Will They Share? Predicting Location Sharing Behaviors of Smartphone Users through Self-Reflection on Past Privacy Behaviors

Muhammad Irtaza Safi*, Abhiditya Jha*, Malak Eihab Aly†, Xinru Page‡, Sameer Patil§ and Pamela Wisniewski*

*University of Central Florida
safim772@gmail.com, abhiditya@knights.ucf.edu, pamwis@ucf.edu

†New York University
mea433@nyu.edu

‡Bentley University
xpage@bentley.edu

§Indiana University
patil@indiana.edu

***Abstract*—Location sharing is a particularly sensitive type of online information disclosure. To explain this behavior, we compared the effectiveness of using self-report measures drawn from the literature, behavioral data collected from mobile phones, and a new type of measure that represents a hybrid of self-report and behavioral data to contextualize users' attitudes toward their past location sharing behaviors. This new measure was based on a reflective learning paradigm, where one reflects on past behavior to inform future behavior. Based on a study of Android smartphone users (N=114), we found that the construct 'FYI About Myself' and our new reflective measure of one's comfort with sharing location with apps on the smartphone were the best predictors of location sharing behavior. Surprisingly, Behavioral Intention, a commonly used proxy for actual behavior, was not a significant predictor. These results have important implications for privacy research and designing systems to meet users' location sharing privacy needs.**

I. Introduction

Mobile users now account for the majority of Internet traffic (52%) [8], and mobile app revenue is projected to reach $188 billion by 2020 [4]. However, the wealth of personal information that mobile apps access has a considerable impact on their usage. A recent Pew study on smartphone app usage in the U.S. reported that 90% of smartphone owners indicated that knowing how their personal information is used is important when deciding whether to install an app [27]. In fact, 60% of app users reported deciding not to install an app because of the personal information it requested, and 46% had uninstalled an app after discovering the extent to which it collected personal information [27]. To encourage app adoption and use, researchers and app designers need to be able to understand and predict people's willingness to disclose various types of personal information to apps and use this knowledge to help align the system with users' privacy needs [39].

However, predicting personal information disclosure online is not straightforward. Research has revealed discrepancies between people's stated concerns and their actual disclosure behaviors. This widely acknowledged 'privacy paradox' [6], [26] has made it difficult to predict user behavior based on stated privacy preferences. Instead, social science researchers have developed scales to predict and explain users' privacy intentions [25], [28], [36], [43]. Yet, for the most part, there is still a gap when mapping these self-reported measures to actual behavior. Meanwhile, computational scientists have taken a different approach of making predictions based on behavioral data, such as the number of apps installed, the kinds of apps installed, the types of permissions granted to the apps, etc. However, these studies have typically used this behavioral data to detect problematic behavior or to alert users regarding potential data leaks [2], [3], [20]. The collected behavioral data has not been used to help users reflect on their past privacy behaviors to inform and guide future privacy decisions to align better with their true privacy preferences [40].

We focus on understanding and predicting location sharing behavior in the context of mobile apps. We compare the effectiveness of attitudinal measures and behavioral data, as well as newly developed measures of attitude toward past privacy behavior, to answer the following research questions:

RQ1: *What factors best predict users' app location sharing behavior for three types of predictor variables?*:
(A) pre-validated perceived constructs from the literature,
(B) scraped behavioral data, and
(C) hybrid measures capturing users' perceptions of their past location sharing behavior?

RQ2: *Across the three types of predictor variables above, what combination of factors best predicts users' app location sharing behavior?*

To this end, we conducted a study with 114 Android users recruited via Amazon Mechanical Turk. We developed an Android app that participants installed on their smartphones. The app measured perceived constructs (via a questionnaire embedded in the app), scraped behavioral data such as the number of apps with permission to access location (collected unobtrusively by the app), and then asked users whether they were comfortable with the current location permissions for each of the apps installed on their phones (a construct we refer to as a perceived measure in the context of past behavior). We

Workshop on Usable Security (USEC) 2019
24 February 2019, San Diego, CA, USA
ISBN 1-891562-57-6
https://dx.doi.org/10.14722/usec.2019.23014
www.ndss-symposium.org

asked participants to share their location with our app and treated their sharing decision as our dependent variable.

For RQ1A, we leveraged several pre-validated perceived measures from the literature that have been shown to be relevant for predicting behavioral intent toward location-sharing: Behavior Intention [35], Perceived Surveillance [43], Perceived Intrusion [42], Secondary Use of Personal Information [34], FYI About Myself [28], and Power Usage [25]. These pre-validated constructs were collected via a questionnaire embedded in the study app. To investigate RQ1B, we scraped behavioral data in the background, including the app manifest of the apps installed on the phone, the 'Dangerous Permissions' [13] granted to these apps, and specifically, whether location permissions were granted to these apps. According to Google, Dan gerous Permissions are those that provide access to a user's personal information or stored data or control an app's access to the operation of other apps [13]. To answer RQ1C, we created our own measures by asking participants to reflect on their comfort level sharing their location with existing apps on their smartphones and whether they would revoke this permission for apps they were not comfortable having access to their location (even though location access had been previously granted to these apps). To analyze this data, we carried out binary logistic regression models in a step-by-step fashion to derive the strongest model to predict whether the participants granted the location permission to our study app (RQ2).

We found that most perceived measures (except FYI About Myself) were weak predictors of participant location sharing behavior on their own, and none of the behavioral data scraped from the phone was significantly linked to the decision to grant location access to our app. We also discovered that using a perceived in context of past behavior variable (specifically, the percentage of existing apps to which the participant was comfortable giving location access) could significantly improve the predictive power of our model. The model that used only perceived variables explained 17.2% of the variance, the model with only scraped behavioral data explained 2.4% of the variance, and the combined model with perceived measures and the perceived in context of past behavior variable explained 21.5% of the variance in our dependent variable. The difference in explanatory power between these models was statistically significant.

Improving the ability to predict people's location sharing behaviors can help researchers understand the factors behind users' privacy decisions. App developers can also anticipate user attitudes, and perhaps, ask for the location permission only when users are likely to grant it for a better user experience. Overall, we make the following contributions to the field of end-user mobile privacy; we: 1) identify the perceived measures which best predict users' actual location sharing behavior, 2) show that scraped behavioral data is not a good predictor of location sharing behavior, and 3) uncover that a hybrid perceived measure of users reflecting on their past behavior can significantly improve models for predicting location sharing behaviors.

II. Related Work

Approaches to studying privacy on mobile devices have generally been divided into 'social' and 'computational' approaches. Social science researchers often conduct survey studies with attitudinal measures to gauge user behavior, whereas computational researchers often rely on behavioral data scraped from the device to predict future user privacy decisions. In the sections that follow, we describe these two approaches and argue for the need of merging the two in order to advance interdisciplinary research on understanding and predicting users' privacy behaviors.

A. Mobile Privacy and Behavioral Intention

Survey based studies on mobile privacy have been conducted mainly in the fields of Information Systems (IS) and Human-Computer Interaction (HCI). A key theme of these studies has been to try to predict user behavior based on self-reported survey responses. Such questionnaires generally try to measure the user's attitude toward specific topics using specially designed constructs (e.g., for a review of various informational privacy measures, refer to Preibusch [30]). For instance, a commonly accepted practice in this research is that user behavior can be predicted based on users' attitudes and that Behavioral Intention (based on the Theory of Planned Behavior [1]) is the strongest predictor of actual behavior. The Theory of Planned Behavior states that behavioral beliefs inform user attitudes toward behaviors which then lead to behavioral intention which directly impacts user behavior [1]. Therefore, researchers have developed various constructs to quantify different aspects of user beliefs, feelings, intentions, and attitudes toward mobile privacy in order to predict behavioral intention as a dependent variable.

For example, Smith et al.'s [34] work was one of the first to develop scales to measure user concern for information privacy. The work introduced a fifteen-item instrument that measured concerns regarding data collection, unauthorized secondary use, improper access, and errors in data handling. Our research focuses on the context of mobile privacy. Therefore, we draw from the Mobile Users' Information Privacy Concern (MUIPC) framework [43] that identifies factors that influence mobile phone users' behavioral intention to use mobile apps and share their personal information with the apps. MUIPC is a three-factor scale that determines users' concern for information privacy in a mobile context by measuring their concern regarding misuse of shared data, degree of intrusion, and perceived surveillance. MUIPC has been utilized in several user behavior studies [17], [32].

In this research, we draw from a number of studies that specifically examined users' location sharing behaviors. For example, Guha et al. [14] studied the practice of deceptive location sharing (i.e., users deliberately sharing incorrect location information) and found that users engaged in this practice as acts of boundary and impression management due to various concerns about privacy. Page et al. [28] showed that location sharing decisions were influenced by a specific communication style called 'For Your Information' (FYI), where users would rather infer availability and social information about others than interact and ask them this information; those preferring an FYI Communication style were more likely to use location sharing social networks. Further work by Page et al. [29] found that the desire to preserve relationship boundaries was the main source of privacy concerns regarding location sharing social networks; when people felt that location sharing would

change their relationships with others, they expressed a wide range of social privacy concerns, such as worrying about feeling compelled to interact with others or being inundated with information from other people. Heavy users of location-sharing social media were found to be less concerned that location sharing could impact their relationship boundaries. Other location sharing studies reinforce the potential privacy concerns induced by sharing location. Barkhuus et al. [5] distinguished between location tracking services (i.e., services that send out a user's location) and position aware services (i.e., services that rely on the device's knowledge of its location). Even though users perceived both of these services as equally useful, Barkhuus et al. [5] found that location tracking services produced far greater concern for privacy.

Additionally, researchers have profiled power users as those who use technology to the fullest extent, can easily adapt to technological changes, and feel that technology is an integral part of their lives [25]. Power users were found to prefer customization (i.e., tailoring interfaces to learned user preferences) as long as the customization was self-initiated. If the customization was done automatically by the system, power users felt a loss of agency and did not feel as positively toward the presented content [36]. Therefore, we incorporated Power Usage as a construct in our model, which is described in more detail in our Research Framework.

B. Computational Approaches to Mobile Privacy

In contrast, computational approaches to mobile privacy have generally focused on designing tools to detect dangerous behavior, designing systems to increase user awareness about privacy threatening behavior, and using machine learning to identify malware threats to user data [2], [3], [20]. For instance, researchers have built a system that automatically maps each app permission to the part of the app user interface requesting it to make users more aware of what permissions are being requested [20]. Almuhimedi et al. [2] designed an Android permissions manager that periodically sends 'nudges' to remind users of the different kinds of data being collected and the frequency with which it is collected. The nudges resulted in 58% of the participants restricting some of their permissions. Fawaz et al.'s [11] LP Guardian tool worked intelligently to make sure that apps accessed location only when the user expected, anonymized location information for certain services in the background, and limited user profiling by different apps without significantly affecting the user experience. Such research raises users' awareness and helps them consider privacy risks of their decisions, but it does less in terms of helping researchers and designers understand users' privacy behaviors. More importantly, these efforts assume that using behavioral data from the past is well-aligned with users' actual privacy preferences, which may not necessarily be the case. As such, these approaches may potentially propagate past mistakes or nudge users to make future privacy decisions that they might regret [41].

C. An Interdisciplinary Approach to Mobile Privacy

There is a dearth of research that has merged the two aforementioned approaches by collecting both survey data and scraping behavioral data to try to understand and predict user privacy behaviors. Lin et al. [21] downloaded apps through the Google play store and performed static analysis of the app source code to identify specific sections within the code where permissions were used. Separately, they recruited participants to answer questions about the downloaded apps, using these responses to derive a set of privacy profiles based on correlating self-reported preferences with app permissions within the source code. Similarly, Ghosh et al. [12] used phone metadata, such as call frequency and call length, to predict user privacy concerns. They found that higher call response rate, higher missed call rate, and higher number of new contacts were associated with a low concern for privacy. Lin et al. [22] found that clearly revealing the purpose of requesting sensitive permissions made it more likely that users will feel positively about granting permissions to an app.

These studies demonstrate the potential benefit of adopting interdisciplinary approaches to understand and predict users' privacy behaviors and/or attitudes. Unlike Lin et al. [21], who created privacy profiles, or Ghosh et al. [12], who predicted privacy concerns, we predict user privacy *behavior* (specifically, whether participants chose to grant location access to our study app). While Lin et al. [22] asked participants general questions regarding apps, we tailored our questions to the specific apps installed on the participants' devices. For example, instead of asking participants whether they are comfortable sharing their location with mobile apps in general, we scraped the app manifest of all apps installed their mobile phone and asked specifically about location sharing comfort regarding each app to which they had already granted location access. We consider such contextually prompted responses as perceived measures in the context of past behavior. We believe this new type of variable represents a hybrid construct that combines the strengths of the social and computational sciences in a way that helps us better understand end-user privacy behaviors.

III. Research Framework

Our goal was to predict whether a participant will share their location with our app based on the following distinct classes of independent variables: 1) perceived measures, 2) scraped behavioral variables/data, and 3) perceived measures in context of past behavior. To ensure a baseline understanding of the key differences among these measures, we provide a brief introduction to each.

Perceived variables measure how a user feels about a certain topic, behavior, or action. Perceived measures are self-reported user data which is grouped into sets of related items called constructs. Constructs are often considered latent variables, i.e., variables that are inferred from other observed variables. A construct must be carefully designed to be statistically valid [10], i.e., it must pass various construct validity tests so that it is confirmed to be measuring what it claims to measure. For the purposes of this study, we used pre-validated constructs from prior research. An example of a pre-validated construct is the Behavioral Intention construct used by Xu et al. [44] to quantify the degree to which users plan on disclosing personal information and using mobile apps.

In contrast, behavioral data is content or meta data that is collected unobtrusively from the users' devices. This data is considered 'objective' in that it is based on scraped data

as opposed to the users' subjective self-reports. Studies by computer scientists often use objective variables since such data can be accessed via software [31]. An example of using scraped behavioral data to make predictions is Seneviratne et al.'s [31] work that predicted user traits, such as religion, gender, and relationship status, based on the kinds of apps installed apps on the phone.

In the subsections below, we describe our dependent variable and the three classes of predictor variables, including our new perceived measures in the context of past behavior.

A. Dependent Variable: Location Sharing

The dependent variable we chose as a measure of privacy behavior was whether a person agreed to provide their location to our study app. We framed the variable as a 'Yes/No' choice based on whether the location permission was granted. Location services are the cornerstone of personalized mobile content and play an important role in delivering targeted advertisements. Social networks, such as Foursquare, and mapping applications, such as Google Maps, rely on users sharing their location with the community or the app. Therefore, finding models to improve the prediction of whether users will share their location can go a long way in helping such applications design better user experiences. A number of studies have attempted to predict end-user privacy behaviors in mobile contexts. For instance, research has linked attitudes and behaviors regarding granting permissions with the clarity of describing the purpose of the requested access [22], [33]. If software designers can anticipate these privacy concerns, they can design better experiences that help people meet their privacy expectations, thus helping retain users.

B. Perceived Measures

For our perceived measures, we included relevant constructs from the mobile and location disclosure literature. The constructs are described in more detail in Appendix A. All measures were used in their original form, except Behavior Intention (to use a new app); we slightly adapted the items originally developed by Xu et al. [44] by shortening the intention-to-use time frame from 12 months down to 3 months. Since our users had already installed our study app, we deemed a time frame of 3 months to be more realistic for the context of our study.

Behavior Intention: We used this measure developed by Xu et al. [44] to measure the intention to disclose personal information and to use mobile apps in the next 3 months. The Theory of Planned Behavior states that behavior can be predicted using attitudes toward the behavior and behavioral intent [1]. Xu et al. [43] showed that increased user privacy concern reduced a user's behavioral intention to disclose personal information and to use mobile apps. Therefore, we wanted to measure the intention to use new apps and to share data with new apps as we expected these factors to correlate with location sharing decisions.

Perceived Surveillance: The perceived surveillance construct developed by Xu et al. [43] quantifies users' perception of being surveilled and having too much information collected about them. Perceived surveillance is one of the factors in MUIPC [43] and is rooted in the dimension of 'collection' from Malhotra et al.'s [24] Internet Users' Information Privacy Concerns (IUIPC) scale. Malhotra et al. [24] noted that data collection is the starting point of various privacy concerns. Therefore, the dimension of collection in IUIPC measures the degree to which a person is concerned about the specific data that others have relative to the value of benefits received. Since user privacy decisions are often based on an assessment of the perceived benefits and risks associated with the decision, we added the perceived surveillance measure to quantify the perception of the balance between surveillance and benefit [18].

Perceived Intrusion: Developed by Xu et al. [42], perceived instrusion quantifies the perception of intrusion caused by using mobile apps. Xu et al. [42] conducted a survey and found that perceived intrusion shaped individuals' views about the privacy practices of specific websites. Therefore, we used the perceived intrusion measure in order to correlate it with the privacy related decision of sharing location. Our intent was to examine if the perceived intrusion of mobile app use had an effect on location sharing behavior.

Secondary Use of Personal Information: The secondary use of personal information construct quantifies people's concerns about their information being used for purposes other those for which it was collected [34]. Solove et al. [35] noted that "the potential for secondary use generates fear and uncertainty over how one's information will be used in the future, creating a sense of powerlessness and vulnerability." Xu et al. [43] also make secondary use a factor of their MUIPC scale. Therefore, we included this measure in order to study the impact of the perception of secondary use on location sharing behavior.

FYI About Myself: The FYI About Myself construct quantifies the 'FYI Communication Style' identified by Page et al. [28]. People with the FYI Communication Style prefer to keep in touch with others without direct interaction and were shown to be more willing to share their location in location sharing social networks [28]. We included this measure as it has already been shown to impact location sharing decisions.

Power Usage: The Power Usage scale of Marathe et al. [25] measures the degree to which a user is a 'power user.' Power users are technologically adept and use their gadgets to the fullest potential. Kang et al. [15] found that power users are less likely to share personal information on personalized mobile sites but reveal more when interacting with non-personalized mobile content. Since our study app is highly personalized, we used this measure to analyze whether being a power user affects location sharing choices.

C. Behavioral Scraped Variables

We collected the following data from the devices of our study participants:

Number of installed apps: The number of apps installed on the device based on the device app manifest.

Total Dangerous Permissions granted: The number of Dangerous Permissions granted to the apps installed on the device. We included this measure as it indicates that someone more willing to grant Dangerous Permissions is perhaps less concerned about privacy. Access to an Android user's location (both fine and coarse) is considered a Dangerous Permission

along with others, such as calendar, call logs, camera, contacts, microphone, phone, sensors, SMS, storage, etc. [13].

Location ratio: We calculated location ratio as the number of apps with location permission divided by the total number of apps installed on the device. This variable serves to quantify a participant's past behavior about location sharing. We included this measure because it provides a snapshot of past location sharing behavior.

D. Perceived Measures in the Context of Past Behavior

We combined our scraped behavioral variables with the attitudinal variables to create new perceived measures contextualized to past behavior:

Location comfort (percentage): This variable was calculated as the percentage of apps installed on the device for which the participant had granted location permission (i.e., behavioral data) and expressed comfort with the granted location access (i.e., attitudinal data). This variable captures a person's feelings regarding their past privacy related behavior. The Theory of Planned Behavior [1] suggests that behavior is based on attitude toward that behavior. This variable captures the attitude toward past location sharing behavior.

Location revoke (percentage): We calculated this variable as the percentage of apps for which the participant reported a desire to revoke location access because of being uncomfortable sharing location (despite having granted location access to these apps). Thus, the location revoke percentage variable captures the intention to revoke location access from apps with which location sharing was found to be uncomfortable. The Theory of Planned Behavior states that behavioral intent is the basis for actual behavior [1].

IV. Methods

Our goal was to verify whether existing perceived measures (RQ1A) and behavioral data (RQ1B) are suitable for predicting *actual* user behavior regarding location sharing. We further wanted to examine if these prediction models could be strengthened by creating a hybrid of these two classes of variables (RQ1C). To collect data on the variables relevant to tackling our research questions, we implemented an app for smartphones running the Android operating system. The following subsections describe the steps involved in the study deployment and participant recruitment. All study procedures were reviewed and approved by the Institutional Review Board (IRB) of the University of Central Florida (UCF).

A. Study Design and App Flow

As mentioned earlier, we were interested in three types of measures: 1) perceived, 2) behavioral, and 3) perceived in context of past behavior. To collect all of these measures simultaneously and seamlessly, we implemented a smartphone app that incorporated an in-app questionnaire to collect the perceived measures. While the questionnaire was being answered, the app collected information on the permissions granted to each app to serve as the behavioral measures. In addition, we asked an 'in context' question about location sharing as our perceived in context of past behavior measure.

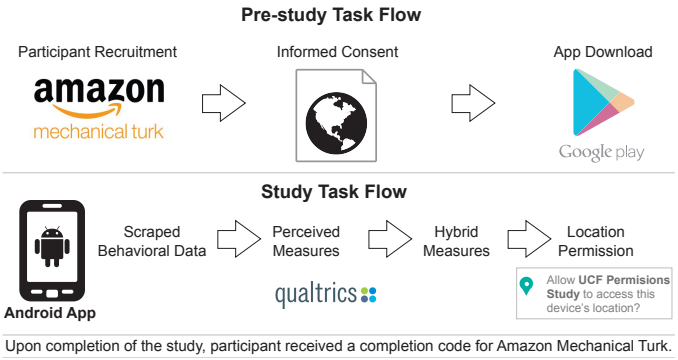

Fig. 1. The flow of the various steps involved in the study.

Figure 1 illustrates the various steps involved in the study. Android users from Amazon Mechanical Turk who were interested in participating in the study were directed to a webpage that introduced the study and sought informed consent for participation. To avoid priming, we did not use the term 'privacy' anywhere within the study description. After reading the study description, those who consented to participate in the study were provided with a randomly generated unique 'Consent ID' and directed to a link to install our study app from the Google Play app store. Upon installing and launching the app, each participant was first required to enter the Consent ID to verify completion of the informed consent procedures. All data collected by the app was transmitted to our database over a secure channel.

After verification of the Consent ID, the app presented the participant with a questionnaire to be completed within the app via the Qualtrics platform. The questionnaire included the pre-validated perceived measures described in our Research Framework. In addition, to help us spot inattentive participants, the questionnaire contained an attention check question: *Please select 'Disagree Somewhat' for this question.*

While the participant was answering the questionnaire, the app ran a background process to collect data on the apps installed on the device along with the permissions granted to each app. Note that the study description that sought informed consent explicitly disclosed the background data collection. Given the privacy-sensitive nature of this information, we minimized the extent of the collected permissions data by capturing information only for permissions classified by Google as Dangerous Permissions. For each app on a participant's device, we collected the list of all granted Dangerous Permissions. We further recorded whether each of the granted Dangerous Permissions was 1) present in the respective app's manifest file but not explicitly requested from the device user, 2) requested but denied by the device user, or 3) requested and granted.

Upon completing the questionnaire, participants were presented with a list of all apps on their device that had been granted access to the location permission. As shown in Figure 2, participants were asked to indicate whether they were comfortable sharing their location with each of the apps in the list. Once participants completed this step and chose to continue, our study app asked for access to the location permission using the standard permissions dialog of the Android operating

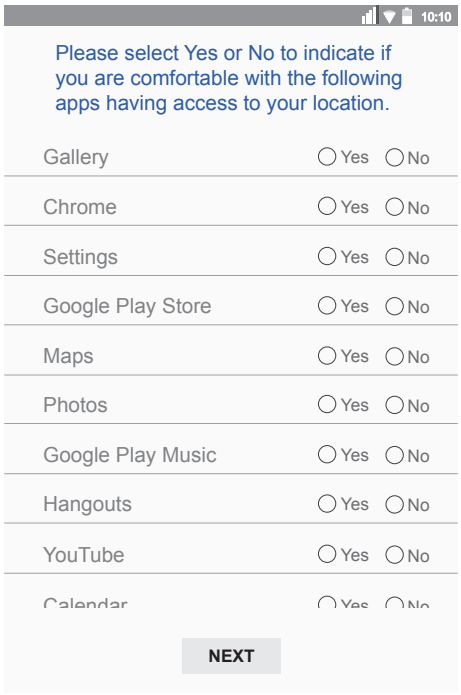

Fig. 2. Screenshot of the study app showing the screen that asked whether the participant was comfortable sharing location with the various apps on the phone. This list was dynamically generated based on the apps installed on the participant's phone.

system. Ostensibly, the study app's location request was made in order to enable us to collect participant location as one of the pieces of demographic information requested by the app at this point in the study. We used the participants' location sharing decisions to record our dependent variable 'Location given.' After recording the participants' choices regarding providing location access to the study app, we requested demographic information and concluded the study.

B. Data Analysis Approach

To prepare our data for analysis, we created composite variables for our constructs by averaging across multiple items and calculated Cronbach's alpha to assess construct validity. Cronbach alpha measures the internal consistency of an individual measure [10]. Cho and Kim [9] explain that a Cronbach's alpha value of 0.6–0.7 may provide adequate reliability, and all of our constructs were above this threshold. The descriptive statistics for all of our variables are provided in Table I. Next, we calculated the percentage of participants who reported that they would revoke location permission to an app for which they indicated in the previous step that they were uncomfortable with the app having their location. The screen we used for gathering this information was similar to the one in Figure 2. We also measured the opposite, i.e., apps to which participants chose to allow location access to an app that did not previously have such access. We did not differentiate between permissions for "coarse location" and "fine location" and included both when calculating location comfort percentages. The dialog requesting location sharing uses the same wording to request access without specifying which of the two types of location permissions is sought. Nonetheless, we calculated location comfort only for apps which had the "access fine location" permission and found that it made a negligible difference to our final model. In Appendix B, we include a correlation matrix of the Pearson's bi-variate correlations between all of our study variables.

After preparing the data, we conducted binary logistic regression analyses to answer each of our high-level research questions. We used binary logistic regression because our dependent variable is dichotomous [19] (i.e., grant/deny location permission). First, we examined separate models for each of the three classes of variables to identify which perceived measures (RQ1A), behavioral data (RQ1B), and perceived in context of past behavior measures (RQ1C) were significant in predicting our dependent variable. Finally, we performed a step wise logistic regression that included the statistically significant variables from each model to combine them into a single model (RQ2). We present the outcomes of these analyses in Section V.

C. Participant Recruitment and Sample Characteristics

We recruited participants by posting the study as a 'Human Intelligence Task' (HIT) on the Amazon Mechanical Turk crowd work platform. To avoid the impact of cultural variance, we limited participation to U.S. adults (18 years of age or older). For adequate response quality, we restricted our HIT to workers who had HIT approval rates greater than 95% with at least 50 approved HITs. Since our study app could run only on the Android operating system, all participants were required to be users of Android devices. Upon completing the study, the app provided each participant with a randomly generated unique completion code to be entered on Amazon Mechanical Turk as the proof of completion of the study task. All participants who demonstrated successful completion of the study task by entering a valid completion code were compensated $1.

Overall, 135 people accepted the HIT and completed the study task. After discarding the responses of those who failed the attention check, we had valid data from 114 participants. This included 63 (56%) males and 51 (44%) females. Most (82%; N=93) of our participants identified themselves as White Caucasian. Most (82%; N=21) of the participants lived in urban or suburban areas (40% urban; N=45 and 42% suburban; N=47) with the remaining 18% (N=25) coming from rural areas. Nearly 3/5^{ths} (58%; N=66) of the participants reported completing at least a 4-year college degree. Nearly 3/5^{ths} of the participants (58%; N=66) were employed full time covering a diversity of occupations from Software Engineering to Food Management and earning a median income in the $40,000–$60,000 range.

V. Results

The following subsections describe the app use practices reported by our participants followed by the results of the analyses we carried out to tackle the research questions we set forth in Section I.

A. Participants' Mobile App Use

Our participants had an average of 94 apps installed on their devices with a minimum of 29 and a maximum of 239, with a standard deviation of 43. The top five most

TABLE I. DESCRIPTIVE STATISTICS OF ALL MODEL VARIABLES AND INTERNAL CONSISTENCY FOR PERCEIVED MEASURES.

Variable	Variable Type	Mean	Median	SD	Skewness	Kurtosis	Cronbach's Alpha
Behavioral Intention	Perceived Measure	4.09	4.33	0.786	−1.230	1.680	0.81
Perceived Surveillance	Perceived Measure	4.16	4.33	0.711	−0.829	0.163	0.64
Perceived Intrusion	Perceived Measure	3.83	4.00	0.999	−0.968	0.574	0.94
Secondary Use of Personal Information	Perceived Measure	4.13	4.00	0.967	−1.360	1.770	0.93
FYI About Myself	Perceived Measure	1.39	1.33	1.022	0.328	−0.740	0.83
Power Usage	Perceived Measure	4.23	4.25	0.524	−0.690	0.790	0.82
Number of Installed Apps	Behavioral Scraped	93.91	82.00	43.200	1.143	1.170	N/A
Total Dangerous Permissions Granted	Behavioral Scraped	230.44	218.50	91.500	0.870	1.185	N/A
Location Ratio	Behavioral Scraped	26.51	25.52	8.590	0.294	−0.386	N/A
Location Comfort (percentage)	Hybrid Measure	33.50	25.00	27.900	0.784	−0.496	N/A
Location Revoke (percentage)	Hybrid Measure	26.88	13.27	30.700	1.073	−0.170	N/A

TABLE II. BINARY LOGISTIC REGRESSION: USING PERCEIVED MEASURES TO PREDICT LOCATION SHARING. NAGELKERKE R^2 = 17.3%

Variable	Odds Ratio	P-value
Behavioral Intention	1.369	0.316
Perceived Surveillance	0.423	0.137
Perceived Intrusion	1.150	0.737
Secondary Usage of Personal Information	1.013	0.977
FYI About Myself	1.759	0.034*
Power Usage	0.977	0.962

* p < 0.05; ** p < 0.01; *** p < 0.001

TABLE III. BINARY LOGISTIC REGRESSION: USING SCRAPED BEHAVIORAL MEASURES TO PREDICT LOCATION SHARING. NAGELKERKE R^2 = 2.4%

Variable	Odds Ratio	P-value
Number of Installed Apps	1.005	0.606
Total Dangerous Permissions Granted	1.000	0.972
Location Ratio	1.020	0.620

TABLE IV. BINARY LOGISTIC REGRESSION: USING PERCEIVED IN CONTEXT OF PAST BEHAVIOR MEASURES TO PREDICT LOCATION SHARING. NAGELKERKE R^2 = 16.0%

Variable	Odds Ratio	P-value
Location Comfort (percentage)	1.024	0.036*
Revoke Location (percentage)	0.988	0.094

* p < 0.05; ** p < 0.01; *** p < 0.001

TABLE V. STEP-WISE BINARY LOGISTIC REGRESSION FOR COMPARING MODELS

Variable	Odds Ratio	P-value	Nagelkerke R^2
STEP 1			
FYI About Myself	2.029	0.004**	11.7%
STEP 2			
FYI About Myself	1.945	0.008**	21.5%
Location Comfort Percentage	1.028	0.009**	

χ^2(Step1, Step2)=8.632, Degrees of Freedom=1, p=0.003*

* p < 0.05; ** p < 0.01; *** p < 0.001

common non-system apps (i.e., apps which are not typically pre-installed on devices) were social media apps: 1) Instagram (43%), 2) Facebook (41%), 3) Snapchat (24%), 4) Hangouts (21%), and 5) Twitter (21%). On average, participants had granted 230 Dangerous Permissions to the various apps on their devices with an average of 25 of these apps having access to their location (fine or coarse). Yet, on average, participants were comfortable with only 33% of the location utilizing apps actually having location access. However, when asked if they would revoke location access for apps with which they were uncomfortable sharing location, only 27% (N=30) of the participants wished to do so. When asked for location access by our study app, 73% (N=84) of participants granted the permission to access their location.

B. Binary Logistic Regression Analyses

RQ1A, Perceived Measures: Our first research question explored if perceived measures can predict actual user location sharing behavior. The results of our logistic regression are shown in Table II. FYI About Myself (p=0.034, e^{β}=1.759) was the only significant predictor. While FYI About Myself performed better than the rest of the perceived measures, the overall model with all of the variables explained 17.3% of the variance in location sharing, compared to 11.7% for FYI by itself, as shown in Table V.

RQ1B, Behavioral Measures: Next, we carried out a binary logistic regression using the number of installed apps, total Dangerous Permissions granted, and location ratio as the independent variables. The results are shown in Table III. Overall, we found these variables to be poor predictors of location sharing with the study app. Location ratio (p=0.620, e^{β}=1.020), number of installed apps (p=0.606, e^{β}=1.005), and total number of Dangerous Permissions granted (p=0.972, e^{β}=1.000) were all found to have no significant effect on the location sharing decision. The overall Nagelkerke R^2 value of this model was very low at 2.4%.

RQ1C, Perceived in Context of Past Behavior Measures: Perceived in context of past behavior variables did nearly as well at predicting the participants' location sharing decision as the perceived constructs. Our binary logistic regression model (Table IV) shows that location comfort percentage was significant (p=0.036, e^{β}=1.024), but revoke location percentage (p=0.094, e^{β}=0.988) was not. Overall, this model explained 16.0% of the variance in location sharing.

RQ2, Combined Model: RQ2 investigated if past user behavior or perceived in context of past behavior variables augment purely perceived measures in predicting user behavior. In order to answer this research question, we combined all significant variables from the above mentioned individual regression models (so as not to inflate our R^2) to achieve the best model for predicting our dependent variable. We found that adding the perceived in context of past behavior variable for location comfort percentage to FYI About Myself helped explain significantly more variance in location sharing behavior than FYI About Myself alone. For a one point increase in location comfort percentage, the probability of sharing location increased by a factor of 1.028. For a one point increase in FYI

About Myself, the probability of sharing location increased by a factor of 1.945. The model which used only FYI About Myself had an R^2 of 11.7%, whereas adding location comfort percentage improved it to 21.5%, which was a statistically significant (χ^2=8.632, p=0.003) change in predictive power. We summarize these models and the R^2 change in Table V.

VI. Discussion

For RQ1A, we found that perceived measures could predict actual privacy behavior (i.e., location sharing) to some extent. Our model showed that FYI About Myself was the most important factor in predicting whether the participants shared their location with our study app. For RQ1B, we discovered that scraped behavioral data proved to be worse at predicting privacy behavior than perceived measures. No significant predictors were found in our model that included only behavioral data and R^2 value was very low. For RQ1C, perceived in context of past behavior variables also had some predictive power for explaining location sharing behavior. Location comfort percentage was found to be a significant predictor variable for the decision to share location. Finally, the combined model of perceived measures and perceived in context of past behavior measures was the best predictor model for location sharing. FYI About Myself and location comfort percentage were found to be the best predictors in this model (RQ2). This suggests that perceived in context of past behavior variables might be able to augment traditional surveys and scraped behavioral data to produce stronger predictive models of user behavior. We discuss the implications of these findings in more detail below.

A. Implications for Mobile Privacy Location Sharing Research

In our review of the literature, most of the mobile privacy research that we encountered was strictly divided between the social sciences or computational sciences, with few studies at the intersection of the two disciplines. We found that the best model for predicting smartphone users' app location sharing behavior was a hybrid of the two. In this subsection, we reflect on the implications of our results for these different privacy research communities individually and suggest a path forward that leverages the strengths of both approaches.

1) Social Science Privacy Research: Similar to Page et al. [28], we found that the FYI About Myself communication style significantly influenced participants' location sharing behavior. FYI Communication Style pertains to how one communicates location information with others. Since FYI About Myself was also a significant predictor of location sharing behavior, it suggests that one's preferred communication style is an important factor to consider when predicting whether a user will grant an app the permission to access location. Designers should consider personal preferences for whether a given user values the convenience of letting others (apps in this case) decide when location access is needed versus wanting a more hands-on explicit approach to disclosing location on a case-by-case basis. Further, we found that the FYI About Myself personal trait was a stronger predictor of location sharing behavior than any of the perceived constructs in Xu et al.'s earlier work, which was based on MUIPC [43] and the Theory of Planned Behavior [1]. This suggests that personal traits play a bigger role than intention when deciding whether to grant location access.

Interestingly, Behavioral Intention [44], which is cited in the social sciences as the strongest predictor of actual behavior [1], was not a significant factor in any of our models. This may have been because the behavioral intention to use mobile apps and share location with them may have been a foregone conclusion for our participants. All of our participants were willing to install our study app and most (73%; N=83) granted our app access to their location. Overall, Behavioral Intention (M=4.09, SD=0.79) and Power Usage were high among our participants (M=4.23, SD=0.52); therefore, the lack of significance in some of these variables may also have been due to the proclivity of our participants to use mobile apps and share information with the apps.

Regardless of the reason for Behavioral Intention not correlating with participants' actual location sharing behavior, our results call into question the common practice of using Behavioral Intention as a proxy measure for actual privacy behavior [25], [28], [36], [43]. While Behavioral Intention may be useful for predicting some behaviors, such as technology adoption, privacy researchers should consider using more contextualized measures for the specific type of information disclosure being studied, similar to the FYI About Myself personal trait that was created based on empirical work specifically on mobile location privacy sharing [28]. Further, researchers should also consider correlating perceived measures with a proxy for actual behavior. Because privacy behaviors are often paradoxical, and possibly more nuanced, than other technology-related behaviors, such as technology adoption [37], [38], it is possible that the Behavioral Intention construct [43] and the Theory of Planned Behavior [1] may not be the best proxies or approaches for predicting actual privacy behaviors. Similarly, Power Usage was not a significant predictor of participants' location sharing behavior. Upon further reflection, this lack of an effect may also be because the Power Usage construct focuses on using technology in general, whereas FYI About Myself is more directly connected to the the concept of location sharing. Overall, these results indicate that location-specific perceived constructs should be used to predict users' future location sharing behavior, rather than more general measures about privacy or mobile apps.

2) Computational Privacy Research: In contrast, the finding that scraped behavioral data was not a good predictor of location disclosure at all suggests that a user's future location disclosure behavior is not tied to past behavior. The number of installed apps and the number of Dangerous Permissions granted may be too broad and thus not tied to attitudes about location information. However, we were surprised that past location sharing behavior (i.e., location ratio) was also not a significant predictor of future location sharing behavior. While Ghosh et al. [12] found that device metadata, such as call duration and ignored calls, could predict self-reported user privacy concerns, our findings suggest that phone metadata might not be the best predictor of actual privacy behavior, such as mobile location sharing. In fact, past behavior may not even reflect desired behavior. Granting permissions may be performed either as a condition of using the app or may have been done without much thought or understanding of the permissions that are granted. As a result, currently granted

permissions do not predict the desired behavior when the user is explicitly asked to grant location permissions. Therefore, we offer a word of caution to computational researchers and system designers who create predictive models, design intelligent user defaults, and recommend privacy choices against using past privacy behavior as a proxy for determining users' privacy preferences.

3) The Importance of Self-Reflection on Past Behavior: Our key novel findings were that people's reflection on their past behavior (i.e., their comfort with their past location sharing decisions) was a predictor of their future behavior (i.e., whether they granted location access to our study app). Incorporating such reflection greatly improved on using just future intentions (i.e., perceived measures) or past actions (i.e., scraped behavioral data). Of the perceived measures where participants reflected on their past location sharing behavior, location comfort was the most influential in improving the prediction of location sharing behavior, more so than the revoke location measure, even though both were statistically significantly correlated with our dependent variable (see Appendix B). This suggests that users' comfort level with their past location permissions translates directly to their future decisions about sharing their location. Yet, even when participants were uncomfortable sharing their location with some of the apps installed on their smartphone, very few (27%; N=30) said they would revoke the location permission after knowing these apps did indeed have access to their location. Thus, future research needs to look beyond raising users' awareness of their privacy behaviors and try to understand factors that could effectively motivate behavior change to help users feel comfortable about their mobile privacy settings.

Overall, the results of our study suggest that combining social science approaches with computer science approaches can yield stronger predictive models. Our work encourages future privacy research to identify and measure relevant perceived in context of past behavior variables for the privacy related phenomenon or behavior being studied. In our case, we identified location comfort percentage as the perceived in context of past behavior variable strongly correlated with location sharing decisions. This suggests that how users *feel* about their past behavior is a better predictor of future behavior than the actual past behavior itself.

B. Implications for Design

Our work shows that user perceptions of their past behavior (in our case, location sharing decisions) greatly improves on using just attitudes about future behavior to predict future disclosure. We suggest that a combination of attitudinal measures asked in context of one's actual past behavior is more useful for understanding attitudes that lead to action. In fact, this technique mirrors a 'reflective learning' [7] approach, which has been shown to produce positive learning outcomes [23]. Namely, by reflecting on one's past choices, one can become more aware of one's actions and make better future choices, which might be incongruous with one's past choices. This capability to 'learn reflectively' could be supported in the design of apps as a context aware feature that could periodically remind users of their past decisions and give them the opportunity to reflect on and change their decisions based on a new context or bad experiences with previous decisions. Such a design could support a more dynamic conception of privacy that matches what users want, as opposed to what they think they want or what they did in the past. Prior research has shown that achieving the right 'privacy fit' can lead to higher user engagement with the service and help users feel more socially connected with others [39].

C. Limitations and Future Research

Limiting our participants only to the U.S. limits our ability to generalize our results to other populations since privacy decisions and experiences can be shaped by the cultural environment. Further, some research has shown that Amazon Mechanical Turk workers have unique privacy profiles compared to average users [16]. Moreover, users recruited from Amazon Mechanical Turk are likely to be more technically savvy than the general population. As a result, the use of a participant sample recruited from Amazon Mechanical Turk may also constrain the generalizability of our results. Since we limited participation to adults of ages 18 and above, the applicability of our results to younger populations needs to be verified. It may also be useful to verify whether our results generalize to those who use devices with operating systems other than Android, such as Apple's iOS. Future research should consider replicating our study with samples drawn from other populations that are more diverse in terms of ages, cultures, and technical abilities and perhaps try to identify more culture-specific perceived in context of past behavior variables. An interesting future direction could be to track changes in location sharing comfort for specific apps over the duration of time they are installed on the phone. Such tracking could help app makers identify location sharing behavior trends and take corrective action.

VII. Conclusion

The rapid growth of portable communications devices has meant that the boundaries of privacy are being tested in new ways. It is therefore crucial to understand and be able to predict user behavior in order to design experiences which respect the user's expectations of privacy. We contribute to the field of mobile privacy by shedding light on the kinds of perceived measures that can explain user location sharing behavior. We show that scraped behavioral data might not be the best indicator of future user behavior. However, augmenting perceived measures with perceived in context of past behavior variables can help strengthen prediction models.

Acknowledgment

This research was partially supported by the U.S. National Science Foundation (NSF) under grant CNS-1814439. Any opinion, findings, recommendations, and conclusions expressed in this material are solely those of the authors and do not necessarily reflect the views of the U.S. National Science Foundation. This work was also partially supported by a grant from the Bentley Data Innovation Network. We would like to thank Heather Lipford and Bart Knijnenburg for their input on our study design.

REFERENCES

[1] I. Ajzen, "The theory of planned behavior," *Organizational Behavior and Human Decision Processes*, vol. 50, no. 2, pp. 179–211, 1991. [Online]. Available: http://www.sciencedirect.com/science/article/pii/074959789190020T

[2] H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal, "Your location has been shared 5,398 times!: A field study on mobile app privacy nudging," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '15. New York, NY, USA: ACM, 2015, pp. 787–796. [Online]. Available: http://doi.acm.org/10.1145/2702123.2702210

[3] B. Amos, H. Turner, and J. White, "Applying machine learning classifiers to dynamic Android malware detection at scale," in *9th International Wireless Communications and Mobile Computing Conference*, ser. IWCMC 2013, July 2013, pp. 1666–1671.

[4] App Annie, "App Annie's app monetization report: Publishers to earn $189 billion from stores and ads in 2020," November 2016. [Online]. Available: http://go.appannie.com/report-app-annie-app-monetization-2016

[5] L. Barkhuus and A. Dey, "Location-based services for mobile telephony: A study of users' privacy concerns," in *Proceedings of Interact 2003*, 2003, pp. 709–712.

[6] S. Barth and M. D. de Jong, "The privacy paradox – investigating discrepancies between expressed privacy concerns and actual online behavior – a systematic literature review," *Telematics and Informatics*, vol. 34, no. 7, pp. 1038–1058, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0736585317302022

[7] E. M. Boyd and A. W. Fales, "Reflective learning: Key to learning from experience," *Journal of Humanistic Psychology*, vol. 23, no. 2, pp. 99–117, 1983. [Online]. Available: https://doi.org/10.1177/0022167883232011

[8] D. Chaffey, "Global social media research summary 2018," *Smart Insights*, November 2018. [Online]. Available: https://www.smartinsights.com/social-media-marketing/social-media-strategy/new-global-social-media-research/

[9] E. Cho and S. Kim, "Cronbach's coefficient alpha: Well known but poorly understood," *Organizational Research Methods*, vol. 18, no. 2, pp. 207–230, 2015. [Online]. Available: https://doi.org/10.1177/1094428114555994

[10] L. J. Cronbach and P. E. Meehl, "Construct validity in psychological tests," *Psychological Bulletin*, vol. 52, no. 4, pp. 281–302, 1955.

[11] K. Fawaz and K. G. Shin, "Location privacy protection for smartphone users," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '14. New York, NY, USA: ACM, 2014, pp. 239–250. [Online]. Available: http://doi.acm.org/10.1145/2660267.2660270

[12] I. Ghosh and V. K. Singh, "Predicting privacy attitudes using phone metadata," in *Social, Cultural, and Behavioral Modeling*, ser. Lecture Notes in Computer Science, K. S. Xu, D. Reitter, D. Lee, and N. Osgood, Eds., vol. 9708. Cham: Springer International Publishing, 2016, pp. 51–60.

[13] Google Inc. (2019) Permissions overview. [Online]. Available: https://developer.android.com/guide/topics/permissions/overview

[14] S. Guha and S. B. Wicker, "Spatial subterfuge: An experience sampling study to predict deceptive location disclosures," in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '15. New York, NY, USA: ACM, 2015, pp. 1131–1135. [Online]. Available: http://doi.acm.org/10.1145/2750858.2804281

[15] H. Kang and W. Shin, "Do smartphone power users protect mobile privacy better than nonpower users? Exploring power usage as a factor in mobile privacy protection and disclosure," *Cyberpsychology, Behavior, and Social Networking*, vol. 19, no. 3, pp. 179–185, 2016, pMID: 26783875. [Online]. Available: https://doi.org/10.1089/cyber.2015.0340

[16] R. Kang, S. Brown, L. Dabbish, and S. Kiesler, "Privacy attitudes of Mechanical Turk workers and the U.S. public," in *Proceedings of the 10th Symposium On Usable Privacy and Security*, ser. SOUPS 2014. Menlo Park, CA: USENIX Association, 2014, pp. 37–49. [Online]. Available: https://www.usenix.org/conference/soups2014/proceedings/presentation/kang

[17] M. J. Keith, J. S. Babb, and P. B. Lowry, "A longitudinal study of information privacy on mobile devices," in *47th Hawaii International Conference on System Sciences*, ser. HICSS 2014, Jan 2014, pp. 3149–3158.

[18] C. D. Kennedy-Lightsey, M. M. Martin, M. Thompson, K. L. Himes, and B. Z. Clingerman, "Communication privacy management theory: Exploring coordination and ownership between friends," *Communication Quarterly*, vol. 60, no. 5, pp. 665–680, 2012. [Online]. Available: https://doi.org/10.1080/01463373.2012.725004

[19] D. G. Kleinbaum and M. Klein, *Logistic regression: A self-learning text*, 3rd ed. Springer Science & Business Media, 2010.

[20] Y. Li, Y. Guo, and X. Chen, "PERUIM: Understanding mobile application privacy with permission-UI mapping," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '16. New York, NY, USA: ACM, 2016, pp. 682–693. [Online]. Available: http://doi.acm.org/10.1145/2971648.2971693

[21] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, "Expectation and purpose: Understanding users' mental models of mobile app privacy through crowdsourcing," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, ser. UbiComp '12. New York, NY, USA: ACM, 2012, pp. 501–510. [Online]. Available: http://doi.acm.org/10.1145/2370216.2370290

[22] J. Lin, B. Liu, N. Sadeh, and J. I. Hong, "Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings," in *Proceedings of the 10th Symposium On Usable Privacy and Security*, ser. SOUPS 2014. Menlo Park, CA: USENIX Association, 2014, pp. 199–212. [Online]. Available: https://www.usenix.org/conference/soups2014/proceedings/presentation/lin

[23] R. Loo and K. Thorpe, "Using reflective learning journals to improve individual and team performance," *Team Performance Management: An International Journal*, vol. 8, no. 5/6, pp. 134–139, 2002. [Online]. Available: https://doi.org/10.1108/13527590210442258

[24] N. K. Malhotra, S. S. Kim, and J. Agarwal, "Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model," *Information Systems Research*, vol. 15, no. 4, pp. 336–355, 2004. [Online]. Available: https://pubsonline.informs.org/doi/abs/10.1287/isre.1040.0032

[25] S. Marathe, S. Sundar, M. Nije Bijvank, H. van Vugt, and J. Veldhuis, "Who are these power users anyway? Building a psychological profile," 2007. [Online]. Available: https://research.vu.nl/en/publications/b349ddc2-97fe-4a2a-85e9-e6e58b0f0478

[26] P. A. Norberg, D. R. Horne, and D. A. Horne, "The privacy paradox: Personal information disclosure intentions versus behaviors," *Journal of Consumer Affairs*, vol. 41, no. 1, pp. 100–126, 2007. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1745-6606.2006.00070.x

[27] K. Olmstead and M. Atkinson, "Apps permissions in the Google Play store," in *Pew Research Center*, October 2015. [Online]. Available: PewResearchCenter

[28] X. Page, B. P. Knijnenburg, and A. Kobsa, "FYI: Communication style preferences underlie differences in location-sharing adoption and usage," in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '13. New York, NY, USA: ACM, 2013, pp. 153–162. [Online]. Available: http://doi.acm.org/10.1145/2493432.2493487

[29] X. Page, A. Kobsa, and B. P. Knijnenburg, "Don't disturb my circles! Boundary preservation is at the center of location-sharing concerns." in *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, ser. ICWSM 2012, 2012, pp. 266–273.

[30] S. Preibusch, "Guide to measuring privacy concern: Review of survey and observational instruments," *International Journal of Human-Computer Studies*, vol. 71, no. 12, pp. 1133–1143, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1071581913001183

[31] S. Seneviratne, A. Seneviratne, P. Mohapatra, and A. Mahanti, "Predicting user traits from a snapshot of apps installed on a smartphone," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 18,

no. 2, pp. 1–8, Jun. 2014. [Online]. Available: http://doi.acm.org/10.1145/2636242.2636244

[32] S. Sharma and R. E. Crossler, "Disclosing too much? Situational factors affecting information disclosure in social commerce environment," *Electronic Commerce Research and Applications*, vol. 13, no. 5, pp. 305–319, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1567422314000350

[33] F. Shih, I. Liccardi, and D. Weitzner, "Privacy tipping points in smartphones privacy preferences," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ser. CHI '15. New York, NY, USA: ACM, 2015, pp. 807–816. [Online]. Available: http://doi.acm.org/10.1145/2702123.2702404

[34] H. J. Smith, S. J. Milberg, and S. J. Burke, "Information privacy: Measuring individuals' concerns about organizational practices," *MIS Quarterly*, vol. 20, no. 2, pp. 167–196, 1996. [Online]. Available: http://www.jstor.org/stable/249477

[35] D. J. Solove, "A taxonomy of privacy," in *University of Pennsylvania Law Review*, vol. 154, no. 3, January 2006, pp. 477–560, GWU Law School Public Law Research Paper No. 129. [Online]. Available: https://ssrn.com/abstract=667622

[36] S. S. Sundar and S. S. Marathe, "Personalization versus customization: The importance of agency, privacy, and power usage," *Human Communication Research*, vol. 36, no. 3, pp. 298–322, 2010. [Online]. Available: http://dx.doi.org/10.1111/j.1468-2958.2010.01377.x

[37] V. Venkatesh and H. Bala, "Technology acceptance model 3 and a research agenda on interventions," *Decision Sciences*, vol. 39, no. 2, pp. 273–315, 2008. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-5915.2008.00192.x

[38] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," *MIS Quarterly*, vol. 27, no. 3, pp. 425–478, 2003. [Online]. Available: http://www.jstor.org/stable/30036540

[39] P. Wisniewski, A. K. M. N. Isam, B. P. Knijnenburg, and S. Patil, "Give social network users the privacy they want," in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ser. CSCW '15. New York, NY, USA: ACM, 2015, pp. 1427–1441. [Online]. Available: http://doi.acm.org/10.1145/2675133.2675256

[40] P. Wisniewski, A. K. M. N. Islam, H. R. Lipford, and D. C. Wilson, "Framing and measuring multi-dimensional interpersonal privacy preferences of social networking site users," in *Communications of the Association for Information Systems*, vol. 38, no. 10, 2016, pp. 235–258.

[41] P. J. Wisniewski, B. P. Knijnenburg, and H. R. Lipford, "Making privacy personal: Profiling social network users to inform privacy education and nudging," *International Journal of Human-Computer Studies*, vol. 98, pp. 95–108, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1071581916301185

[42] H. Xu, T. Dinev, H. J. Smith, and P. Hart, "Examining the formation of individual's privacy concerns: Toward an integrative view," in *Proceedings of the Twenty-Ninth International Conference on Information Systems*, ser. ICIS 2008, no. 6, 2008, pp. 1–16. [Online]. Available: https://aisel.aisnet.org/icis2008/6

[43] H. Xu, S. Gupta, M. B. Rosson, and J. M. Carroll, "Measuring mobile users' concerns for information privacy," in *Proceedings of the Thirty Third International Conference on Information Systems*, ser. ICIS 2012, 2012.

[44] H. Xu and H.-H. Teo, "Alleviating consumers' privacy concerns in location-based services: A psychological control perspective," in *Proceedings of the Twenty-Fifth International Conference on Information Systems*, ser. ICIS 2004, no. 64, 2004, pp. 793–806. [Online]. Available: https://aisel.aisnet.org/icis2004/64

APPENDIX A
PERCEIVED MEASURES

A. Behavioral Intention

Taken from Xu et al. [44], these items were measured on a 5-point Likert scale from 1 - Agree Strongly. 2 - Agree Somewhat. 3 - Neutral. 4 - Disagree Somewhat. 5 - Disagree Strongly.

Items:

1) I am likely to disclose my personal information to use mobile apps in the next 3 months.
2) I predict I will use new mobile apps in the next 3 months.
3) I intend to use mobile apps in the next 3 months.

B. Perceived Surveillance

Taken from Xu et al. [43], these items were measured on a 5-point Likert scale from 1 - Agree Strongly. 2 - Agree Somewhat. 3 - Neutral. 4 - Disagree Somewhat. 5 - Disagree Strongly.

Items:

1) I believe that the location of my mobile device is monitored at least part of the time.
2) I am concerned that mobile apps are collecting too much information about me.
3) I am concerned that mobile apps may monitor my activities on my mobile device.

C. Perceived Intrusion

Taken from Xu et al. [42], these items were measured on a 5-point Likert scale from 1 - Agree Strongly. 2 - Agree Somewhat. 3 - Neutral. 4 - Disagree Somewhat. 5 - Disagree Strongly.

Items:

1) I feel that as a result of my using mobile apps, others know about me more than I am comfortable with.
2) I believe that as a result of my using mobile apps, information about me that I consider private is now more readily available to others than I would want.
3) I feel that as a result of using mobile apps, information about me is out there that, if used, will invade my privacy.

D. Secondary Use of Personal Information

Taken from Smith et al. [34], these items were measured on a 5-point Likert scale from 1 - Agree Strongly. 2 - Agree Somewhat. 3 - Neutral. 4 - Disagree Somewhat. 5 - Disagree Strongly.

Items:

1) I am concerned that mobile apps may use my personal information for other purposes without notifying me or getting my authorization.
2) When I give personal information to use mobile apps, I am concerned that apps may use my information for other purposes.
3) I am concerned that mobile apps may share my personal information with other entities without getting my authorization.

E. FYI About Myself

Taken from Page et al. [28], these items were measured on a 5-point Likert scale from 1 - Agree Strongly. 2 - Agree Somewhat. 3 - Neutral. 4 - Disagree Somewhat. 5 - Disagree Strongly.

Items:

1) I want others to know what I am up to without my having to bother to tell them.
2) Others should be able to find out about me when they feel they need to.
3) I would prefer to share about myself with everyone in case anyone wants to know.

F. Power Usage

Taken from Marathe et al. [25], these items were measured on a 5-point Likert scale from 1 - Agree Strongly. 2 - Agree Somewhat. 3 - Neutral. 4 - Disagree Somewhat. 5 - Disagree Strongly.

Items:

1) I think most technological gadgets are complicated to use.
2) I make good use of most of the features available in any technological device.
3) I have to have the latest available upgrades of technological devices that I use.
4) Use of information technology has almost replaced my use of paper.
5) I love exploring all the features that any technological gadget has to offer.
6) I often find myself using many technological devices simultaneously.
7) I prefer to ask friends how to use any new technological gadget instead of trying to figure it out myself.
8) Using any technological device comes easy to me.
9) I feel like information technology is a part of my daily life.
10) Using information technology gives me greater control over my work environment.
11) Using information technology makes it easier to do my work.
12) I would feel lost without information technology.

APPENDIX B
CORRELATIONS MATRIX OF ALL MODEL VARIABLES

N = 114 in all cases	Location Sharing (DV)	Behavioral Intention	Perceived Surveillance	Perceived Intrusion	Secondary Use of Personal Information	FYI About Myself	Power Usage	Number of Installed Apps	Total Dangerous Permissions Granted	Location Ratio	Location Comfort Percentage	Revoke Location Percentage
Location Sharing (DV)	1.000	0.148	-0.239^{*}	-0.164	-0.0185^{*}	0.279^{**}	0.059	0.100	0.115	0.074	0.281^{**}	-0.275^{**}
Behavioral Intention		1.000	-0.108	-0.078	-0.094	0.169	0.467^{**}	0.135	0.034	0.016	0.328^{**}	-0.144
Perceived Surveillience			1.000	0.737^{**}	0.686^{**}	-0.344^{**}	-0.072	-0.135	-0.101	0.023	-0.321^{**}	0.353^{**}
Perceived Intrusion				1.000	0.735^{**}	-0.317^{**}	-0.035	-0.135	-0.042	0.042	-0.357^{**}	0.282^{**}
Secondary Use of Personal Information					1.000	-0.353^{**}	-0.003	-0.188^{*}	-0.142	0.031	-0.280^{**}	0.273^{**}
FYI About Myself						1.000	0.053	0.051	0.077	0.060	0.138	-0.276^{**}
Power Usage							1.000	0.219^{*}	0.255^{**}	0.086	0.178	-0.090
Number of Installed Apps								1.000	0.726^{**}	-0.033	0.373^{**}	-0.300^{**}
Total Dangerous Permissions Granted									1.000	0.505^{**}	0.190^{*}	-0.284^{**}
Location Ratio										1.000	-0.030	-0.117
Location Comfort (percentage)											1.000	-0.409^{**}
Revoke Location (percentage)												1.000

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$