location to be shared or *denied*); and

- An auditing *response* refers to the subject's answer (*satisfied* with the outcome, *unsatisfied*, or *not there*).

Subjects' auditing responses illustrated their satisfaction or dissatisfaction with how their rules would have handled their friends' hypothetical requests for their location. Although participants' auditing behaviors sometimes varied from day to day, we expected their responses to contain trends in sharing and satisfaction over the three weeks of the study. Figure 5 shows the mean percentages of the four key varieties of outcome-response pairings for the two conditions for each of the three weeks. For clarity, audits to which subjects responded *not there* are not included.

To compare subjects' satisfaction in each condition, mean satisfaction rates were calculated per participant per week. This was a simple computation using each participant's audits for a given week:

$$\text{satisfaction rate} = \frac{\text{satisfied audits}}{\text{satisfied audits} + \text{unsatisfied audits}}$$

As the figures show, both conditions appear to have an upward trend in satisfaction over the course of the study,

Figure 5. Weekly mean percentages of outcome-response pairs from subjects in the rule condition (above) and the profile condition (below).

although for the rule condition it is very slight and for the profile condition it is most marked between the first and second weeks. The difference between the Week 1 and Week 3 distributions of subjects' mean satisfaction rates is statistically significant[5] for the profile condition (means of 0.70 and 0.86, respectively; using a matched one-tailed t-test, p=0.02) but not for the rule condition (p=0.23). Comparing the two conditions by week (using unmatched t-tests), the subjects in the rule condition show a mild indication of being more satisfied in Week 1, but the difference was not significant (p=0.09). No significant differences were found between the conditions in Weeks 2 and 3 (p=0.40 and p=0.45, respectively).

Independently of subjects' responses to requests, the audit outcomes also served as a mechanism for measuring each subject's willingness to disclose their location to their friends. This is because each participant's sharing rules were responsible for the outcomes to their location requests, which were generated with the identities of friends they specified during the introductory session. To compare subjects' sharing in each condition, mean sharing rates were calculated from the requests, with one statistic per subject per week. This calculation was similar to that for the satisfaction rate:

---

[5] We assume a threshold of p=0.05 for significance tests throughout the paper.

$$\text{sharing rate} = \frac{\text{request allowed audits}}{\text{request allowed audits} + \text{request denied audits}}$$

Using a matched one-tailed t-test, the increase in the sharing rate from Week 1 to Week 3 of subjects in the rule condition was found to be significant (p=0.0005). For the profile condition, the increase from Week 1 to Week 3 fell short of the significance threshold (p=0.08), as did the increase from Week 1 to Week 2 (p=0.11), which (as shown in Figure 5) saw the greatest mean sharing rate. Comparing the two conditions by week (using unmatched one-tailed t-tests), the profile condition shared significantly more during Week 2 (p=0.01), and there were mild indications of the same for Week 1 (p=0.13) and Week 3 (p=0.093).

We observed a "gap" between participants' comfort with sharing and the rate that their settings actually shared. This gap was indicated by the *request denied, unsatisfied* audits that subjects produced in both conditions.[6] Using these audits and those of the *request allowed, satisfied* variety, we calculated a "comfort rate" that was indicative of a participant's willingness to share, independently of the effects of their settings:

$$\text{comfort rate} = \frac{\text{denied and unsatisfied audits} + \text{allowed and satisfied audits}}{\text{all (what - if) audits}}$$

Over the entire study, the mean of rule condition participants' comfort rates was 0.55, and for the profile condition it was 0.67. Although the difference was not statistically significant (p=0.08, using an unpaired one-tailed t-test), it was one of several indications that the profile wizard made participants feel more inclined to share their location than their counterparts in the rule condition.

Finally, by calculating participants' satisfaction and sharing rates for the entire length of the study, it was possible to make overall comparisons between the two conditions. Including data from all weeks, participants in the profile condition shared significantly more than those in the rule condition (means of 0.51 and 0.36, respectively; p=0.04). However, the overall difference in satisfaction between the profile and rule conditions was not significant (means of 0.81 and 0.82, respectively; p=0.36).

Privacy Preferences
Most subjects made changes to their privacy settings during the study; 61% of those in the Rule Condition made at least one change to their rules, and 76% did so in the Profile Condition. The mean number of changes a subject made to their rules (including rule creation, modification, or deletion) was 1.9 for the rule condition and 1.4 for the profile condition, and for both conditions the median was one. Using an unmatched one-tailed T-test, the difference between the conditions' distributions of total changes was not statistically significant (p=0.27).

---

[6] The persistence of audits in this category across all three weeks was perhaps a symptom of acclimation to privacy settings, an explanation we revisit in the Discussion.

Figure 6. Weekly mean changes to subjects' sets of sharing rules, including rule creation, modification, and deletion; for the rule condition (above) and the profile condition (below).

Figure 7. Weekly mean quantities of rules containing each variety of articulation, for the rule condition (above) and the profile condition (below).

Figure 6 shows the volume of rule changes per week for each condition, excluding activity during the introductory session. As shown, rule condition subjects did not delete any rules, while profile condition subjects deleted a few (a total of two) during only the first week. At a high level, the two conditions appear to be roughly equal when comparing specific weeks and changes to rule sets, but examining the qualities of settings changes showed differences. In the rule condition, the most common change to settings was to add a new rule; all such new rules shared with manually chosen sets of five or fewer friends, with varying time and location restrictions. In the profile condition, the most common change to settings was to add friends to an existing rule. Nearly all such additions were to rules that shared with friends all the time and at all locations.

Participants' settings appeared to reach stable states by the final week of the study, during which only five made changes to their rules. For each of these participants, the change consisted of either adding one friend to an existing rule (one participant in the rule condition and one in the profile condition) or creating a new rule for one friend (one in the rule condition and two in the profile condition).

Examining subjects' rules in quantity, those subjects in the rule condition appeared to start with slightly more than those in the profile condition (means of 3.1 and 2.75 rules,

respectively). However this difference was not statistically significant ($p=0.19$, using a one-tailed unmatched t-test). By the end of the study, these means had risen to 3.3 and 2.9 respectively, though again no significant difference was found ($p=0.19$). The profile-centric wizard produced between zero and two rules for the user, although the user was given the option of producing more rules afterward. No similar prior constraint existed for the rule-centric wizard, but the outcome chiefly in terms of rule quantity appeared to be roughly the same across both conditions.

Prior to the study we had expected trends to appear in subjects' sharing rules over the three-week period, and we wished to determine whether those trends would differ by condition. Figure 7 shows the composition of subjects' rules for each of the three weeks. In both figures, *location*, *time*, and *location and time* are mutually exclusive, but *all friends* does not share this property (e.g., a rule counted as *all friends* rule could also be a *location* rule). Many of the week-to-week trends are slight, involving the action of just one or two participants, but a few stand out as substantial. For the rule condition, the drop in the mean number of *all friends* rules from Week 1 (0.71) to Week 3 (0.32) was significant ($p=0.02$); for the profile condition, a similar drop was observed, but it was not statistically significant ($p=0.14$). The increase in *time* rules from Week 1 to Week 3 was also significant for the rule condition ($p=0.04$) but

not for the profile condition (p=0.11). These differences suggest that subjects in the rule condition felt a somewhat greater need to further qualify their rules after the introductory session than those in the profile condition.

DISCUSSION
Overall, these results are consistent with the conjectures that users in the profile condition would share more than those in the rule condition, and that their satisfaction would be roughly equal. Although profile condition participants showed signs of being less satisfied with their settings during the first week, by the end of the study the gap in satisfaction between the conditions had disappeared. Our observations suggest that the profile wizard performed worse initially in capturing users' preferences, but profile condition participants' subsequent edits to their settings brought them to the same level of satisfaction.

Equality in sharing, however, was less evident: participants in the profile condition appeared more comfortable sharing their location than those in the rule condition. One plausible explanation for the difference in comfort may concern the relative cognitive burden associated with specifying rules. Although participants in the rule condition could create fully expressive policies, doing so is complex, and may leave participants with less of a clear understanding of how the system will behave. Conversely, default profiles are not fully expressive, but they are easy for users to understand and adapt to. The observed difference in sharing and the similarity in satisfaction suggest that the initial set of privacy rules a user is presented with has a lasting impact on their sharing behavior.

Weekly trends in the auditing data and declining numbers of settings changes suggest that users in both conditions gradually became acclimated to their privacy settings over the study period. One of the intended effects of the nightly auditing was to keep subjects aware of the effects of their preferences, so that they would be inclined to change those preferences if they did not seem satisfactory. Moreover, during the introductory session, subjects were advised to change their sharing rules at any time if they were unhappy with them, and after the initial wizards both conditions were exposed to the same interface to make any changes to their rule sets. Responses to the exit survey suggested that participants did not struggle to edit their sharing settings, as nearly all participants agreed or strongly agreed that it was easy to do so. Two participants described Locaccino as "intuitive and simple to use" and "pretty polished". Still, the two conditions retained unequal levels of sharing, suggesting that the wizards had a lasting impact on the level of sharing that users felt comfortable with. We also observed that *request denied, unsatisfied* audits persisted (with some diminishment) through the length of the study. Subjects either found the sharing settings to be insufficiently expressive, or they declined to change their settings to increase their satisfaction. The tendency to acclimate should be carefully noted by designers of interfaces for privacy settings: design decisions chiefly

intended to simplify users' initial privacy choices may have lasting effects on what they deem to be an acceptable level of disclosure.

We observed that subjects in the profile condition appeared to be more engaged in the study than subjects in the rule condition. Those in the profile condition completed substantially more audits and also provided more *not there* auditing responses, perhaps indicating greater attention to the location information shown to them. This greater engagement may suggest that, by capturing users' likely preferences, the privacy profiles provided a more compelling introduction to Locaccino. Designers of privacy controls for online social networks may take note that such tools can encourage greater engagement if they are carefully designed around users' needs.

Although our study was conducted with a continuous location sharing system, we believe our results are applicable to check-in based systems as well. As these systems add features that promote sharing with increasingly diverse social connections, privacy concerns will necessitate more expressive controls for users to manage their preferences. Accordingly, we expect simplifying elements such as privacy profiles to appear in social networks based on both models of location sharing, in an effort to help users quickly arrive at settings close to their ideal preferences.

In summary, the results of this study suggest that efforts to simplify privacy choices can have a significant impact on the levels of privacy that users select. This observation correlates with a similar one by Acquisti and Gross [1] that default privacy settings tend to be retained by users. However, both wizards in the present experiment required the user to make key choices about how they would share their location. It appears that, even when users are compelled to make deliberate choices about their privacy, relatively small differences in the presentation of settings can lead users to arrive at significantly different and persistent sharing choices.

CONCLUSIONS
We have conducted a study of the effects of privacy profiles on the preferences, sharing tendencies, and satisfaction levels of users of location sharing applications. Our results demonstrate that privacy profiles for location sharing settings can have an enduring impact on how users view their privacy, even in light of continuing opportunities to reflect on the sharing outcomes that result from their chosen settings. More broadly, our results suggest that efforts to simplify privacy settings must be undertaken carefully, as such simplification can easily influence the users whom the settings are designed to engage and inform.

REFERENCES

1. Acquisti, A., and Gross, R. Information Revelation and Privacy in Online Social Networks. In *Proc. WPES 2005*, ACM Press (2005).

2. Anthony, D., Kotz, D., and Henderson, T. Privacy in location-aware computing environments. *IEEE Pervasive Computing*, 6(4) (2007), 64–72.

3. Barkhuus, L., Brown, B., Bell, M., Sherwood, S., Hall, M., and Chalmers, M. From awareness to repartee: sharing location within social groups. In *Proc CHI 2008*, ACM Press (2008), 497–506.

4. Barkhuus, L. and Dey, A. K. Location-Based Services for Mobile Telephony: a study of users' privacy concerns. In *Proc. Interact 2003*, 207-212.

5. Benisch, M., Kelley, P. G., Sadeh, N., and Cranor, L. F. 2011. Capturing Location-Privacy Preferences: Quantifying Accuracy and User-Burden Tradeoffs. *Personal and Ubiquitous Comp.* 15(7), 2011.

6. Beresford, A. and Stajano, F. Location privacy in pervasive computing. *Pervasive Comp.* 2(1) (2005).

7. Brush, B., Krumm, A. J., and Scott, J. Exploring End User Preferences for Location Obfuscation, Location-Based Services, and the Value of Location. In *Proc. Ubicomp 2010*.

8. Brown, B., Taylor, A., Izadi, S, Sellen, A., Kaye, J., and Eardley, R. Location family values: A field trial of the whereabouts clock. In *Proc. Ubicomp 2007*, 354–371.

9. Consolvo, S., Smith, I. E., Matthews, T., LaMarca, A., Tabert, J. & Powledge, P. Location Disclosure to Social Relations: Why, When, & What People Want to Share. In *Proc CHI 2005*, ACM Press (2005), 81-90.

10. Facebook: http://facebook.com.

11. Egelman, S., Oates, A., and Krishnamurthi, S. Oops, I did it again: Mitigating repeated access control errors on Facebook. In *Proc. CHI 2011*, 2295–2304.

12. Find My Friends: https://www.apple.com/icloud/features/find-my-friends.html.

13. Foursquare: http://foursquare.com.

14. Gandon, F., and Sadeh, N. Semantic Web Technologies to Reconcile Privacy and Context Awareness. *Journal of Web Semantics* 1(3) (2004).

15. Google Latitude: www.google.com/latitude.

16. Gruteser, M., and Grunwald, D. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. In *Proc. MobiSys 2003*, ACM Press (2003).

17. Johnson, M., Egelman, S., and Bellovin, S. M. Facebook and Privacy: It's complicated. In *Proc. SOUPS 2012*.

18. Kelley, P. G., Drielsma, P. H., Sadeh, N., and Cranor, L. F. User-controllable learning of security and privacy policies. In *Proc. AISec*, ACM Press (2008).

19. Kim, R. Report: Potential Foursquare investors raising eyebrows at paltry revenue. http://gigaom.com/2012/11/21/report-potential-foursquare-investors-raising-eyebrows-at-paltry-revenue/, 2012.

20. Krumm, J. A survey of computational location privacy. *Personal and Ubiquitous Comp.* 13(6) (2009), 391-399.

21. Lin, J., Benisch, M., Sadeh, N., Niu, J., Hong, J., Lu, B., and Guo, S. A Comparative Study of the Location-Sharing Privacy Preferences in the U.S. and China. Technical report CMU-CyLab-12-003: Carnegie Mellon University, 2012. Available at http://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab12003.pdf.

22. Lin, J., Xiang, G., Hong, J., and Sadeh, N. Modeling People's Place Naming Preferences in Location Sharing. In *Proc. Ubicomp,* ACM Press (2012).

23. Lindqvist, J., Cranshaw, J., Wiese, J., Hong, J. I., Zimmerman, J. I'm the Mayor of My House: Examining Why People Use foursquare - a Social-Driven Location Sharing Application. In *Proc. CHI 2011*, ACM Press (2011).

24. Locaccino: http://locaccino.org.

25. Mugan, J., Sharman, T., and Sadeh, N. Understandable Learning of Privacy Preferences Through Default Personas and Suggestions. Technical report CMU-ISR-11-112: Carnegie Mellon University, 2011. Available at http://reports-archive.adm.cs.cmu.edu/anon/isr2011/CMU-ISR-11-112.pdf.

26. Patil, S. and Lai, J. Who Gets to Know What When: Configuring Privacy Permissions in an Awareness Application. In *Proc. CHI 2005*, 101-110.

27. Ravichandran, R., Benisch, M., Kelley, P. G., and Sadeh, N. Capturing Social Networking Privacy Preferences: Can Default Policies Help Alleviate Tradeoffs between Expressiveness and User Burden?. *In Proc of 2009 Privacy Enhancing Tech. Symp*.

28. Sadeh, N., Hong, J., Cranor, L, F., Fette, I., Kelley, P. G., Prabaker, M., Rao J. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal and Ubiquitous Comp.* 13 (6) (2009), 401-412.

29. Tang, K. P., Lin, J., Hong, J. I., Siewiorek, D. P., and Sadeh, N. Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing. In *Proc. Ubicomp 2010*. 85-94.

30. Tsai, J., Kelly, P., Hankes, P., Cranor, L., Hong, J., and Sadeh, N. Who's Viewed You? The Impact of Feedback in a Mobile Location-Sharing Application. In *Proc. CHI 2009*, ACM Press (2009).