Received: 20 October 2020 | Revised: 25 August 2021 | Accepted: 13 September 2021

DOI: 10.1002/asi.24580

RESEARCH ARTICLE

JASIST WILEY

"Not all my friends are *friends*": Audience-group-based nudges for managing location privacy

Isha Ghosh | **Vivek Singh**

Department of Communication, Information, and Media Studies, School of Communication and Information, Rutgers University,
New Brunswick, New Jersey, USA

Correspondence
Isha Ghosh, Department of Communication, Information, and Media Studies, School of Communication and Information, Rutgers University,
4 Huntington Street, New Brunswick, NJ 08901, USA.
Email: isha.ghosh@rutgers.edu

Abstract

The popularity of location-based features in social networks has been increasing over the past few years. Location information gathered from social networks can threaten users' information privacy through granular tracking and exposure of their preferences, behaviors, and identity. In this 6-week study ($N = 35$), we investigate the effect of "audience-group"-based interventions on Facebook check-in behavior of participants. These "audience-group"-based nudges help close the gap between the users' *perceived* audiences and those that are *permitted* to view their check-ins. The nudges remind users that their real-time location information may be visible to a larger group of friends than they expect. Based on both quantitative and qualitative data analyses, we report that reminding users of the unexpected audiences that have access to their location check-ins could be a promising way to help users manage their privacy in online location sharing. These findings motivate several recommendations for app designers as well as information privacy researchers to better design and evaluate location sharing in online social networks.

1 | INTRODUCTION

Location-based services and online social networks (OSNs) have grown rapidly in recent years. As a result, many social networks like Facebook allow users to "check-in" to a particular location, which allows the sharing of a user's location within a group of friends. However, this sharing of location data poses significant risks of violation of personal privacy. Wicker (2012) classifies areas with high risks for users as home location, doctor's location, and frequently visited locations. Information about these places can offer much more additional information about the user and can easily be gathered from OSNs like Facebook. Studies have shown that individuals overwhelmingly value location information over media (photos, videos) and communications (messaging friends) (Almuhimedi et al., 2015; Fuller, 2019; Staiano et al., 2014). Location information posted and shared on OSNs are bundled with social information about the user. This combination of real-time spatial information, profile information, and social network information allows for the inference of not only where an individual is but who they are with and potentially where they are likely to be at a particular date or time. For instance, if an individual is prone to check-in at a movie theater with the same group of friends every weekend, a potential adversary could easily infer their future presence at the movie theater and act accordingly. Analysis of a user's location history may reveal the number of times they went to the hospital within a certain time period, thus providing a certain amount of sensitive health-related information. If a user checks in at a particular spot regularly, it may help infer their home location as within a radius and knowledge of home location may allow implications about annual income level, race, or ethnicity to be drawn.

Facebook's design of friend lists also "flattens" the friend hierarchy, giving all friends the same level of access by default, leading to context collapse (Raynes-Goldie, 2010).

Information shared within an individual's Facebook friend group is heavily dependent on the way Facebook and its users interpret the meaning of online friendship and the social norms that go with it. For instance, a recent report finds that the average Facebook user has 338 friends; however, only 28% of this number are considered close or genuine friends (Osman, 2021). This implies that concerns over privacy of information posted on Facebook are highly contextual, that is, information that is easily shared with a friend may be extremely sensitive to a boss or colleague (Mantouvalou, 2019). Viewing this through the lens of Nissenbaum's analysis of privacy in terms of "contextual integrity" (Nissenbaum, 2004) allows us to realign the privacy problem to one of the maintaining context-specific information flow. The same information could be perceived as sensitive or not based on the audience it is being shared with. The communication privacy management (CPM) theory specifically discusses the management of information sharing within different audiences (Petronio, 2012). The CPM theory conceptualizes privacy as an iterative process of creating and re-creating boundaries around information disclosed within a group. However, the Facebook interface does not make it easy for a user to establish and maintain the many separate audience groups that characterize offline life. If these separate audiences are difficult to explicitly create on Facebook, then managing information flows among them is likely to be very difficult.

There are also more explicit harms to an individual's location information being shared within a large group of known and unknown friends (Min, 2016). The user can be constantly tracked or stalked by anyone in their friends list potentially resulting in physical or emotional harm. Statistics on cyber-stalking show that almost 88% of Facebook users spy on their romantic ex-es via Facebook (Short, 2016). While Facebook has privacy and security settings that can be used to protect against cyber-stalking, these tools are often ineffective when the infraction is committed by the user's "friend."

Studies on Facebook "friending behavior" have found that accepting friend requests is a highly contextual process and most users are likely to friend people with whom they had shared interests or mutual friends but never met (Eslami et al., 2015; Li & Kobsa, 2020). Research also states that social media users underestimate their audience size, guessing that their audience is just 27% of its true size (Bartsch & Dienlin, 2016; Croom et al., 2015). This implies that at any given time, a group of (potentially unknown) people can access an individual's real-time locations. Gaining permissions to view an individual's real-time and past locations would be of great value to malicious users that become part of the individuals' friend network. These "friends" with access to user data can combine location data with other information, such as the date and time of posting, device posted from, and so on, to trace, in real time, the movement of a user. Similarly, "friends" with unfriendly intents may be facilitated in their activities (from stalking to kidnapping to domestic violence) as geolocation data may reveal personal information such as home, work, and school addresses. This implies that even within a closed group of friends, the audience that has permissions to view the user's location information is often much larger than expected or wanted.

There are, therefore, multiple scenarios in which disclosing location information to other individuals on a "friends" list could lead to harm, and this work aims to prevent such harm. Specifically, in this article, we test the effects of audience-based nudges on location disclosure in OSNs. Through statistical analysis of participants' behavioral data and follow-up interviews, we find evidence that nudges reminding individuals of their *permitted but unexpected* audience can have significant effect on their location disclosure behavior and reduce unintended disclosures and regret.

2 | RELATED WORK AND RESEARCH QUESTIONS

Research in the fields of psychology, behavioral economics, and behavioral decision making has discussed the role of biases and heuristics such as "nudges" that influence privacy decision making (Wang et al., 2013). Nudges are "soft-paternalistic" behavioral interventions that do not restrict choice but attempt to account for bounded rationality in decision making (Thaler & Sunstein, 2009). Within privacy and security contexts, researchers have examined the use of nudges to empower users to make better privacy decisions and reduce information asymmetry in online settings (Almuhimedi et al., 2015; Wang, Leon, Chen, & Komanduri, 2013). Other nudging examples include hints on encouraging users to create stronger passwords in online sign-up interfaces (Forget et al., 2008) and security dialogs for warning users about potential malicious email attachments (Brustoloni & Villamarín-Salomón, 2007). Most relevant is the work of Wang et al. (2013) who proposed three intervention mechanisms (audience, timer, and sentiment) that could be integrated into Facebook. The audience nudge presents the user with profile pictures of five of the users' friends with the notification that the post would be visible to these friends. The second mechanism, timer nudge, includes a time delay of 10 s before the users' posts on Facebook. The third nudge, sentiment nudge, gives the user feedback on whether the post would be perceived as positive or negative.

Although Wang et al.'s (2013) nudge uses the idea of audiences, it does not consider the different audience groups within an individual's friends list. This implies that the nudge is based on the idea that the user has similar privacy concerns about everyone in their friends list. As explained by Nissenbaum's (2004) theory of contextual integrity, there is no universal privacy norm. Taking this into account, we split audience nudges into permitted audience nudges and perceived audience nudges. This small but highly significant difference allows us to gain an understanding of the contextual nature of user's privacy concerns. We also focus our research purely on location privacy. As outlined in the previous section, location information on Facebook has specific features that make it highly sensitive to users.

Most Facebook users have a large friend networks consisting of more than few hundreds of friends; however, the number of people they interact with regularly is much smaller (Ahmed et al., 2019). Scholars have found that individuals are only able to recollect a small percentage (<30%) of their Facebook friends (Bernstein et al., 2013; Croom et al., 2015) and stronger ties are more easily recalled than weaker ones (Brewer, 2000; Brewer & Webster, 1999). For instance, Brewer (2000) finds that individuals are more likely to recall people who they know very well and had seen at least once in the past week. This implies that people are likely to have a biased recollection of the size and composition of their audience that is very different from their actual audience. This *perceived* audience is very different (and much smaller) from the audience *permitted* to access this information. These two distinct audiences, that is, the audience the user perceives versus the audience that is permitted to view their posts, have been termed as expected and intended audiences in research literature (Stutzman & Kramer-Duffield, 2010; Taddicken, 2014). In their work on different audiences existing within an individual's friend network, Stutzman and Kramer-Duffield (2010) explored the cognitive processes and mechanisms through which audience segmentation occurs. The audience-based nudges in this study use this biased perception of audiences to help individuals manage their location privacy. To the best of our knowledge, *our work is the first attempt to use interventions based on (permitted and perceived) audience groups to help users better manage their location privacy*. The first research question and hypothesis we examine are:

> **RQ1.** Can audience-group-based nudges reduce location information disclosure in OSNs?

> **H1.** People receiving reminders of their permitted audience will have lower number of check-ins compared to those receiving reminders about perceived audience.

It is also important to address the notion of control over information when discussing location privacy. A longitudinal analysis of information sharing on Facebook shows that as users grew more familiar with the interface and settings, individuals' consciously enacted boundary management strategies like making a profile "friends-only" to protect their information (Wisniewski et al., 2015). However, this notion of control can also be counter-intuitive where a perception of control over one's personal information results in the disclosure of more sensitive content (Brandimarte et al., 2013). For instance, the knowledge that real-time location is only being shared with a closed group may result in a false sense of security motivating higher rates of disclosure. In such a scenario, we hypothesize that a reminder of the permitted audience will cause individuals to exert greater control by changing visibility settings on their check-ins.

> **H2.** People receiving reminders of their permitted audience will be more likely to change visibility settings on their check-ins compared to other groups.

Information about a user's location can also be gathered from data entered in a user's Facebook profile or in response to questions such as "Where do you live?", "Where did you go to high school?", and so on. This combination of spatial and social data allows for the creation of large-scale, longitudinal data sets, which could result in lifestyle profiling, demographic bias, and surveillance, with considerable privacy implications. Most users upload location information on their user profiles without clarity on how much of this information is visible and to whom. This information is often tagged onto status messages or user posts leading to unintentional location disclosure. We therefore hypothesize that reminding individuals of this unexpected audience will encourage them to re-check their profiles and make changes if needed.

> **H3.** People receiving reminders of their permitted audience will be more likely to make changes to location settings on their profiles compared to other groups.

Although knowledge of the audience is closely related to managing information privacy online, concerns associated with unexpected location sharing (e.g., being stalked) make the question of audiences even more pertinent. For example, a user may be comfortable using the check-in functionality to communicate their locations to family or close friends, but when these mobility patterns are accessible to a larger group of (not very well-known) friends, the consequences can be dangerous. This implies

that the frequency and intensity of interactions vary across different audience groups in an individual's friend networks. Being reminded of these audience groups might result in changes in the quality and quantity of interactions. We, therefore, present a second research question to investigate the effect (if any) that nudges have on an individual's interactions with their friends' network.

> **RQ2.** Do audience-group-based nudges influence users' interactions with their audiences in OSNs?

3 | STUDY DESCRIPTION

We use Facebook as a testing application domain for this study because of its popularity and complexity of privacy issues associated with it. A unique affordance of Facebook is the ability to form reciprocal connections with "friends." On Instagram and Twitter, there is no option to "friend" people, one can only "follow" and possibly be "followed" back. The modalities of communication also differ across different sites. Social networks like Instagram emphasize visual image sharing, and location-based sites like Foursquare are specifically designed to share location; however, Facebook provides the largest array of functions, including text-based posts, photo sharing, and location sharing. Given the focus of this study on managing location information via audience-group-based nudges, Facebook was thought to be the most appropriate application domain.

Participants for the study were recruited using flyers, email announcements, and social media posts in 2018 in an area surrounding a university in the north-eastern region of the United States. The participants needed to (a) be between 18 and 75 years of age; (b) have a minimum of three Facebook check-ins in the last 2 weeks; and (c) use an Android smartphone. This can hence be considered a convenience sample, which was considered acceptable, given the exploratory nature of this work.

We conducted a 6-week between-subjects study with three experimental conditions (a) permitted audience condition: reminding individuals of the audience that is permitted to view their check-ins, which includes people in their friends list that the individual does not interact with often (e.g., past friends, acquaintances); (b) perceived audience condition: reminding individuals of their perceived audience, that is, people that the individual interacts with frequently (e.g., close friends, classmates); (c) baseline condition: where we simply use icons with the message that location information shared on Facebook is visible to everyone in their friends list.

We used a mixed-method study using quantitative and qualitative methods. Participants filled in an online questionnaire about their concern over online location sharing, and frequency of their social media usage before and after interventions. We provide a more detailed explanation of this data in the description of Phase 1 of the study. We compared the before–after location-sharing behaviors using statistical tests to test if the intervention had any effect on location disclosure (RQ1) or frequency and intensity of Facebook usage (RQ2). We also conducted exit interviews to gain a contextual understanding of the results from quantitative tests. In the next few sections, we present a timeline and detailed description of the study.

3.1 | Timeline

This study lasted for a total of 45 days and was divided into three phases (roughly 2 weeks each) as shown in the timeline (Figure 1).

3.1.1 | Phase 1

We invited participants to the lab to read and sign the consent form as well as confirm that they actively used the location check-in feature on Facebook. We then asked participants to fill in an online questionnaire that asked for a count of their total check-ins in the last 2 weeks (i.e., 2 weeks prior to the date of filling in the questionnaire), the number of Facebook friends they had, frequency of Facebook use, updates made to their profile, and location settings. We included a question asking if participants were concerned about the information they posted online and questions about any updates made to Facebook privacy settings before starting the study. We asked participants "Are you concerned about information available about you online?" This question was included as a sanity check to verify that no single intervention group contained a majority of privacy-sensitive individuals.

We also asked participants "Have you made any edits to your default Facebook privacy settings? If yes, please describe the edits" and "In the "About" section of your Facebook profile, do you have the "lives in" section updated?" to help analyze H2 and H3.

We asked participants to open the Facebook app on their phones while filling in the questionnaire so they could provide an accurate count of their check-ins, total number of Facebook friends, and profile information. As undertaken by similar studies in the space (Li & Kobsa, 2020; Wisniewski et al., 2015), we ask questions not only related to privacy but also general questions about length of Facebook use, frequency of usage, network, usage

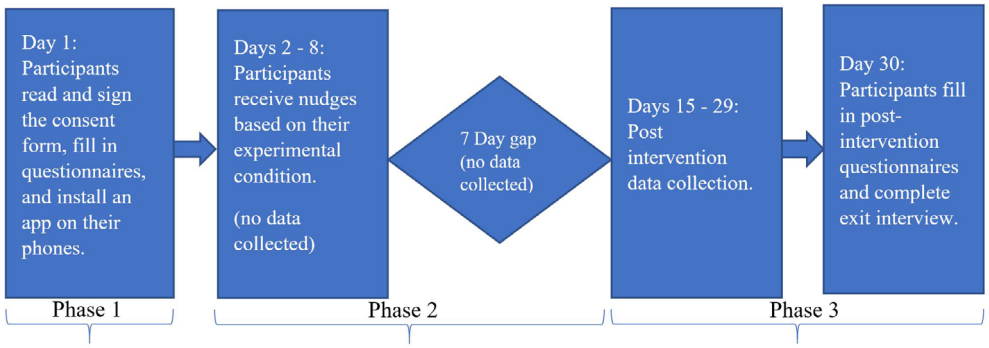

FIGURE 1 Timeline of the study

patterns, and overall profile. Variables collected from this survey (check-in visibility, edits to Facebook profiles) are used in the analysis of Hypotheses 2 and 3. We purposely asked about different aspects of their Facebook profiles as we did not want participants to think deeply about location information privacy in the first phase.

3.1.2 | Phase 2

After completing Phase 1 of the study, participants were randomly split into three different groups. Each group was asked to download and install an app on their Android phones that would send nudges (based on the experimental condition) for a period of 7 days (details of these nudges are explained in the following section). Participants in all experimental conditions received the nudge at the same time every morning (for 7 days) on their mobile phones. The intervention was delivered as a phone notification which participants could view and then click to dismiss. If a participant ignored the notification, it would be stored in the phone status bar until participants viewed the notification and clicked Okay. After installation, participants needed to login to Facebook via the app only once. After this step, they could continue to open and use Facebook as before without interacting with the app. During the first login, the app accessed participant's Facebook friend list and collected profile pictures and names of people in the users' acquaintance suggestions page or close friends' list depending on the experimental condition. For the baseline condition, the app did not access any data but simply used pre-set icons for the intervention.

The app was designed to be used on Android phones and installed via a link provided to users. We planned on building a phone app as people typically use the check-in feature on their mobile phones. We also designed the app such that it only needed a one-time access to the participant's Facebook account. While the app needed to be installed (to deliver nudges), participants had the flexibility to use Facebook on their mobile phones as usual. They did not need to go through the app again. We did not want people who were not part of the study to download or install the app, so the app was not hosted on an app store. While we recognize the difference in data protection and security standards between iPhones and Androids, these securities are mostly designed to protect data from being accessed by malicious apps. Our study, on the other hand, focuses on the privacy implications of location sharing within a user's Facebook friend network; that is, information that the user voluntarily discloses to an audience. Such disclosure does not receive additional protections from the Apple iOS, and therefore, the use of an Android only app was considered acceptable for this exploratory study.

Experimental conditions

We now explain the three different experiment conditions used in this study.

Permitted audience condition

This nudge was designed to remind participants that people with whom they have not interacted with for a while are also part of their friends list and have access to their check-ins. Facebook, by default, provides acquaintance list suggestions to users. These are people in the users' friends list that have not frequently interacted with the user (Coens, 2012). The permitted audience condition nudge displays the profile pictures and usernames of three different people from the acquaintance list to the participant every day (Figure 2).

Perceived audience condition

This condition reminds participants of the perceived audience to their location check-ins. This nudge shows the profile pictures and usernames of three friends that the user interacts most with. The Facebook algorithm is designed to show a snapshot of the friends that the user interacts with most on the left sidebar of their profile page (Neal, 2015). Choosing friends from this list ensures that the nudge reflects the user's perceived audience (Figure 3).

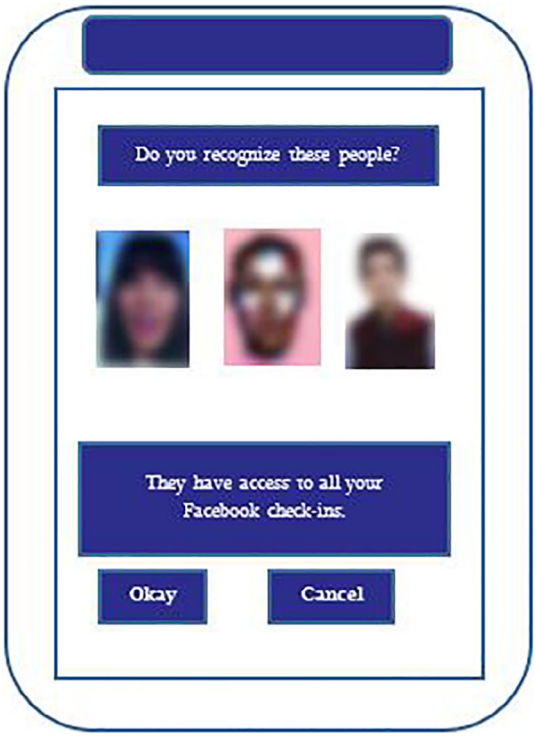

FIGURE 2 Permitted audience condition (images used are for representation purpose only)

FIGURE 3 Perceived audience condition (images used are for representation purpose only)

Baseline condition

Randomized control trial studies have discussed the use of no-treatment groups as well as placebo groups to compare with treatment groups (Cairns, 2019). Studies comparing a treatment and no-treatment group can show that keeping other factors constant the treatment (or intervention) has an effect on the outcome variable. (Cairns, 2019). However, there is an additional concern that an intervention alone, regardless of the content of the intervention, may have a measurable effect (i.e., the placebo effect). It is therefore possible that simply receiving an intervention has some effect on the individuals' disclosure behavior. While there is a case to be made for using a no-treatment or placebo baseline, we chose to use a placebo condition in order to minimize any confounding factors. The placebo intervention closely mirrors the treatment interventions while hiding crucial audience-group information allowing us to measure the effect of audience-group-based interventions on location disclosure.

Research has also documented that participants are more receptive to visual information rather than textual information (Henning & Ewerth, 2018; Singh et al., 2021). To avoid any confounding factors by changing the mode of information presented to the participant, we created a baseline nudge that uses the same combination of visual and textual elements as the treatment groups. This allows us to isolate and measure the impact of audience-group-based reminders on location sharing in OSNs (Figure 4).

3.1.3 | Phase 3

The main purpose of this study is to measure the effect of interventions on location information sharing. During the first phase of the study, we asked participants for a count of the number of check-ins made 2 weeks prior to the start date. In order to compare this with post-intervention data, we invited participants back to the lab 2 weeks after a 7 day gap period and asked them to fill out a similar questionnaire asking for a count of check-ins, number of Facebook friends, frequency of Facebook use, and any updates made to their profile and location settings. After filling the questionnaire, we performed an exit interview where we asked participants about their experience with the nudges, any differences in Facebook use, and overall thoughts on the study. We used open-ended questions, so participants were able to speak at length about interesting aspects of the study. Participants who completed the exit interview were considered to have successfully finished the study.

3.2 | Participants

We invited 42 prospective participants via email to the study location to read and sign the consent form and participate in the study. However, due to some participants dropping out, we had a total of 35 participants who completed the study. We had a similar distribution of age and gender in all three conditions. Participation in this study was voluntary and participants received a total of $25 for completing the study. All personnel involved with the study underwent human subject training and IRB certification. All data were anonymized before analysis.

A summary of the data collected in the 2 weeks before interventions and 2 weeks after interventions is shown in Table 1. On an average, participants across all three conditions had ~7 check-ins in the 2-week period before receiving the nudges and ~5 check-ins in the 2-week period after interventions. We did not measure the check-ins made when participants were receiving the nudges or during the gap period to avoid skewing the results (Table 1).

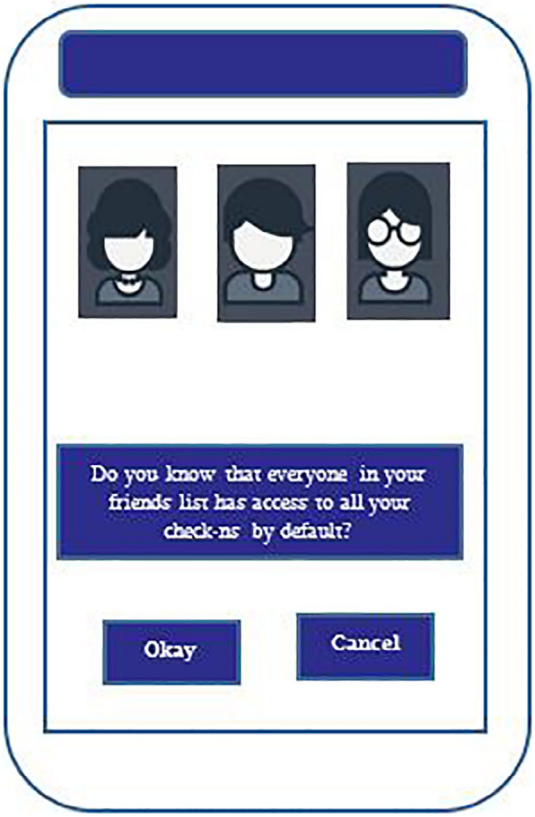

FIGURE 4 Baseline intervention

TABLE 1 Average number of check-ins before and after receiving nudges

Experiment conditions	Check-ins	
	Before	After
Permitted audience group	6.38	3.38
Perceived audience group	7.41	5.92
Baseline group	5.5	7

3.3 | Measures

We asked participants to fill in a demographic survey and a questionnaire about their Facebook usage and location check-in habits at the start and end of the study. Participants were asked to provide a count of their check-ins in the last 2 weeks, total number of Facebook friends, frequency of Facebook use (number of hours a day), as well as open-ended questions about their comfort level with Facebook privacy settings, and location information shared on their profiles. In order to quantify the effect of nudges on location sharing habits, we identified the following metrics in the Table 2.

4 | RESULTS

There was a total of 35 participants with 13 in the permitted audience group, 12 in the perceived audience group, and 10 in the control group. Fifteen participants were full-time graduate or undergraduate students while 13 were employed full time and 7 worked part time. There was a distribution of males and females of different age groups and races in all three experimental conditions. A visual representation of the demographic data is presented in Figure 5.

All participants were active Facebook users and its location check-in feature with 31 out of 35 participants

TABLE 2 Summary of metrics considered in the study

Metric name	Definition
Total number of check-ins	$\sum$ (check-ins) The total number of check-ins made by all participants in each experimental condition.
Changes to visibility settings on posts	*n* (users who made changes to visibility settings of current or past check-ins) Measures the change in visibility settings of each location check-in.
Lives In section empty or restricted	*n* (Users who made changes to location settings in their profile) The number of users who removed or edited privacy settings on location information displayed in their profile (e.g., Hometown, lives in)

reporting that they checked Facebook at least once a day and all participants had at least 3 check-ins in 2 weeks. We wanted to make sure that there were no significant differences in participant privacy concerns that may impact their information disclosure. We therefore performed a preliminary analysis comparing participant privacy concerns over the three experiment groups. Five out of 12 participants in the permitted audience condition, 4 out of 12 participants in the perceived audience condition, and 4 out of 10 participants in the baseline condition said they were concerned about information shared on Facebook. An analysis of variance (ANOVA) test across the three groups yielded a nonsignificant difference ($p > .05$) in privacy concerns across the three experiment groups.

Roughly 40% participants had over 8 check-ins in a 2-week period and nearly 60% of the participants reported that they rarely posted "content" on Facebook or only posted on special occasions. We asked the participants "What posts would qualify as being content posts?" and their responses suggest that status messages, posting photos or videos were "content" while forwarding messages, liking pages or posts, and notably, checking-in, was thought of as browsing Facebook. A summary of the before–after intervention data is shown in the Table 3.

As seen from the Table 3, though the nudges were designed to create awareness about location check-ins

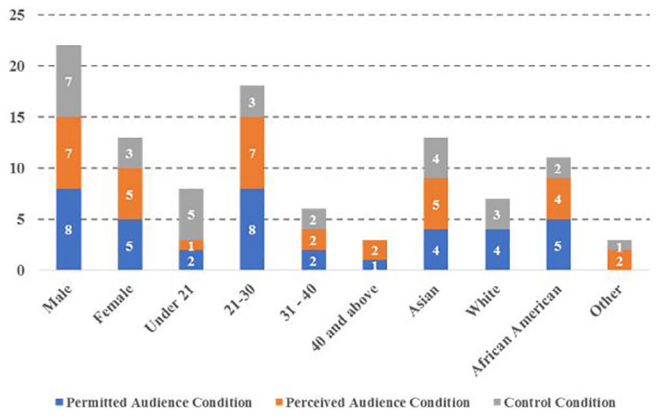

FIGURE 5 Descriptive statistics of the sample population

participants also made other changes to their Facebook profile and settings (especially those related to location). We describe these changes in detail in the next section. Next, we present an analysis of the research questions and hypotheses (Table 4).

4.1 | Location disclosure

The goal of this study was to understand the effect of audience-based nudges on Facebook -ins. We examined multiple hypotheses (Table 4) and here we present the results for our first hypothesis: people receiving reminders of their permitted audience will have lower check-ins compared to other groups.

As seen in Table 3, there was a 46.98% decrease in posting check-ins in the permitted audience-group condition compared to 20.22% decrease in the perceived audience-group condition and 7.69% in the control group (increase in number of check-ins). In order to check if these differences were significant, we needed to run a mixed-model ANOVA with time (before-, after-) as within-subject factors and experiment condition (permitted, perceived, and control) as between-subject factors. The results show a significant effect of the intervention on check-ins $F(2, 32) = 5.017$, $p < .05$. Pairwise comparisons show a significant difference in the reduction of check-ins between the permitted and perceived audience groups ($p < .05$) as well as between the permitted and baseline groups ($p < .01$) There was no significant difference ($p > .05$) in reduction of check-ins between the perceived audience condition and the baseline condition. We find that nudges about the permitted audience significantly reduced the number of location check-ins when compared to nudges about the perceived audience or a baseline nudge and H1 holds true.

During exit interviews, we found that participants were largely aware of their location sharing and were enthusiastic about the benefits gained from sharing information (e.g., social bragging, taking advantage of promotional offers). However, they expressed concern when they were notified of acquaintances in their friends list who had access to their location information. This is

TABLE 3 Summary of results across three experimental groups comparing before and after conditions

	Permitted audience ($n = 13$)		Perceived audience ($n = 12$)		Baseline group ($n = 10$)	
	Before	After	Before	After	Before	After
Total number of check-ins	83	44	89	71	65	70
Changes to visibility settings on posts	0	5	0	2	0	0
Lives In section empty or restricted	1	10	2	8	1	2

TABLE 4 Result of hypotheses testing using analysis of variance

	Key variable	*p* value	Significance
H1	Difference in check-ins	.013	Significant
H2	Difference in post visibility	.067	*Not* significant
H3	Difference in profile edits	.003	Significant

consistent with Altman's (1975) conceptualization of privacy as a process of negotiation through which individuals control access to themselves. Participant MA180 stated "I check-in like almost daily, and I know the people that like these posts [...] It was weird and kinda embarrassing to see my friends from high school and stuff seeing me at my fancy new job. Maybe they felt like I was bragging or showing off to them. That's like the opposite of me [...] I realized I don't need to do that every single day." Another participant FA120 mentioned "Most of my check-ins are at the gym.... it's my way of staying on track with my New Year resolution but seeing random people that could know where I was everyday at a particular time really creeped me out. That was something I stopped doing. I just put a check mark on my calendar now." We find that providing nudges reminding participants of this unexpected audience, rather than nudges asking them to stop using the check-in feature, results in a re-evaluation of their information sharing habits and reduces disclosure.

4.2 | Check-in visibility setting

The second hypothesis states that reminding individuals of their permitted audience will cause them to change visibility settings on their check-ins. For a friends-only profile, the default visibility of location check-ins is all friends. Although OSNs allow users to edit this default, it is not something users often think of when posting. We asked participants if they made any edits to this default setting of their current or previous check-in posts after receiving the interventions.

From data collected in the exit interview, we found that 5 out of 13, that is, 38.46% participants in the permitted audience section either deleted previous check-ins or changed visibility to "Only Me", "Friends except...", or particular friends for past and current check-ins (Table 3). This percentage decreased to 16.67% in the perceived audience group while no participants in the control group changed visibility settings. All these changes were made in the weeks following the interventions.

In order to compare the number of people who updated their location settings across the three experiment conditions, we use a nonparametric mixed-model ANOVA (Brunner et al., 2002). We first coded participants who changed the visibility settings of their location posts as Y and those who did not make a change as N. The model described by Brunner et al. (2002) can handle categorical data as well as between- and within-subject factors. *This test however, did not find a significant relationship* ($p > 0.05$) *between the intervention and the number of people that changed their visibility settings.*

The qualitative exit interviews helped add a layer of nuance to these findings. Participants across the three experiment conditions mentioned the use of check-in as a "quick way of staying in touch..." with all their close friends or family members rather than the effort of individually contacting each person. Counter-intuitively, participants found it easier to reduce their check-ins rather than edit the visibility settings on each post they made. Users would need to have custom lists of friends or family they wanted to include in their audience and then edit the settings of each post to include or exclude friends from these lists. A user who prioritizes time and convenience might simply skip checking-in rather than configuring these settings.

4.3 | Location information on profile

Facebook allows users to put in their hometowns, current location, school/work details, and contact information along with a short bio on their profiles. The third hypothesis: people receiving reminders of their permitted audience will be more likely to make changes to location settings on their profiles, tests the effect of nudges on location information contained in user profiles.

In the pre-intervention session, we asked participants if their current town of residence was visible in the "Lives In" section of their Facebook profile. This field is automatically updated by Facebook and the user location is often tagged to any posts made by the user therefore revealing their location often without the user being aware of it. Individuals' have to navigate to the specific section and either delete the information or change visibility settings in order to remove location information from their profile. In the post-intervention session, we asked participants if they had made any changes to the "Lives In" section of their profile. For the analysis, we only considered cases where participants made changes to the "Lives In" section after receiving interventions. There were 4 out of 35 participants (1 in permitted audience condition, 2 in perceived audience condition, and 1 in control condition) (Table 3), who had already removed this information before joining the study. We therefore excluded their data from our analysis. Of the remaining 31 participants, a total of 16 people (9 in permitted audience condition, 6 in perceived audience condition, and 1 in control condition) removed or restricted visibility of this information during the experiment.

We again needed to compare categorical variables and therefore use the same Y/N coding and nonparametric mixed-model ANOVA as in the previous section. A participant who *removed or restricted visibility* of the Lives In section was coded as Y and others were coded as N. We find a significant association ($p < .01$) between the intervention and changes made to the users Lives In section. Further pairwise comparisons using Bonferroni's corrections show a significant difference between the number of people changing the Lives In section in the permitted audience condition and the perceived and baseline conditions. There was also a significant difference between the perceived audience and baseline conditions.

The results obtained from all three hypotheses point to the multifaceted nature of location sharing in OSNs. It is no longer enough to limit voluntary location disclosure, *users' must also pay attention to the visibility of past location sharing as well as location information automatically included in their profiles and shared within their network*. Together, the analysis helps shed light on our RQ1. Receiving the audience-group-based interventions caused significant number of users to rethink and reduce the frequency of their check-ins. At the same time, it also caused users to rethink the different ways in which their location information could be exposed and take steps (removing location from profile, editing past check-ins, etc.) to better manage location privacy in OSNs.

4.4 | Other effects of interventions

Our second research question was designed to help us dive deeper into the effects of the nudges and uncover latent effects (if any) of these notifications. We asked if audience-group-based nudges affected users' interactions with their OSN audiences. This study brought out some innovative practices adopted by participants to protect location privacy when reminded of the audience of their disclosure. While asking participants about the frequency of OSN use, we also asked participants to fill in the number of Facebook friends they had. Table 5 shows changes made to the friends' list before and after receiving interventions.

As seen in Table 5, there was a reduction in the number of friends' participants that had across all three groups. Participants in the permitted audience condition had the most change in the number of friends as well as the most people (7 out of 13) that edited their friends list. Two out of 12 participants in the perceived audience condition and 1 out of 10 in the control group made similar edits resulting in a reduction of 59 and 107, respectively.

The data were found to be non-normal, and therefore we again use the nonparametric mixed-model ANOVA (Brunner et al., 2002). Although this test did not show significant associations ($p > .05$), we cannot deny the effect of the intervention on users' interactions within their network. Wisniewski et al. (2015) highlighted the connection between interactions with a large number of friends and information disclosure in their work. The authors found that users with high levels of friend expected Facebook's interface to protect their information from unwanted others. The results from our study showing a pruning of friend lists may be related to an awareness that this was not the case. The audience-group interventions may have resulted in users feeling uncomfortable about the number of people in their friends list and motivated them to edit it down to a more manageable number. While we did not investigate if there were any external circumstances that resulted in pruning friends list, the data shown in Table 5 point to a connection between the intervention and edited friend lists. This unforeseen effect of the intervention needs to be studied further and presents an interesting direction for future work.

Previous research has shown that individuals are only able to recollect a small percentage (<30%) of their Facebook friends (Bartsch & Dienlin, 2016; Croom et al., 2015). Our study showed similar results with participants sometimes expressing surprise when they received notification. Notably, participant MA150 in the acquaintance condition was unaware that the notifications they received were from their friends' list. Another participant MA210 mentioned adding friends of friends, and "anyone that sent [me] a friend request and looked sort of decent" to their friends list. However, not recognizing most of the friends shown in the daily

TABLE 5 Average number of Facebook friends before and after receiving the intervention

Experiment conditions	Average number of friends		Average difference in number of friends	No. of users who edited friend lists
	Before	After		
Permitted	1,300.69 ($\sigma = 617.15$)	1,244.07 ($\sigma = 619.69$)	56.61	7
Perceived	1,105.75 ($\sigma = 473.28$)	1,100.83 ($\sigma = 469.45$)	4.91	2
Baseline	1,063.10 ($\sigma = 566.34$)	1,052.40 ($\sigma = 566.04$)	10.7	1

Note: *SD* σ is shown within parentheses.

interventions led them to edit down their friends list to people they actually knew and wanted to stay in touch with. Participants in the perceived audience-group condition on the other hand, were unsurprised by the notifications since the profiles they were shown were more likely to be from their active friends list. The information contained in the intervention matched the users perception of the audience and therefore did not result in any additional audience management changes.

Our interventions were designed to focus on location-sharing behavior; however, we found that participants tended to adopt a range of privacy protection practices according to their comfort level. We also found that the permitted audience nudge caused users to re-evaluate other aspects of their online selves as well. For some participants, the simplest way to manage location privacy was to clean out their friends' list. Within a contextual perspective, privacy of information depends heavily on the environment it is shared in. While users may have been aware that they had a large friends' network, the nudge presented the user with specific information regarding the presence of unexpected members within this network. This result presents an interesting area for future research. Users may not only limit their posts but also proactively reduce the audience to their disclosure if the platform does not provide strategies for more nuanced sharing. During exit interviews, participants did not report feelings of isolation or unhappiness; in fact, they felt more in control of their information ("This is a good remainder for me to keep like my online stuff on track"—MA0110), their friend network (*they are on my friends list, but they aren't friends, so it felt good to be more selective—FA0170*), and happier with sharing information. While expanding friend networks can result in a number of advantages, often times these networks become too large and unwieldy for users to handle. Balancing these conflicting priorities of network building and personal information management is an important research question for information science researchers, which should be studied in more detail in future work.

5 | DISCUSSION

In an attempt to help individuals better manage location privacy in OSN, this study uses the concepts of permitted versus perceived audience groups to remind participants that their location information may be accessible to a larger audience than they expect. When a user checks-in and uses the default friends-only visibility settings, all of their friends gain access to this information and can potentially misuse it. In this scenario, receiving reminders of different audience groups caused people to "think about the number of people that could see where I was" (MA110), and make changes to their location sharing habits as well as privacy settings. Location information is often sensitive and individuals' may not be comfortable sharing this information with all their friends. One of the participants expressed this sentiment saying "...if a 1000 people see me forward or repost or like something...that's okay, but when they see where I go clubbing...that's like much weirder," (FA190).

According to Nissenbaum's theory of contextual integrity the *context and environment* in which information exchange occurs is closely related to individual privacy concerns (Wu, Vitak, & Zimmer, 2020). In the case of OSNs, the context and audience to disclosure are both constantly changing. The ability of OSN to record information has made it possible for the audience of a location check-in to grow substantially months and sometimes years after the event. This implies that a person added to an individual's friends list now has access to all of their previous check-ins, compounding audience management issues. For instance, a person may have been regular at a pub during their college years. This may not be representative of the person's current habits, but it may be the impression communicated to future recruiters. In such a scenario, reminding individuals of the unexpected but permitted audience to their location sharing acts as a cue to re-check and sometimes delete previous check-ins. This change in the context in which information was shared can result in users feeling of discomfort about information disclosure. At the time of posting, individuals may not have had enough information or may have underestimated the future impact of their check-ins. As shown by our results, receiving daily nudges about location check-ins motivated users to adopt multiple privacy protection strategies from reducing the number of check-ins to editing visibility settings to deleting past check-ins.

Although our study was conceptualized to help users better manage their location check-ins, we found that receiving daily notifications caused users to think about location sharing beyond check-in posts. A number of users reported thinking about the "Lives In" and "Hometown" information contained in their profile and the visibility settings on this information. An unexpected phenomenon we saw was the number of participants that removed people from their friends' list as an effect of the nudges. One participant stated, "When I didn't recognize the people showing up in the notifications, I decided my Facebook page needed to be cleaned. I removed everyone I had not personally met" (MA240). This also points to the importance of the circumstances in which information is shared and disclosed when discussing individual privacy concerns. Seeing that their location information can be accessed by unexpected and perhaps unknown members of their network led to users feeling concerned not only about the

information that was available but also about the composition of their friends network. For some users, managing their privacy was connected with managing the context in which information flow occurred.

The editing of friend lists is an important implication for the design of future interventions. It is important to note that interventions often have unintended consequences. For instance, in the current study, some of the friendship ties were removed as an indirect consequence of the intervention. This may not be significant at an individual level but viewed over time and across a large population, an intervention to increase privacy awareness could also result in the increase of isolation. In cases where friends were removed without explanation or notice, it may result in a souring of relationships. While none of the users mentioned a feeling of loss, it reminds us to be mindful of any unintended social effects of our interventions in future. We therefore present this result as an additional take-away with implications to the intervention design literature.

Designers should also keep in mind the difference in privacy concerns associated with different modalities of information. OSNs today allow users to share various modalities of information (text, visual, audio, location, etc.), but present users with the same set of privacy protection mechanisms. While a friends-only profile may be an efficient way to manage privacy for text-based or visual information, it was not enough to mitigate concerns over location sharing. As participants noted, they were aware and comfortable with the idea that some of their status updates or forwards would be available to a large audience but when this same audience had access to their location information, the risk quotient transfers from the online to the physical space resulting in heightened privacy concerns. This highlights a profound difference in individual privacy concerns over location versus other communicative information and shows a potential direction for future privacy interfaces.

It was also easier for users to make one-time edits to their location disclosures or friend lists than to remember to edit the audience of each post. While Facebook gives users the option to create custom lists and choose audiences, these settings might be too cumbersome for users to manage. Designers must therefore think creatively about building privacy mechanisms (e.g., geo-filters that is, the audience need to be within a certain radius of the check-in to access the information) specifically to help users manage location privacy.

Finally, designers must always check for the unintended side-effects of interventions. For instance, an intervention that encourages users to avoid frequent check-ins at a pub or bar, may cause them to miss out on making valuable new connections or gaining a sense of connection leading to negative social consequences. While these effects are, by definition, difficult to foresee or predict, it is important for app designers to be aware of the possibilities of unintended side-effects and consider the effects of the intervention at a societal level.

5.1 | Limitations

As with most works, this study is also not free from limitations. Firstly, we would like a larger sample size and to study the effect of nudges over an extended period of time. The data collection was also based solely on recruitment through online channels, and the final sample yielded a large proportion of college-age highly educated respondents. We attempted to have an even distribution of participants in terms of age and gender across the three experiment conditions, however, due to some participants dropping out after the first stage, we had fewer females in the baseline group when compared to other groups. While privacy concerns did not vary significantly across the three experiment groups, future research may find it valuable to examine a wider range of age and education groups.

6 | CONCLUSION

As OSNs become increasingly powerful mechanisms for information sharing, a considerable proportion of users report sharing information and feelings that they later regret disclosing (Wang et al., 2013). Our study of location check-ins on Facebook showed that people often check-in as a "spur-of-the-moment" decision without thinking deeply about the consequences of this disclosure. We also found that people tend to have a mix of different audiences in their Facebook friend lists and are often unable to recall all the different groups that make up their Facebook friends network.

Drawing on existing research in the area of location sharing and privacy nudges, we designed audience-group-based nudges to test the effectiveness of reminders based on permitted versus perceived audiences on an individuals' location disclosures. The results obtained from statistical analysis and qualitative interviews show that the use of unexpected audience in reminders can be a powerful tool to help people avoid regrettable disclosure. The use audience-based nudges could hence be a useful tool for privacy researchers, app designers, and policy-makers to help people protect their location information privacy.

ORCID

Isha Ghosh https://orcid.org/0000-0002-0879-3514
Vivek Singh https://orcid.org/0000-0002-8194-2336

REFERENCES

Ahmed, J., Villata, S., & Governatori, G. (2019). Information and friend segregation for online social networks: A user study. *Ai & Society*, *34*(4), 753–766.

Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L.F., & Agarwal, Y. (2015). Your location has been shared 5,398 times! A field study on mobile app privacy nudging. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 787–796).

Altman, I. (1975). *Environment and social behavior: Privacy, personal space, territory and crowding*. Monterey, California: Brooks Cole, 1975.

Bartsch, M., & Dienlin, T. (2016). Control your Facebook: An analysis of online privacy literacy. *Computers in Human Behavior*, *56*, 147–154.

Bernstein, M. S., Bakshy, E., Burke, M., & Karrer, B. (2013). Quantifying the invisible audience in social networks. *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (pp. 21–30).

Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, *4*(3), 340–347.

Brewer, D. (2000). Forgetting in the recall-based elicitation of personal and social networks. *Social Networks*, *22*(1), 29–43.

Brewer, D., & Webster, C. (1999). Forgetting of friends and its effects on measuring friendship networks. *Social Networks*, *21*(4), 361–373.

Brunner, E., Domhof, S., & Langer, F. (2002). *Nonparametric analysis of longitudinal data in factorial experiments*. John Wiley & Sons.

Brustoloni, J. C., & Villamarín-Salomón, R. (2007). Improving security decisions with polymorphic and audited dialogs. *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS)* (pp. 76-85). ACM.

Cairns, P. (2019). *Doing better statistics in human-computer interaction*. Cambridge University Press.

Coens, J. (2012). See posts that matter to you [Blog post]. https://newsroom.fb.com/news/2012/03/see-posts-that-matter-to-you/

Croom, C., Gross, B., Rosen, L. D., & Rosen, B. (2015). What's her face (book)? How many of their Facebook "friends" can college students actually identify? *Computers in Human Behavior*, *56*, 135–141.

Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K., & Sandvig, C. (2015). I always assumed that I wasn't really that close to [her]: Reasoning about Invisible Algorithms in News Feeds. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. (pp. 153-162). ACM.

Forget, A., Chiasson, S., van Oorschot, P.C., & Biddle, R., (2008). Improving text passwords through persuasion. *Proceedings of the 4th Symposium on Usable Privacy and Security* (SOUPS'08), 1-12.

Fuller, M. (2019). Big data and the Facebook scandal: Issues and responses. *Theology*, *122*(1), 14–21.

Henning, C., & Ewerth, R. (2018). Estimating the information gap between textual and visual representations. *International Journal of Multimedia Information Retrieval*, *7*(1), 43–56.

Li, Y., & Kobsa, A. (2020). Context and privacy concerns in friend request decisions. *Journal of the Association for Information Science and Technology*, *71*(6), 632–643.

Mantouvalou, V. (2019). 'I lost my job over a Facebook post: Was that fair?' Discipline and dismissal for social media activity. *International Journal of Comparative Labour Law and Industrial Relations*, *35*(1), 101–125.

Min, J. (2016). Personal information concerns and provision in social network sites: Interplay between secure preservation and true presentation. *Journal of the Association for Information Science and Technology*, *67*(1), 26–42.

Neal M., (2015). Facebook's magic formula for determining your 9 top friends [Blog Post]. https://motherboard.vice.com/en_us/article/ezp4bj/facebooks-magic-formula-for-determining-your-9-top-friends

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, *79*, 101–139.

Osman, P. (2021). Wild and Interesting Facebook Statistics and Facts (2021). Kinsta. https://kinsta.com/blog/facebook-statistics/

Petronio, S. (2012). *Boundaries of privacy: Dialectics of disclosure*. Suny Press.

Raynes-Goldie, K. (2010). Aliases, creeping, and wall cleaning: Understanding privacy in the age of Facebook. *First Monday*, *15*(1), 1–4 https://firstmonday.org/ojs/index.php/fm/article/download/2775/2432

Short, E. (2016). New prevention orders won't do enough to stop online stalking. *The Guardian* https://www.theguardian.com/commentisfree/2016/dec/12/prevention-orders-wont-stop-online-stalking-amber-rudd

Singh, V. K., Ghosh, I., & Sonagara, D. (2021). Detecting fake news stories via multimodal analysis. *Journal of the Association for Information Science and Technology.*, *72*, 3–17. https://doi.org/10.1002/asi.24359

Staiano, J., Oliver, N., Lepri, B., de Oliveira, R., Caraviello, M., & Sebe, N. (2014). Money walks: a human-centric study on the economics of personal mobile data. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 583-594). ACM.

Stutzman, F., & Kramer-Duffield, J. (2010). Friends only: examining a privacy-enhancing behavior in Facebook. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1553-1562). ACM.

Taddicken, M. (2014). The 'privacy paradox' in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure. *Journal of Computer-Mediated Communication*, *19*(2), 248–273.

Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, CT: Yale University Press.

Wang, Y., Leon, P. G., Chen, X., & Komanduri, S. (2013). From facebook regrets to facebook privacy nudges. *Ohio State Law Journal*, *74*(6), 1307–1334.

Wang, Y., Leon, P. G., Scott, K., Chen, X., Acquisti, A., & Cranor, L. F. (2013). Privacy nudges for social media: an exploratory Facebook study. *Proceedings of the 22nd International Conference on World Wide Web* (pp. 763-770). ACM.

Wicker, S. B. (2012). The loss of location privacy in the cellular age. *Communications of the ACM*, *55*(8), 60–68.

Wisniewski, P., Xu, H., Lipford, H., & Bello-Ogunu, E. (2015). Facebook apps and tagging: The trade-off between personal privacy and engaging with friends. *Journal of the Association for Information Science and Technology*, *66*(9), 1883–1896.

Wu, P. F., Vitak, J., & Zimmer, M. T. (2020). A contextual approach to information privacy research. *Journal of the Association for Information Science and Technology*, *71*(4), 485–490.

How to cite this article: Ghosh, I., & Singh, V. (2022). "Not all my friends are *friends*": Audience-group-based nudges for managing location privacy. *Journal of the Association for Information Science and Technology*, *73*(6), 797–810. https://doi.org/10.1002/asi.24580