

# Examining Power Use and the Privacy Paradox between Intention vs. Actual Use of Mobile Applications

Moses Namara  
Clemson University  
Clemson, SC, USA  
mosesn@clemson.edu

Reza Ghaiumy Anaraky  
Clemson University  
Clemson, SC, USA  
rghaium@clemson.edu

Pamela Wisniewski  
University of Central Florida  
Orlando, FL, USA  
pamwis@ucf.edu

Xinru Page  
Brigham Young University  
Provo, UT, USA  
xinru@cs.byu.edu

Bart P. Knijnenburg  
Clemson University  
Clemson, SC, USA  
bartk@clemson.edu

## ABSTRACT

The prevalence of smartphones in our society warrants more research on understanding the characteristics of users and their information privacy behaviors when using mobile apps. This paper investigates the antecedents and consequences of “power use” (i.e., the competence and desire to use technology to its fullest) in the context of informational privacy. In a study with 380 Android users, we examined how gender and users’ education level influence power use, how power use affects users’ intention to install apps and share information with them versus their actual privacy behaviors (i.e., based on the number of apps installed and the total number of “dangerous permission” requests granted to those apps). Our findings revealed an inconsistency in the effect of power use on users’ information privacy behaviors: While the intention to install apps and to share information with them *increased* with power use, the actual number of installed apps and dangerous permissions ultimately granted *decreased* with power use. In other words, although the self-reported intentions suggested the opposite, people who scored higher on the power use scale seemed to be more prudent about their informational privacy than people who scored lower on the power use scale. We discuss the implications of this inconsistency and make recommendations for reconciling smartphone users’ informational privacy intentions and behaviors.

## CCS CONCEPTS

• Security and privacy → Privacy protections; • Human-centered computing → Smartphones; Empirical studies in HCI; • Social and professional topics → Gender.

## KEYWORDS

Privacy; Smartphones; Mobile Apps; Power Use; Individual Differences

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*EuroUSEC '21, October 11–12, 2021, Karlsruhe, Germany*

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8423-0/21/10...\$15.00

<https://doi.org/10.1145/3481357.3481513>

## ACM Reference Format:

Moses Namara, Reza Ghaiumy Anaraky, Pamela Wisniewski, Xinru Page, and Bart P. Knijnenburg. 2021. Examining Power Use and the Privacy Paradox between Intention vs. Actual Use of Mobile Applications. In *European Symposium on Usable Security 2021 (EuroUSEC '21), October 11–12, 2021, Karlsruhe, Germany*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3481357.3481513>

## 1 INTRODUCTION

The ubiquity of smartphone usage has caused a marked increase in the number of people affected and the amount of data collected, analyzed, and categorized by mobile applications (“apps”) [25]. Users are concerned about this development, which causes them to engage in a range of information privacy management tactics, such as deciding to grant certain data access permissions to specific apps only, or to not use—or even not to install—certain apps on one’s smartphone [6, 53]. Prior research has shown that the decision to use a given smartphone app depends on a myriad of factors, including the user’s interests, past negative privacy experiences, expertise using a smartphone (“power use”), self-efficacy, trust in the application developer, and personal traits, such as gender and education [13, 32, 69]. In this study, we examine the effect of power use on Android users’ information privacy management strategies. Given the often-reported gap between users’ privacy intentions and their actual behaviors (i.e., *the privacy paradox*) [11, 24], we measure this effect both in terms of their intention to install apps and their intention to share information with these apps (i.e., their behavioral intentions), as well as in terms of the number of apps they install and the total number of permission requests for sensitive information they grant (i.e., their actual behaviors). We focus on Android users compared to iPhone users because Android’s application development environment is open and more permissive than iOS (e.g., Android users can circumvent the Google Play Store and download applications (APK files) from known or unknown sources on the web—which is not the case with iPhone users unless the device is jail-broken [30]). Android users also prominently interact with app permissions and are likely to be more privacy-conscious than iPhone users, given their increased vulnerability and exposure to privacy risks associated with the apps they install and/or use [57]. Nevertheless, even though Android users tend to exhibit more technology knowledge than iOS users, research shows that there

are no significant differences in privacy attitudes between the two platform users [2].

In terms of smartphone privacy, previous research has found that “*power use*” not only impacts privacy protection behaviors (such as deciding whether to install or uninstall an app due to the personal information it collects) directly, but also indirectly through privacy concerns and trust placed in mobile service providers [32, 41, 52]. Power use is defined as the competence, motivation, knowledge, expertise, and desire to use technology to its fullest [45]. Furthermore, power use mediates the effect of individual characteristics such as gender (e.g., females tend to use smartphones, install and use apps like e-commerce apps to a greater extent than men [36]) and education (e.g., people with higher levels of education tend to use smartphones to a greater extent than those with lower education [36]) on privacy behavior [30, 57]. This relationship between power use and privacy warrants further in-depth investigation. Thus, we contribute to this body of literature by studying how “*power use*” influences intentions and behaviors towards permissions management by addressing the following research questions:

- **RQ1:** *How do individual differences (i.e., gender, education) influence power use?*
- **RQ2:** *How does power use influence users’ behavioral intention (e.g., users’ intention to install Android apps and/or share data with those apps)?*
- **RQ3:** *How does power use influence actual user informational privacy behaviors (e.g., the actual number of Android apps installed and/or the total number of dangerous permissions granted)?*

We conducted a study with smartphone users (N=380) using Android OS 6.0 and above to examine the effect of power use on information privacy management intentions and actual privacy behaviors. Our analysis revealed an inconsistency in the effect of power use on users’ information privacy behaviors: while power use was *positively* related to the intention to install apps and to grant them access to their personal information, power use was *negatively* related to the actual installation of apps, and, in turn, the actual number of permissions granted. This study contributes to the privacy research community’s understanding of Android users’ privacy management behaviors and strategies. We also show another variant of the intention-behavior gap [62] by providing insight into the inconsistency of the effect of power on user privacy intentions and behaviors.

## 2 RELATED WORK

In the following subsections we first synthesize the related work on privacy in the context of mobile phones and smartphone apps. Then, we review the related literature on power use and information privacy.

### 2.1 Smartphone App Privacy and “Dangerous Permissions”

Several researchers have compared users’ attitudes towards being prompted for mobile app permissions at run-time (Android’s current strategy [27]) vs. install-time (the pre-Android 6.0 strategy) [7, 46]. Moore et al. [46] highlighted that each type of prompt was

able to accomplish a different purpose: Asking users to grant permissions at install-time kept users better informed about which type of information would be accessed, while a run-time prompt was better at communicating the reason why a certain type of information was being requested. However, their study found no clear evidence of whether the run-time prompts were more effective than the install-time requests in informing users about the implications of the permission requests. Our study is limited to users of Android 6.0 and up, thereby focusing on the run-time request model used by a vast majority of the Android devices currently in use.

Recent research has specifically focused on what Google classifies as “*dangerous permissions*” (e.g., access to location, camera, contacts), which have been shown to pose greater privacy risks [10]. Dangerous permission requests have increased over the years [73], with some apps evidently requesting more dangerous permissions than necessary for the proper operation of the app [29]. Some apps have even been shown to crash or become unusable when ostensibly unnecessary permissions are not granted [21], while other apps completely circumvent the run-time permission model and gain access to protected data without user consent [20]. Indeed, both researchers and journalists have increasingly reported privacy concerns around the covert collection and sharing of user data by Android applications [25, 29, 47, 55].

Much of this existing research around dangerous permissions does not consider the user experience, but instead involves the programmatic analysis of apps to identify the data flow and potential misuse of permissions [4, 19, 40, 64, 70, 70]. Notable exceptions are works evaluating user understanding of dangerous permission prompts and works focused on increasing the usability of such prompts [7, 8, 10]. We extend this body of work by studying factors that influence users’ intention to grant or reject dangerous permissions. In addition to participants’ self-reported intention, we measure the actual total number of dangerous permissions they granted on their smartphone. This is an important addition, given the well-documented finding that privacy intentions and behaviors do not always align (i.e., the “*privacy paradox*” [49]).

### 2.2 Power Users, Privacy, and Smartphones

One factor that arguably influences users’ intention and behavior regarding dangerous permissions is the concept of “*power use*”. Developed by Marathe et al. [45], power use relates to people’s competence, motivation, knowledge, and desire to use technology to its fullest. Power use is an important construct in studies that seek to determine people’s ability to adapt and use technology. For example, people who score high on the power use scale are motivated to learn about new technologies, spend a considerable amount of time using new gadgets, read and write online reviews of devices, push technological devices to their functional limits, and exert a greater amount control over their technology use, e.g., through customizable interface features/settings [32, 65, 80]. Scholars have also found that people who score high on the power use scale tend to enjoy improved outcomes, such as being influential on social media sites such as Instagram, usually due to their expertise and intensive use of the application(s) [60]. With increased usage, these users gain new knowledge, competence, experience and expertise with such technologies. Subsequently, they also prefer to control the

access and use of their personal information by these technologies [45, 65].

In fact, studies have demonstrated that the inherent desire for control among people who score high on the power use scale may partly be driven by their concern for online privacy [45, 65]. For example, Kang et al. [32] found that smartphone users who have higher expertise and seek higher levels of control may be more aware of potential risks presented by app permissions, thereby making them less likely to grant dangerous permissions to apps and subsequently less vulnerable to privacy invasions. Other work has also shown the connection between power use and other individual differences. For example, Zhong [80] found that the time one spends using a mobile device and their ability to multitask is a good predictor of power use. Zhong's work [80] also found that male smartphone users are more likely to know how to operate smartphones and score higher on power use than their female counterparts. However, this finding contrasts with Bonne et al.'s [13] work, which found that female smartphone users were less likely to grant app permissions than their male counterparts, thereby challenging the permission granting (i.e., information privacy management) behaviors of people who score high on power use, especially along gender characteristics. Our research investigates this influence of gender (together with education) on power use to specifically examine whether power use may help explain some of the conflicting results found in prior research [13, 80]. In the next section, we present our research framework.

### 3 RESEARCH FRAMEWORK

The purpose of this study is to examine the relationships between Android users' power use (i.e., participants' ability, expertise, and experience using Android smartphone apps to the fullest capability), their actual information privacy management behaviors (i.e., installing apps and granting permissions) (RQ3), and their expressed intentions toward those behaviors (RQ2). Furthermore, we study the influence of individual differences on power use (RQ1). In the following sections, we introduce the constructs and hypothesized relationships that constitute our research model (Fig. 1).

#### 3.1 Smartphone Privacy Behaviors

In our study, we examine two interrelated information privacy behaviors on smartphone devices: 1) the number of smartphone applications installed on the participant's device, and 2) the total number of dangerous permissions granted to these apps. Importantly, we label these metrics as "actual behaviors" because they are scraped directly from the app manifest on the participant's smartphone. One of the unique contributions of our work is a simultaneous investigation of the effects of power use on users' actual information privacy management behaviors as well as their intentions towards those behaviors. We provide more details about the measured actual behaviors below.

**3.1.1 Number of Apps Installed.** Smartphone users are concerned about the access that apps (or third-party libraries within these apps) have to their sensitive personal information [15, 35, 40] and these concerns significantly influence the number of apps they install [13, 34, 42, 57]. Bonne et al. [13] found that information privacy concerns influence the ratings and reviews of apps to such an extent

that users will uninstall apps if they are uncomfortable granting them permissions or if they feel that the apps should not have access to certain permissions. Therefore, even though the number of apps a user has installed on their smartphone is traditionally considered a "usage" behavior rather than a "privacy" behavior, it is an important behavioral indicator of the user's information privacy management practices [11].

As discussed in Section 2.2, people who score higher on the power use scale tend to be more motivated to learn about or use new technologies to their fullest [32, 65]. As such, they are likely to have more mobile applications installed. In other words, we expect power use to have a significant effect on the number of apps installed on a user's device:

- **H1:** Power use will be positively associated with the total number of apps installed.

**3.1.2 Total Number of Dangerous Permissions Granted.** While permission requests are intended to give users control over a smartphone app's use of personal data, in reality Android users often grant permissions to apps with vague descriptions and unclear purposes [22]. The subsequent access apps have to their personal information stored or accessed through their devices increases users' privacy risk and vulnerabilities. In an attempt to distinguish between less risky "regular" permissions (e.g., access to mobile networks, WIFI networks, Bluetooth, audio settings, etc.) and more dangerous ones, Android created a designation of "dangerous permissions" to requests for access to private user data, such as a user's location, calendar, call logs, camera, contact, microphone, phone, sensors, SMS and storage (see the full list in Appendix B) [27]. The "dangerous permission" classification was intended to help developers understand that these permissions arguably have stronger implications for users' information privacy than regular permissions and thus be cautious in requesting them in their applications. However these permissions are still highly requested [13]. Moreover, Android users are over five times more likely to grant these permissions than to deny them [13]. Therefore, we are interested in understanding the factors that contribute to the number of dangerous permissions Android users grant to apps installed on their smartphones.

Research also indicates that smartphone users who score higher on the power use scale are less likely to be vulnerable to privacy risks due to a stronger desire to control their privacy, which is arguably instilled by their relatively greater knowledge about mobile applications [32]. For example, people who score higher on the power use scale are less likely to share their personal information on personalized mobile sites [65]. In terms of smartphone apps, power use is thus likely negatively related to the total number of permissions granted—although this effect may be countered by the hypothesized (H1) positive association between power use and number of apps installed (since having more apps installed generally means granting more permissions). Therefore, we hypothesize the following:

- **H2:** Power use will be negatively associated with total dangerous permissions granted.
- **H3:** The number of installed apps will be positively associated with the total number of dangerous permissions granted.

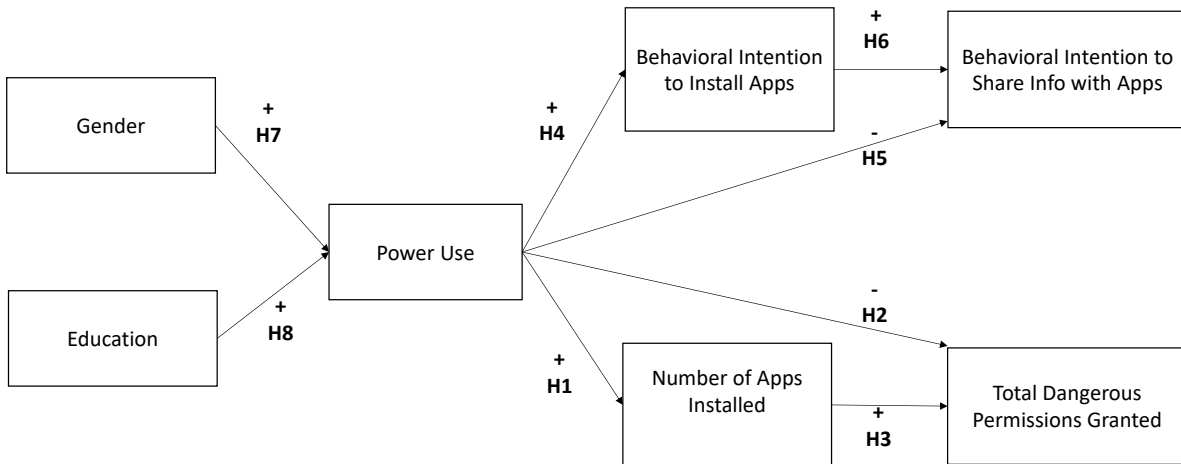


Figure 1: The proposed research model

### 3.2 Behavioral Intention

Behavioral intention is the measure of the strength of an individual’s “self-prediction” or “behavioral expectation” to perform a specified action, and thus one of the most accurate predictors of an individual’s future behavior [71]. Simultaneously, though, one of the most consistent findings in the field of privacy is a gap between users’ intended and actual behaviors (see [11, 24] for an overview). Labeled the “Privacy Paradox” [49], this phenomenon suggests that people tend to share more personal information than they claim to intend to share when asked beforehand (i.e., people behave contrary to what they say). Barth and De Jong [11] suggest that the privacy paradox within the mobile context could be different from other contexts such as social media and other online media where it has predominantly been studied. Concurrently, we suggest that the predominance of the privacy paradox may depend on the users’ expertise and experience (i.e., “power use”), as those who have more expertise/experience may more consistently engage in what would be considered appropriate protective behavior. Therefore, in this study, we assess the effect of power use on both users’ intention to use smartphone apps and share personal information (i.e., “what people say they do”) as well as their actual app installation and permission granting behaviors (i.e., “what people actually do”). The effects of power use on intentions are outlined below.

**3.2.1 Behavioral Intention to Install Apps.** Wang et al. [67] found that function, social, emotional, and epistemic value all play a role in users’ intention to use mobile apps, while Xu et al. [77] found that mobile users’ information privacy concern was a significant factor in their behavioral intention to use mobile apps and share personal information with them. Given that people who score higher on

the power use scale tend to be more motivated to learn about new technologies to their fullest [32], we expect that they similarly have higher intentions to use apps. We thus hypothesize the following:

- **H4:** Power use will be positively associated with users’ intention to install apps.

**3.2.2 Behavioral Intention to Share Information with Apps.** In a study of users’ behavioral intention to disclose information to apps, Keith et al. [33], revealed significant correlations with several privacy risky behaviors, such as sharing registration information, location services, store credit cards, and profiles. However, we argue in line with Xu et al. [77] that behavioral intention to use apps and the behavioral intention to share information with them are two separate issues, especially given that Android’s run-time permissions model now allows for more granular privacy management [7, 46]. Therefore, while we hypothesize that behavioral intention to use apps will be positively associated with behavioral intention to share information with apps, we treat them as two separate constructs that can vary independently from one another.

Research has shown that smartphone users who score high on the power use scale are less likely to be vulnerable to privacy risks due to a desire to control their privacy [32]. Thus, despite the positive association between power use and users’ intention to install apps, we expect that power use is negatively related to users’ intention to share their information with the apps they install. Therefore, we hypothesize the following:

- **H5:** Power use will be negatively associated with users’ behavioral intention to share information with apps.
- **H6:** Users’ intention to use apps will be positively associated with their intention to share information with these apps.

### 3.3 Individual User Differences: Gender and Education

Finally, it is important to consider individual differences that may affect people's level of power use. Prior work has indicated that smartphone and application usage differences can vary based on demographic factors, such as gender and level of education [58, 61, 79]. In the realm of Android permission management, Bonne et al. [13] found that women, on average, deny permissions twice as often as men across all age groups, while Wash and Rader [72] found that less-educated individuals are more likely to believe there is nothing they can do to protect their privacy and subsequently are less likely to act in a privacy-conscious manner. These individual differences (e.g., gender and education) tend to be strong predictors of power use [80]. Thus, we hypothesize the following effects of gender and education on power use among Android users:

- **H7:** Females will score higher on the power use scale than males.
- **H8:** Education will be positively associated with power use.

In the next section, we describe our methods for examining the model (see Fig. 1).

## 4 METHODS

In this section we provide an overview of our study, explain how we operationalized the constructs in our research framework, and describe our data analysis approach.

### 4.1 Study Overview

The objective of our study is to examine the effect of power use on Android users' information privacy behaviors (i.e., installing apps and granting permissions) and expressed intentions toward those behaviors (i.e., users' ability, expertise, and experience using apps to the fullest capability), and the influence of users' individual differences on power use. Our study was administered through an Android app that was made available for installation via Google Play. We recruited participants via Amazon Mechanical Turk, because users on that platform are more tech-savvy than their peers and thus were less likely to experience issues installing our app [56].

Upon installation and with user consent, our "UCF Permissions" app scraped the Android run-time "dangerous permissions." While the data was being scraped in the background, participants answered a brief in-app survey measuring their level of power use, intention to download and use applications within the next three months, intention to grant access to or share information with applications within the next three months, and gender and educational background (see Appendix A for the full survey).

Given the privacy-sensitive nature of scraping user mobile data, we explicitly sought user consent to collect background information about their installed applications before downloading the study app. We also made it clear within the study description that the data collection would be limited to information regarding the dangerous permissions requested by these applications. To avoid priming users, we were careful not to use the term 'privacy' anywhere in the study description. All study procedures were reviewed and approved by the Institutional Review Board of University of Central Florida. In the next subsection we provide a detailed explanation of Android's

run-time permissions framework and how we used our app to scrape participants' dangerous permissions granted to the apps on their phone.

### 4.2 Leveraging Android's Run-time Permissions

Starting with Android 6.0 and above, users have been afforded a more granular ability to selectively and explicitly approve permissions to apps at run time via a system dialog rather than at install time [27]. Android permissions are divided into three protection levels based on the level of risk that they present to the user's privacy or the operation of other apps: *normal*, *signature* and *dangerous* [27], with only the latter requiring explicit user approval. Dangerous permissions cover situations in which apps want to access data or resources that involve the user's private information and/or situations that potentially affect the users' stored data or the operation of other apps on their device [27]. For example, the ability to access a user's exact location, read their contacts, access their device's camera, and read from or write to the external storage are all considered dangerous app operations that users have to consent to explicitly [26].

Our study app used the *PackageManager* class of the Android SDK to get information about the "Number of Apps Installed" (i.e., the number of applications installed on the participant's mobile device as identified by their package names in the *Android Manifest.xml* file) and the "Total Dangerous Permissions Granted" (i.e., the aggregate count of the dangerous permissions granted to the installed applications). Specifically, the study app collected information regarding the installed apps (i.e., app and package name) and the corresponding "dangerous" permissions granted to each app. Permission data was collected by recording whether each of the 23 dangerous permissions was not explicitly requested by the app (-1), requested but denied by the user (0), or requested and granted by the user (1). The total number of "dangerous" permissions for each user was a summation of the number of the permissions that were requested and granted by the user (1).

### 4.3 Survey Design and Operationalization of Constructs

For our subjective measures, we used pre-validated survey scales to assess power use [45], intention to install apps, and intention to share information with apps [78] (see Appendix A for the full survey). For power use, participants were asked 12 items adapted from Marathe et al. [45] (e.g., "I make good use of most of the features available in any technological device"), where each item had response options on a 5-point scale, ranging from 1 (*Strongly Agree*) to 5 (*Strongly Disagree*) (see Appendix A.1).

Our behavioral intention measure was adapted from Xu et al. [78] who modeled it based on earlier works on the Technology Acceptance Model [17]. Unlike Xu et al. [78], the time frame referent of our intention measures was reduced from 12 months to 3 months (e.g., "I am likely to disclose my personal information to use mobile apps in the next 3 months"), because there has been an expeditious increase in the rate of adoption and usage of mobile applications since 2004 when Xu et al. [78] published their work. We divided the behavioral intention measures between *intention*

to install mobile applications (see Appendix A.2) and intention to share information with these apps (see Appendix A.3) because with participants' heightened privacy concerns about their personal information [9] these intentions might have diverged over time and thus no longer represent the same concept. Furthermore, the *behavioral intention to share information with apps* scale was combined with two additional items that specifically inquired about users' willingness to grant dangerous permissions. Rather than asking leading questions that directly mention "dangerous permissions", the *location sharing* permission (one of the most predominantly requested dangerous permissions) was used as an example of a dangerous permission that regular users are likely to understand [50, 51, 78].

#### 4.4 Procedure and Participant Recruitment

The study recruited Android users using the Amazon Mechanical Turk platform<sup>1</sup> as its sampling population. MTurkers were required to have a historical HIT approval rate greater than 95% with at least 50 approved past HITs to ensure satisfactory response quality. Furthermore, participation was limited to U.S. adults (age 18+) with Android devices updated to use Android 6.0 or later for the study app to function appropriately. Upon reading the study description, participants who consented to participate in the study were provided with a link to download and install the app from the Google Play store. The study description explicitly disclosed our study app's background data collection intention and requested user consent. Additionally, participants were advised to delete the app after participation.

Upon completion of the survey and the background scan, a random unique completion code was generated and entered on Amazon Mechanical Turk as proof of study completion. All Participants who completed the study or could not complete the study due to technical difficulties, were compensated \$1. Pilot testing suggested that the study's duration ranged from 15 to 30 minutes on average. In line with Amazon's policy of anonymized data collection [5], the study app did not collect any other personal information apart from the package name of the installed applications (unique across apps) and data regarding the 23 dangerous permissions (see Appendix B, Table 4).

Overall, 429 MTurkers accepted the HIT and participated in the study. Upon analysis, we found that 49 participants failed our attention check questions. After discarding their responses, we were left with valid data from 380 participants. Our sample was relatively gender-balanced (195 males and 181 females). About 33% of the participants reported completing an Associates' degree, and (55.5%) completed at least a four-year college degree. Table 1 further describes the demographics of our participants.

#### 4.5 Data Analysis Approach

For constructs that were measured using prevalidated scales (i.e., power use, behavioral intention to use mobile apps, and behavioral intention to share information with apps (see Appendix A)), we checked the scale reliability (i.e., the extent to which all the items in a scale measure the same construct) using Cronbach's alpha. We then created indices for each of these constructs by averaging

<sup>1</sup><https://www.mturk.com/>

Variables		Total (N=380)	Percent (%)
Gender	Male	195	51.3
	Female	181	47.6
	Other	2	0.52
	Did not specify	2	0.52
Education	High school	40	10.5
	Associate's degree	124	32.6
	Bachelor's degree	67	17.6
	Some Grad, but no degree	94	24.7
	Master's degree	16	4.2
	Doctoral degree	34	8.95
Race	White	253	66.6
	Black or African American	46	12.1
	Hispanic or Latino	26	6.8
	Asian	20	5.3
	American Indian or Alaska Native	2	0.5
	Not Specified	33	8.7

**Table 1: Demographics (gender, education, and race) of our study participants.**

across all scale items (see Table 2). The number of apps installed was computed as the number of apps installed on the participant's device, excluding our study app and any other apps known to come pre-installed on Android devices. The total number of permissions granted was computed as the number of permissions granted to each app, summed over all of the participant's apps.

Variable	Variable type	Mean	Median	SD	Cronbach's Alpha
Power use	Subjective measure	2.21	2.17	0.46	0.79
Behavioral intention to install apps	Subjective measure	1.60	1.50	0.75	0.86
Behavioral intention to share information with apps	Subjective measure	2.24	2.00	0.99	0.80
Number of installed apps	Scraped behavior	51.58	42.00	35.30	N/A
Total dangerous permissions granted	Scraped behavior	67.98	59.00	46.17	N/A

**Table 2: Descriptive statistics of the dependent, independent variables and internal consistency reliability of our survey measures**

The statistical significance of the hypothesized relationships between the constructs in our research model was tested using a path model (see Fig. 2) that included the participant's gender and education level, the self-reported subjective measures on behavioral intention (i.e., intention to install and share information with apps), and the scraped actual behavioral data (i.e., the total number of apps and the total dangerous permissions granted). This path model can be seen as a series of linear regressions that together describe the statistically significant paths between the scraped behaviors and subjective measures [39].

We examined the sign and significance of the path coefficients in R; in our resulting path models (see Fig. 2), the solid incoming arrows ( $\rightarrow$ ) between constructs represent significant relationships while the broken line arrows ( $\dashrightarrow$ ) represent tested relationships that were found to be non-significant. Each linear regression contains a regression coefficient (indicated by the number on the arrow as well as its thickness), the standard error of the regression effect (in parenthesis), and the significance level denoted by asterisks (or "ns" for non-significant effects). The subjective constructs were scaled to have a standard deviation (SD) of one (1.0) so that one SD difference in a construct (e.g., power use) causes a  $\beta$  SD difference in another construct (e.g., behavioral intention to use apps). The

Hypothesis	Supported?
H1: Power Use → number of apps installed (+)	No
H2: Power Use → total dangerous permissions granted (-)	No
H3: Number of apps installed → Total Dangerous Permissions Granted (+)	Yes
H4: Power use → intention to install Apps (+)	Yes
H5: Power use → intention to share information with Apps (-)	No
H6: Intention to install apps → intention to share info with installed Apps (+)	Yes
H7: Females → power use (+)	Yes
H8: Education → power use (+)	No

**Table 3: Hypothesis test results**

results are graphically presented in Figures 2 and the outcomes of our hypotheses are summarized in Table 3.

Finally, we expanded our model in a post hoc analysis and provide additional results regarding the average number of dangerous permissions granted and the likelihood of granting specific permissions (see Section 5.3).

## 5 RESULTS

Below, we describe our study’s findings. We first provide descriptive statistics regarding the apps participants frequently had installed on their Android devices, and the dangerous permissions most commonly granted to those apps. Then, we present our hypotheses test results, followed by a post hoc analysis to further unpack additional nuances in our data.

### 5.1 Descriptive Statistics and Internal Consistency of Model Constructs

Our participants (N=380) had a total of 6727 unique applications installed on their devices. The average number of applications installed per participant was 51, with a standard deviation of 35 apps. Among our participants, the maximum number of installed applications on a single device was 226, and the minimum was two. The ten most common applications installed by our participants are listed in Appendix C, Table 5. We did not consider system apps<sup>2</sup> (e.g., Google Play Store, Camera, Contacts, Gallery, etc) in our analysis, as such apps normally come pre-installed on Android devices by the manufacturer without any explicit user input in the installation decision. Next, we examined the likelihood of an installed app requesting each of the dangerous permissions and the extent to which participants granted them. The most requested dangerous permissions were WRITE\_EXTERNAL\_STORAGE (73%), ACCESS\_COARSE\_LOCATION (32%), CAMERA (29%), and READ\_PHONE\_STATE (29%). Similarly, the most granted dangerous permissions were WRITE\_EXTERNAL\_STORAGE (100%), ACCESS\_FINE\_LOCATION (100%), READ\_PHONE\_STATE (100%), GET\_ACCOUNTS (100%), and CAMERA (99.7%). Table 4 in Appendix B lists the likelihood of each of the permissions being requested and granted.

Prior to testing our model, we calculated Cronbach’s alpha for each of the constructs in our model to ensure scale reliability (i.e., the extent to which the survey items that constitute a scale measure

<sup>2</sup>The apps (i.e., “bloatware”) were excluded based on their known package names associated with the android or phone manufacturer in tandem with app lists such as [1, 54]. Thus, third-party apps (e.g., Gmail, Google Duo or Facebook) that at times also come pre-installed were not excluded given that users have a choice on whether to use them and grant them “dangerous” permissions

the same construct [66]; see Table 2). Acceptable values of Cronbach’s alpha range from 0.70 (acceptable) to 0.95 (excellent) [66]. Our scales for power use and behavioral intention all had good reliability.

### 5.2 Hypotheses Testing Results

The results of our hypothesis tests are summarized in Figure 2 and Table 3. We discuss them in more depth below.

**5.2.1 Effects of Power Use on Privacy Behaviors (H1–H3).** Power use is associated with the number of apps installed. However, this effect is *negative* rather than *positive* as hypothesized ( $\beta = -0.168$ ,  $p < .001$ ; **H1** is not supported). Furthermore, there is no direct effect of power use on the total dangerous permissions granted ( $\beta = -0.085$ ,  $p = .061$ ; **H2** not supported). This effect is instead mediated by the number of apps installed ( $\beta = 1.394$ ,  $p < .001$ ; **H3** supported), with power use in effect being *negatively* associated with the total number of dangerous permissions granted (total effect:  $\beta = -0.234$ ,  $p < .001$ ).

**5.2.2 Effects of Power Use on Behavioral Intention (H4–H6).** Power use is positively associated with users’ intention to install apps ( $\beta = 0.635$ ,  $p < .001$ ; **H4** supported). Power use is also *positively* associated (rather than negatively, as hypothesized) with users’ intention to share information with apps ( $\beta = 0.237$ ,  $p < .01$ ; **H5** not supported). Additionally, we do find support for **H6** in a significant association between participants’ intention to install apps and their intention to share information with these apps ( $\beta = 0.498$ ,  $p < .001$ ). The total effect<sup>3</sup> of power use on users’ intention to share information with these apps is  $\beta = 0.553$  ( $p < .001$ ).

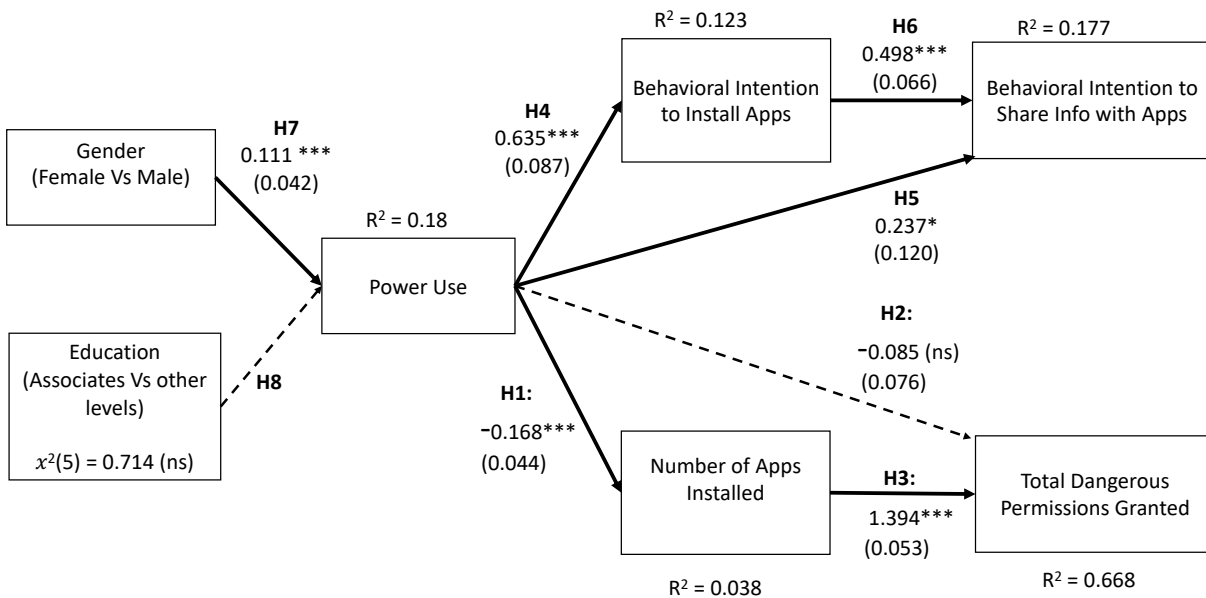
**5.2.3 Effects of Gender and Education on Power Use (H7–H8).** Gender has a significant effect on power use: women score on average higher on the power use scale than men ( $\beta = 0.111$ ,  $p < .001$ ; **H7** supported). This finding dispels the stereotype that women are lower-skilled mobile phone users, especially in contexts where smartphones are viewed as “hi-tech” gadgets [16]. Interestingly, the participants’ education level does not significantly affect power use (**H8** not supported).

Overall, our findings (see Fig. 2) show a positive association between power use privacy intentions—both in terms of users’ intention to install apps as well as their intention to share information with these apps. However, an examination of their actual privacy behaviors shows a negative association between power use with the number of apps installed that subsequently mediates the effect of power use on the total number of dangerous permissions granted (which is also negative). We investigate these results further in our post hoc analyses, wherein we examine this inconsistency in the effect of power use on users’ privacy intentions versus their actual behaviors.

### 5.3 Post Hoc Analyses

Note that while power use is positively related to the intention to install apps (see Fig. 3a) and to grant them access to personal

<sup>3</sup>The total effect is “the sum of the direct and indirect effects of the exogenous variable on the outcome” variable. [28]



**Figure 2: Path modeling results. (The broken line and ns - showcase the non-significant relationships. The straight lines showcase the significant relationships and p-levels: \*\*\*  $p < .001$ , \*\*  $p < .01$ )**

information (see Fig. 3b), it is negatively related to the actual installation of apps (see Fig. 3c), and in turn, the actual number of permissions granted. On a similar note, we find that there is no relationship between power use and the *average* number of dangerous permissions granted *per app* (see Fig. 3d; the Kendall Rank Correlation between power use and average number of dangerous permissions is  $\tau = -0.31$ ,  $p = 0.38$ ). This suggests that participants who score lower on the power use scale are more likely to be exposed to privacy risks due to the high number of apps that they had installed, rather than the number of permissions granted to each app. Conversely, this suggests that participants who score higher on the power use scale either install fewer apps or uninstall apps more frequently in an effort to shield themselves from privacy risks.

Further unpacking the latter result, we ran 23 logit models to test the relationship between power use and the *likelihood* of each specific permission being granted *if requested*. We found no significant relationships except one: power use is significantly negatively related with the likelihood of granting the ACCESS\_FINE\_LOCATION permission ( $b = -0.408$ ,  $p < .01$ ) (see Fig. 4). Specifically, as power use increases by one standard deviation (1.0 SD), the odds of granting this permission decrease by 33.5%. This suggests that participants who score higher on the power use scale are particularly more cautious about applications that request access to their fine-grained location.

## 6 DISCUSSION

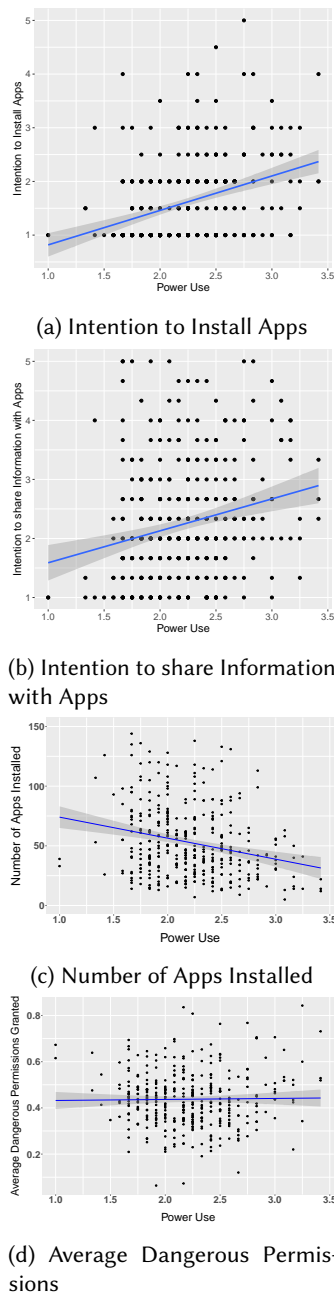
Our results find an interesting inconsistency in the effect of power use on intention versus actual behavior, which has interesting implications for predictive modeling as well as for design and for future research. These implications are discussed below.

### 6.1 The Inconsistency in the effect of Power Use on Privacy Intentions versus Behavior

Our results indicate that despite the positive relationship between power use and users' intention to install apps and share information with these apps, there is a negative relationship between power use and the actual number of apps installed and, subsequently, the total number of dangerous permissions granted. In other words, our findings show an inconsistency in the effect of power use (i.e., users' expertise and experience using mobile applications) on their information privacy intentions versus their actual behaviors: while their intentions seem to suggest otherwise, our results on behavior are actually consistent with prior work that has shown power use to be related to more information privacy-protective behaviors (i.e., the decision to install or uninstall apps due to the personal information they access/collect) [32]. Our findings confirm that indeed the emergence of a privacy paradox could be context-dependent: the mismatch between privacy intentions and behavior is different at different levels of power use [11].

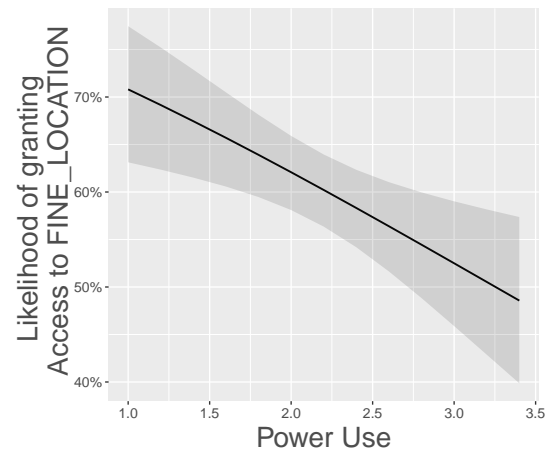
Why does a disconnect between intention and actual behavior matter in this case? A key precept in user-centered privacy research is that the over-arching goal is to help users achieve their privacy goals. In many cases this does not necessarily mean giving





**Figure 3: Relationships between power use and application installations (a) and the intention to (b) install apps (c) share information with these apps (d) number of dangerous permissions granted**

them *more* privacy at all costs, but rather allowing them to reach the right balance between privacy and benefits—e.g., in social networks, between privacy and connectedness [74, 76]. This implies that we should care about users’ intentions and whether those intentions are actualized [38]. Whenever the link between intentions and behaviors is broken, it is thus worth digging deeper: Why did



**Figure 4: Logit model showing the likelihood of granting the ACCESS\_FINE\_LOCATION permission when requested, at different levels of power use.**

participants who scored higher on the power use scale intend to install fewer apps? One reason that aligns with past research [32] could be that users who score higher on the power use scale are privacy conscious; therefore, they are able and motivated to perform a privacy calculus [43] to determine when an app should *not* be installed, despite their intention to do so. This could manifest in a refusal to install certain apps, or a practice of more judiciously uninstalling apps that are no longer deemed useful enough to justify the associated privacy invasion<sup>4</sup>. It is possible that users who score higher on the power use scale perform some level of “garbage clean up” on their devices by removing unused or unwanted apps. Conversely, it is possible that users who score lower on the power use scale, out of habit [12], install and never remove a substantial number of apps, exposing themselves to privacy risks depending on the dangerous permissions they grant to each app.

Another perspective on the concept of power use [65] suggests that a higher level of power use is characterized by an increased *depth* of use (i.e., using the advanced features of apps), rather than an increased *breadth* of use (i.e., using more apps). This perspective could be investigated by a more in-depth analysis of the types of apps used by people with various levels of power use (i.e., advanced apps versus simple apps).

## 6.2 The Predictive Power of Behavioral Intention versus Power Use

By measuring both behavioral intention and actual behavior, our work is able to reveal the privacy paradox between intentions and actual behavior that often manifests (but is less often directly demonstrated) in privacy research. Behavioral intention is cited in the social sciences as the best predictor of actual behavior [3, 17], but as many other researchers who have studied the privacy paradox

<sup>4</sup>research shows that most “disliked” apps are likely to be uninstalled within 2 days [44]

have found (cf. [24]), we also found this claim not to be consistent with our findings. In fact, our study shows power use to be a better predictor of actual privacy-related behaviors (in terms of the number of installed apps and total permissions granted) than intention.

Overall, predicting users' actual privacy behavior is difficult, and the privacy paradox does not provide useful alternatives to the traditional intention-based frameworks [49]. The explanatory power of models that try to predict actual privacy behaviors is traditionally low, commonly ranging from less than one percent to five percent [33]. Our results demonstrate that hybrid models of behavioral intention, other perceived constructs, and behavioral scraped data can potentially work together to improve our understanding and predictive power of information privacy behavior [59]. Specifically, our results imply that researchers who are cognizant of the privacy paradox and looking to update traditional intention-based frameworks of predicting users' behavior through their intentions may consider power usage as a new and perhaps better construct that is more indicative of users' information privacy behaviors.

### 6.3 Implications for Design

There are several implications for design that can be inferred from our results. Such recommendations are most pertinent for users with low levels of power use, who are at a higher risk of privacy violations. Given that these users' exposure to privacy vulnerabilities is mostly due to the number of apps they have installed (rather than the number of permissions granted per app), we particularly emphasize the need for more comprehensive privacy support in installing and managing *apps* rather than individual permissions. In this light, Android's shift to at-runtime permission requests rather than at-installation permission requests could have obfuscated the influence of app installation on user privacy. Given the absence of a clear privacy statement or labels at the time of app installation, we recommend that Android's app store could re-emphasize the permissions an app is likely to request at some point in the app installation process (or perhaps in the app discovery process, cf. [18, 34]). Another suggestion is to run a background process that periodically tries to identify unused apps and recommends that they be uninstalled or that their access to dangerous permissions is revoked. This is a process similar to iPhone's "offload unused apps" setting [63] and could potentially help disrupt the habituation associated with installing apps and granting permissions to them, but not leveraging their full value.

For users with high levels of power use, the approach may be different. When installing a new app, it is likely important to engage them quickly and show the value of granting dangerous permissions to the app in terms of personalized functionality [4, 23] that benefits the user. Otherwise, they are apt to uninstall the app. Further, giving users with higher levels of power use more granular access to customize their information privacy permissions sooner—and being transparent as to why granting such permissions is valuable—will assist them in making a calculated decision on whether to install/keep the app [43]. Note, though, that our results suggest that making the permissions themselves more granular may not work because users—regardless of their level of power use—seem to make privacy decisions on a per-app basis rather than customizing these

permissions more granularly. Therefore, Google may reconsider shifting to a combination of both more granular and install-time privacy permissions in Android, or attempt to find new ways to nudge Android users to more effectively leverage the more granular level of control available [37]. Additionally, similar to the iOS app store, Android's app store can also use "privacy nutrition labels" [34] to help users quickly understand and digest what types of personal data are likely to be requested by the app before they install it [14, 31]. Finally, similar to Wang et al.'s study on Facebook privacy permissions [68], we recommend that smartphone apps request the minimum number of dangerous permissions necessary and appropriate for providing the customized functionality that gives the app its value. Doing so would help users more easily make privacy calculus [43] assessments as to whether the benefits of an app warrant the risks.

### 6.4 Limitations and Future Work

We recruited our participants from Amazon Mechanical Turk and restricted the participant pool to only users within the U.S. Whereas this reinforces the quality of our data (users on MTurk are more tech savvy than their peers [56] and thus less likely to botch the installation of our app), it limits the generalizability of our findings to U.S. Android users. Additionally, we acknowledge that the demographics of Amazon MTurkers may deviate from the general population of all Android users. Specifically, our findings related to gender (i.e., females score higher than males on the power use scale) may be confounded with the fact that these individuals may have simply been MTurkers who just happen to be women. Future studies can build upon our research and examine if these findings translate to more diverse populations of Android users and other demographic traits such as age and race.

We excluded known system applications pre-installed by the device manufacturer or mobile service provider (i.e., "bloatware") [1, 54], because we wanted to ensure that our analysis was based on *user-installed* applications. However, given the fractured and evolving landscape of Android-based phones, we may have mistakenly excluded some user-installed apps or overlooked excluding some system applications that did not come to our purview, which could have perhaps affected our outcomes or the interpretation of our results. Moreover, our analysis only considered the 23 dangerous permissions classified by Google in 2018 (see Appendix B, Table 4) and explicitly declared by apps. Future work could expand our study to include all permissions that can possibly be requested by Android applications. Our analysis was also based on a snapshot of the permissions that our participant's had set at the moment they installed our study app. Whereas the permissions captured are reflective of their actual behavior at the time, future work could examine if the behaviors observed change over time by employing a longitudinal study. This would also resolve whether users with high levels of power use are more likely to remove unused apps, or whether they simply install fewer apps to begin with.

## 7 CONCLUSION

In this paper we examined the effect of power use on users' intention and privacy behavior (i.e., the number of apps a user has

installed on their device and the total number of dangerous permissions granted to those apps). Our research highlights the importance of examining both users' behavioral intentions and their information privacy behaviors as a way to uncover additional privacy paradoxes [11], reinforcing user-centered privacy research [48, 74], rather than pushing a one-size-fits-all privacy-focused agenda [75] that indiscriminately nudges all users towards more restrictive data practices. We did this by unpacking a seeming inconsistency in the effect of power use: While the intention to install apps and to share information with these apps increased with power use, an inspection of participants' smartphone settings revealed that people who scored higher on the power use scale had fewer apps installed and, subsequently, fewer dangerous permissions granted than their counterparts who scored lower on the power use scale.

However, on average, both participants who scored lower and higher on the power use scale granted the same number of dangerous permissions *per app*, indicating that the difference in privacy vulnerability was mostly caused by a difference in the number of apps installed on their devices. To better understand and uncover these kind of privacy behaviors, we encourage researchers to study the diverse range of smartphone users under different contexts.

## ACKNOWLEDGMENTS

We would like to thank the individuals who participated in our study. This research was supported by the U.S. National Science Foundation under grants CNS-1814068, CNS-1814110, and CNS-1814439. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. National Science Foundation. Moses Namara acknowledges support from a Facebook Fellowship Award.

## REFERENCES

- [1] 2021. Pre-installed apps: T-Mobile REVVL. <https://www.t-mobile.com/support/devices/android/t-mobile-revvl/pre-installed-apps-t-mobile-revvl>
- [2] Desiree Abrokwa, Shruti Das, Omer Akgul, and Michelle L Mazurek. 2021. Comparing Security and Privacy Attitudes Among US Users of Different Smartphone and Smart-Speaker Platforms. In *Seventeenth Symposium On Usable Privacy and Security (SOUPS) 2021*. 139–158.
- [3] Icek Ajzen. 1985. From intentions to actions: A theory of planned behavior. In *Action control*. Springer, 11–39.
- [4] Efthimios Alepis and Constantinos Patsakis. 2017. Hey doc, is this normal?: exploring android permissions in the post marshmallow era. In *International Conference on Security, Privacy, and Applied Cryptography Engineering*. Springer, Cham, 53–73.
- [5] Amazon. 2020. Participation Agreement. <https://www.mturk.com/participation-agreement>
- [6] Monica Anderson. 2015. Mobile apps, privacy and permissions: 5 key takeaways. <https://www.pewresearch.org/fact-tank/2015/11/10/key-takeaways-mobile-apps/>
- [7] Panagiotis Andriotis, Shancang Li, Theodoros Spyridopoulos, and Gianluca Stringhini. 2017. A comparative study of android users' privacy preferences under the runtime permission model. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*. Springer, 604–622.
- [8] Panagiotis Andriotis, Martina Angela Sasse, and Gianluca Stringhini. 2016. Permissions snapshots: Assessing users' adaptation to the Android runtime permission model. In *2016 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 1–6.
- [9] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. 2020. Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information. <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>
- [10] Rawan Baalous and Ronald Poet. 2018. How Dangerous Permissions are Described in Android Apps' Privacy Policies?. In *Proceedings of the 11th International Conference on Security of Information and Networks*. 1–2.
- [11] Susanne Barth and Menno DT De Jong. 2017. The privacy paradox—Investigating discrepancies between expressed privacy concerns and actual online behavior—A systematic literature review. *Telematics and informatics* 34, 7 (2017), 1038–1058.
- [12] Grant Blank, Gillian Bolsover, and Elizabeth Dubois. 2014. A new privacy paradox: Young people and privacy on social network sites. In *Prepared for the Annual Meeting of the American Sociological Association*, Vol. 17.
- [13] Bram Bonn , Sai Teja Peddinti, Igor Bilogrevic, and Nina Taft. 2017. Exploring decision making with Android's runtime permission dialogs using in-context surveys. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS) 2017*. 195–210.
- [14] Ian Carlos Campbell. 2020. Apple will require apps to add privacy 'nutrition labels' starting December 8th. <https://www.theverge.com/2020/11/5/21551926/apple-privacy-developers-nutrition-labels-app-store-ios-14>
- [15] Saksham Chitkara, Nishad Gothoskar, Suhas Harish, Jason I Hong, and Yuvraj Agarwal. 2017. Does this app really need my location? Context-aware privacy management for smartphones. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 1–22.
- [16] Francesca Comunello, Mireia Fern andez Ardeol, Simone Mulargia, and Francesca Belotti. 2017. Women, youth and everything else: Age-based and gendered stereotypes in relation to digital technology among elderly Italian mobile phone users. *Media, Culture & Society* 39, 6 (2017), 798–815.
- [17] Fred D Davis. 1985. *A technology acceptance model for empirically testing new end-user information systems: Theory and results*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [18] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. 2009. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proceedings of the 27th international conference on Human factors in computing systems (CHI '09)*. ACM, New York, NY, USA, 319–328. <https://doi.org/10.1145/1518701.1518752>
- [19] William Enck, Machigar Ongtang, and Patrick McDaniel. 2009. Understanding android security. *IEEE security & privacy* 7, 1 (2009), 50–57.
- [20] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopoulos, Ramesh Govindan, and Deborah Estrin. 2010. Diversity in smartphone usage. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*. 179–194.
- [21] Zheran Fang, Weili Han, Dong Li, Zeqing Guo, Danhao Guo, Xiaoyang Sean Wang, Zhiyun Qian, and Hao Chen. 2016. revdroid: Code analysis of the side effects after dynamic permission revocation of android apps. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*. 747–758.
- [22] Johannes Feichtner and Stefan Gruber. 2020. Understanding Privacy Awareness in Android App Descriptions Using Deep Learning. In *10th ACM Conference on Data and Application Security and Privacy*.
- [23] Carol J Fung, Bahman Rashidi, and Vivian Genaro Motti. [n.d.]. Multi-View Permission Risk Notification for Smartphone System. ([n. d.]).
- [24] Nina Gerber, Paul Gerber, and Melanie Volkamer. 2018. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security* 77 (Aug. 2018), 226–261. <https://doi.org/10.1016/j.cose.2018.04.002>
- [25] Stylianos Gisdakis, Thanassis Giannetsos, and Panos Papadimitratos. 2016. Android Privacy C(R)Ache: Reading Your External Storage and Sensors for Fun and Profit. In *Proceedings of the 1st ACM Workshop on Privacy-Aware Mobile Computing (Paderborn, Germany) (PAMCO '16)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/2940343.2940346>
- [26] Google. 2020. Manifest.permission. <https://developer.android.com/reference/android/Manifest.permission>
- [27] Google. 2020. Permissions overview:Android Developers. [https://developer.android.com/guide/topics/permissions/overview#dangerous\\_permissions](https://developer.android.com/guide/topics/permissions/overview#dangerous_permissions)
- [28] Douglas Gunzler, Tian Chen, Pan Wu, and Hui Zhang. 2013. Introduction to mediation analysis with structural equation modeling. *Shanghai archives of psychiatry* 25, 6 (2013), 390.
- [29] Catherine Han, Irwin Reyes,  lvaro Feal, Joel Reardon, Primal Wijesekera, Narseo Vallina-Rodr guez, Amit Elazari Bar On, Kenneth Bamberger, Serge Egelman, et al. 2020. The Price is (Not) Right: Comparing Privacy in Free and Paid Apps. (2020).
- [30] Diarmuid Harkin and Adam Molnar. 2020. Operating-System Design and Its Implications for Victims of Family Violence: The Comparative Threat of Smart Phone Spyware for Android Versus iPhone Users. *Violence Against Women* (2020), 1077801220923731.
- [31] Apple Inc. [n.d.]. App Privacy Details - App Store. <https://developer.apple.com/app-store/app-privacy-details/>
- [32] Hyunjin Kang and Wonsun Shin. 2016. 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 Computing*

- Networking* 19, 3 (2016), 179–185.
- [33] Mark J Keith, Samuel C Thompson, Joanne Hale, Paul Benjamin Lowry, and Chapman Greer. 2013. Information disclosure on mobile devices: Re-examining privacy calculus with actual user behavior. *International journal of human-computer studies* 71, 12 (2013), 1163–1173.
- [34] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3393–3402.
- [35] Nishtha Kesswani and Frank Lin. 2016. How privacy invasive Android apps are?. In *2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom)*. IEEE, 3731–3734.
- [36] Yeolil Kim, Daniel A Briley, and Melissa G Ocepek. 2015. Differential innovation of smartphone and application use by sociodemographics and personality. *Computers in Human Behavior* 44 (2015), 141–147.
- [37] Bart P Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013. Preference-based location sharing: are more privacy options really better?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2667–2676.
- [38] Bart P Knijnenburg, Saadhika Sivakumar, and Darcia Wilkinson. 2016. Recommender systems for self-actualization. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 11–14.
- [39] Bart P Knijnenburg and Martijn C Willemsen. 2015. Evaluating recommender systems with user experiments. In *Recommender Systems Handbook*. Springer, 309–352.
- [40] Konrad Kollnig, Reuben Binns, Pierre Dewitte, Max Van Kleek, Ge Wang, Daniel Omeiza, Helena Webb, and Nigel Shadbolt. 2021. A Fait Accompli? An Empirical Study into the Absence of Consent to Third-Party Tracking in Android Apps. In *Seventeenth Symposium On Usable Privacy and Security (SOUPS) 2021*. 181–195.
- [41] Lydia Kraus, Ina Wechsung, and Sebastian Möller. 2017. Psychological needs as motivators for security and privacy actions on smartphones. *Journal of Information Security and Applications* 34 (2017), 34–45.
- [42] Martin Kuehnhausen and Victor S Frost. 2013. Trusting smartphone apps? To install or not to install, that is the question. In *2013 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE, 30–37.
- [43] Robert S Laufer and Maxine Wolfe. 1977. Privacy as a concept and a social issue: A multidimensional developmental theory. *Journal of social Issues* 33, 3 (1977), 22–42.
- [44] Huoran Li, Xuan Lu, Xuanzhe Liu, Tao Xie, Kaigui Bian, Felix Xiaozhu Lin, Qiaozhu Mei, and Feng Feng. 2015. Characterizing smartphone usage patterns from millions of android users. In *Proceedings of the 2015 Internet Measurement Conference*. 459–472.
- [45] Sampada Marathe, S Shyam Sundar, M Nije Bijvank, Henriette C van Vugt, and Jolanda Veldhuis. 2007. Who are these power users anyway? Building a psychological profile. (2007).
- [46] Scott R Moore, Huangyi Ge, Ninghui Li, and Robert W Proctor. 2019. Cybersecurity for android applications: Permissions in android 5 and 6. *International Journal of Human-Computer Interaction* 35, 7 (2019), 630–640.
- [47] Jack Morse. 2020. As coronavirus spreads, yet another company brags about tracking you. <https://mashable.com/article/coronavirus-location-data-tracking-mobile-phones/>
- [48] Moses Namara, Henry Sloan, Priyanka Jaiswal, and Bart P Knijnenburg. 2018. The Potential for User-Tailored Privacy on Facebook. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC)*. IEEE, 31–42.
- [49] Patricia A Norberg, Daniel R Horne, and David A Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs* 41, 1 (2007), 100–126.
- [50] Xinru Page and Alfred Kobsa. 2009. The circles of latitude: Adoption and usage of location tracking in online social networking. In *2009 International Conference on Computational Science and Engineering*, Vol. 4. IEEE, 1027–1030.
- [51] Xinru Page, Alfred Kobsa, and Bart P Knijnenburg. 2012. Don't disturb my circles! Boundary preservation is at the center of location-sharing concerns. In *Sixth International AAAI Conference on Weblogs and Social Media*.
- [52] Hyanghee Park, Jinsu Eun, and Joonhwan Lee. 2018. Why do smartphone users hesitate to delete unused apps?. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. 174–181.
- [53] Andrew Perrin. 2020. Half of Americans have decided not to use a product or service because of privacy concerns. <https://www.pewresearch.org/fact-tank/2020/04/14/half-of-americans-have-decided-not-to-use-a-product-or-service-because-of-privacy-concerns/>
- [54] Rakesh and Rakesh. 2021. Samsung Bloatware List (2020): Remove Samsung Bloatware Safely. <https://technastic.com/remove-samsung-bloatware-safe-to-remove-apps/>
- [55] Joel Reardon, Álvaro Feal, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodríguez, and Serge Egelman. 2019. 50 Ways to Leak Your Data: An Exploration of Apps' Circumvention of the Android Permissions System. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 603–620.
- [56] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2019. How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1326–1343.
- [57] Lena Reinfelder, Zinaida Benenson, and Freya Gassmann. 2014. Differences between Android and iPhone users in their security and privacy awareness. In *International Conference on Trust, Privacy and Security in Digital Business*. Springer, 156–167.
- [58] Pam Royse, Joon Lee, Baasanjav Undrahbuyan, Mark Hopson, and Mia Consalvo. 2007. Women and games: Technologies of the gendered self. *New media & society* 9, 4 (2007), 555–576.
- [59] Muhammad Irtaza Safi, Abhiditya Jha, Malak Eihab Aly, Xinru Page, Sameer Patil, and Pamela Wisniewski. 2019. Will They Share? Predicting Location Sharing Behaviors of Smartphone Users through Self-Reflection on Past Privacy Behaviors. In *The 2019 NDSS Workshop on Usable Security and Privacy*.
- [60] Noam Segev, Noam Avigdor, and Eytan Avigdor. 2018. Measuring influence on Instagram: a network-oblivious approach. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 1009–1012.
- [61] Christina Shane-Simpson, Adriana Manago, Naomi Gaggi, and Kristen Gillespie-Lynch. 2018. Why do college students prefer Facebook, Twitter, or Instagram? Site affordances, tensions between privacy and self-expression, and implications for social capital. *Computers in Human Behavior* 86 (2018), 276–288.
- [62] Paschal Sheeran and Thomas L Webb. 2016. The intention-behavior gap. *Social and personality psychology compass* 10, 9 (2016), 503–518.
- [63] Drew Smith. 2020. iOS Client Administration. In *Apple macOS and iOS System Administration*. Springer, 109–144.
- [64] Ryan Stevens, Clint Gibler, Jon Crussell, Jeremy Erickson, and Hao Chen. 2012. Investigating user privacy in android ad libraries. In *Workshop on Mobile Security Technologies (MoST)*, Vol. 10. Citeseer.
- [65] S Shyam Sundar and Sampada S Marathe. 2010. Personalization versus customization: The importance of agency, privacy, and power usage. *Human Communication Research* 36, 3 (2010), 298–322.
- [66] Mohsen Tavakol and Reg Dennick. 2011. Making sense of Cronbach's alpha. *International journal of medical education* 2 (2011), 53.
- [67] Hsiu-Yu Wang, Chechen Liao, and Ling-Hui Yang. 2013. What affects mobile application use? The roles of consumption values. *International Journal of Marketing Studies* 5, 2 (2013), 11.
- [68] Na Wang, Pamela Wisniewski, Heng Xu, and Jens Grossklags. 2014. Designing the default privacy settings for facebook applications. In *Proceedings of the companion publication of the 17th ACM conference on Computer supported cooperative work & social computing*. 249–252.
- [69] Na Wang, Bo Zhang, Bin Liu, and Hongxia Jin. 2015. Investigating effects of control and ads awareness on android users' privacy behaviors and perceptions. In *Proceedings of the 17th international conference on human-computer interaction with mobile devices and services*. 373–382.
- [70] Yang Wang, Jun Zheng, Chen Sun, and Srinivas Mukkamala. 2013. Quantitative security risk assessment of android permissions and applications. In *IFIP Annual Conference on Data and Applications Security and Privacy*. Springer, 226–241.
- [71] Paul R Warshaw and Fred D Davis. 1985. Disentangling behavioral intention and behavioral expectation. *Journal of experimental social psychology* 21, 3 (1985), 213–228.
- [72] Rick Wash and Emilee Rader. 2015. Too much knowledge? security beliefs and protective behaviors among united states internet users. In *Eleventh Symposium On Usable Privacy and Security ({SOUPS} 2015)*. 309–325.
- [73] Xuetao Wei, Lorenzo Gomez, Iulian Neamtii, and Michalis Faloutsos. 2012. Permission evolution in the android ecosystem. In *Proceedings of the 28th Annual Computer Security Applications Conference*. 31–40.
- [74] Pamela Wisniewski, AKM Najmul Islam, Bart P Knijnenburg, and Sameer Patil. 2015. Give social network users the privacy they want. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 1427–1441.
- [75] Pamela Wisniewski, Jessica Vitak, Xinru Page, Bart Knijnenburg, Yang Wang, and Casey Fiesler. 2017. In whose best interest? Exploring the real, potential, and imagined ethical concerns in privacy-focused agenda. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 377–382.
- [76] Pamela J Wisniewski, Bart P Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies* 98 (2017), 95–108.
- [77] Heng Xu, Sumeet Gupta, Mary Beth Rosson, and John M Carroll. 2012. Measuring mobile users' concerns for information privacy. (2012).
- [78] Heng Xu and Hock-Hai Teo. 2004. Alleviating consumers' privacy concerns in location-based services: a psychological control perspective. *ICIS 2004 proceedings* (2004), 64.
- [79] Sha Zhao, Julian Ramos, Jianrong Tao, Ziwen Jiang, Shijian Li, Zhaohui Wu, Gang Pan, and Anind K Dey. 2016. Discovering different kinds of smartphone

users through their application usage behaviors. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 498–509. [80] Bu Zhong. 2013. From smartphones to iPad: Power users' disposition toward mobile media devices. *Computers in human behavior* 29, 4 (2013), 1742–1748.

## A THE PERCEIVED SURVEY MEASURES USING PRE-VALIDATED SCALES:

### A.1 Power Use Scale (Adapted from [45])

[Measured on a 5-point Likert Scale: Strongly Disagree - Strongly Agree]

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

### A.2 Behavioral Intention - Install Mobile Apps (Both items adapted from [78])

[Measured on a 5-point Likert Scale: 1- Agree Strongly, 2 - Agree Somewhat, 3 - Neutral, 4 - Disagree Somewhat, 5 - Disagree Strongly]

- I predict I will use new mobile apps in the next 3 months.
- I intend to use mobile apps in the next 3 months.

### A.3 Behavioral Intention - Share Information with Apps (Adapted from [78])

[Measured on a 5-point Likert Scale: 1- Agree Strongly, 2 - Agree Somewhat, 3 - Neutral, 4 - Disagree Somewhat, 5 - Disagree Strongly]

- I am likely to disclose my personal information to use mobile apps in the next 3 months.
- I am likely to grant permission to share [my location] with my existing mobile apps in the next 3 months.
- I am likely to grant permission to share [my location] with new mobile apps in the next 3 months

## B DANGEROUS PERMISSION REQUESTS (N = 6688 INSTALLED APPLICATIONS)

Permission Group	Dangerous Permission	App Request (%)	Grant (%)
STORAGE	WRITE_EXTERNAL_STORAGE	73.0	100
	READ_EXTERNAL_STORAGE	1.1	81.3
LOCATION	ACCESS_FINE_LOCATION	35.0	100
	ACCESS_COARSE_LOCATION	32.0	99.7
CAMERA	CAMERA	29.0	99.73
PHONE	READ_PHONE_STATE	29.0	100
	GET_ACCOUNTS	24.0	100
	CALL_PHONE	8.0	97.6
	ADD_VOICEMAIL	2.0	59.2
	PROCESS_OUTGOING_CALLS	0.0	8.95
	USE_SIP	0.0	6.3
CONTACTS	READ_CONTACTS	14.0	99.5
	WRITE_CONTACTS	4.0	96.3
MICROPHONE	RECORD_AUDIO	14.0	99.7
SMS	READ_SMS	4.0	95.8
	RECEIVE_WAP_PUSH	4.0	96.3
	RECEIVE_SMS	4.0	95.8
	SEND_SMS	0.0	28.4
	RECEIVE_MMS	0.0	28.42
CALL_LOG	READ_CALL_LOG	3.0	87.9
CALENDAR	READ_CALENDAR	4.0	92.1
	WRITE_CALENDAR	4.0	86.6
SENSORS	BODY_SENSORS	2.0	83.2

Table 4: The percentage requests for Dangerous Permissions by the Installed Apps

## C THE TOP 10 COMMON APPLICATIONS PARTICIPANTS' HAD INSTALLED

Application Name	Package Name	Installations	Approx. %
Messenger	com.facebook.orca	228	60.00
Facebook	com.facebook.katana	223	58.68
Instagram	com.instagram.android	193	50.79
Netflix	com.netflix.mediaclient	146	38.42
Spotify	com.spotify.music	128	33.68
Snapchat	com.snapchat.android	124	32.63
Uber	com.ubercab	119	31.31
eBay	com.ebay.mobile	112	29.47
PayPal	com.paypal.android.p2pmobile	110	28.94
Pinterest	com.pinterest	95	25.00

Table 5: The top ten most common installed applications among our participants