

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law. Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing [scholarworks-group@umbc.edu](mailto:scholarworks-group@umbc.edu) and telling us what having access to this work means to you and why it's important to you. Thank you.

# Towards Understanding Connections between Security/Privacy Attitudes and Unlock Authentication\*

Adam J. Aviv  
United States Naval Academy  
aviv@usna.edu

Ravi Kuber  
University of Maryland, Baltimore County  
rkuber@umbc.edu

February 2, 2018

## Abstract

In this study, we examine the ways in which user attitudes towards privacy and security relating to mobile devices and the data stored thereon may impact the strength of unlock authentication, focusing on Android's graphical unlock patterns. We conducted an online study with Amazon Mechanical Turk ( $N = 750$ ) using self-reported unlock authentication choices, as well as Likert scale agreement/disagreement responses to a set of seven privacy/security prompts. We then analyzed the responses in multiple dimensions, including a straight average of the Likert responses as well as using Principle Component Analysis to expose latent factors. We found that responses to two of the seven questions proved relevant and significant. These two questions considered attitudes towards general concern for data stored on mobile devices, and attitudes towards concerns for unauthorized access by *known* actors. Unfortunately, larger conclusions cannot be drawn on the efficacy of the broader set of questions for exposing connections between unlock authentication strength (Pearson Rank  $r = -0.08$ ,  $p < 0.1$ ). However, both of our factor solutions exposed differences in responses for demographics groups, including age, gender, and residence type. The findings of this study suggests that there is likely a link between perceptions of privacy/security on mobile devices and the perceived threats therein, but more research is needed, particularly on developing better survey and measurement techniques of privacy/security attitudes that relate to mobile devices specifically.

---

\*Article to appear at the Workshop on Usable Security (USEC), Feb. 2018

# 1 Introduction

In order to protect sensitive data stored on or accessed via a mobile device, such as a smartphone, knowledge-based authentication methods (e.g. use of PINs, stroke-based patterns) or biometric methods (e.g. fingerprints) can be used to verify the identity of the user, thereby enabling the device to be unlocked. Research suggests that attitudes to securing mobile devices are known to differ among users [19, 17, 14, 18, 33, 1, 34]. Although locking behavior is often practiced, annoyances have been found with the process [19, 1, 34]. While not averse to locking a mobile device, Egelman et al. [14] identified lack of motivation as one factor deterring participants in their study from securing their technologies. Calculated decisions were made by some participants, depending on the perceived effort required to lock devices. Deterrents have also been attributed in part to the absence of threats to security, along with negative perceptions of the security of lock screens [17].

For purposes of unlocking a mobile device, many Android users have the option of using stroke-based patterns which can be entered via the interface. The user is required to connect a sub-set of four or more pre-selected contact points arranged in a 3x3 grid in the correct sequence to authenticate entry. However, studies have shown that in practice, only a small number of distinct patterns are selected from the 389,112 choices available [6]. The majority of these user-selected patterns can be easily guessed with roughly the same difficulty as guessing random 3-digit PINs [4, 39]. While the threat of observer and smudge attacks exist when using a locking mechanism, the benefits often outweigh the impact of losing sensitive data for potential abuse by third parties. Prior work in the area has often focused on the factors impacting locking compliance or lack thereof [19, 17, 14, 18]. Although limited in number, studies have been conducted examining security and privacy attitudes associated with locking behavior [1, 34]. However, research has yet to be undertaken to directly connect security/privacy attitudes with passcode strength.

In this paper, we describe a study examining the ways in which users select authentication stimuli, notably Android’s graphical password system or “pattern unlock” based on their attitudes towards a set of security and privacy prompts concerning information privacy on mobile devices. We collected and analyzed a large dataset ( $n = 750$ ) of self-reported Android unlock patterns/statistics regarding pattern choice and correlated that data with agreement/disagreement responses to a set of privacy prompts related to the privacy of data stored on mobile devices.

The goals of this study are twofold: first we hypothesized that there is a relationship between privacy attitudes regarding mobile devices and the security of the authentication stimuli selected (H1), and second we hypothesized that there may be demographic factors (e.g., age, location of residence) that may affect security and privacy responses (H2). As our questionnaire is short in design, the hope is that these prompts can be used as starting point to design more advanced measurement techniques, or more directly, a subset of these questions could be used at device set-up time (or during a security audit) to better understand user motivation and potentially nudge users towards secure choices.

We conducted per-prompt analysis, an analysis of averages of the scores, and used principal component analysis (PCA) to expose any latent factors in the data. The PCA analysis explained 55% of the variance for a two factor solution with high internal

consistency for each factor solution (Cronbach's alpha  $> 0.7$ ) and low correlation (likely independence) between the factors (Spearman Rank  $\rho = 0.09$ ,  $p < 0.05$ ). The two different latent factors identified different aspects of participant response, particularly around aspects of *present vs. external* threats to security/privacy. A similar dichotomy was found when identifying connections to unlock authentication strength.

The overall results, using the entirety of the questions to predict users with weaker unlock strengths, are unfortunately *inconclusive* as the correlation effect is small (Pearson Rank  $r = -0.08$ ,  $p < 0.1$ ) with a weak effective size. We argue, though, that the results are encouraging for the potential of such a link between security/privacy attitude and unlock authentication may be present, if more advanced measurement methods were developed.

The results of the survey do provide a path forward in identifying better honed measurement methods when analyzing the response to individual privacy/security prompts. Two such prompts had statistically significant differences with regard to the reported strength of the authentication patterns. In particular, an agreement with general concern for privacy of data stored on mobile devices and concern of unauthorized access from known actors (e.g., friends and family) led to stronger authentication choices ( $p < 0.05$ ).

In analyzing other relationships between the participants and the privacy prompts, we show that our two factor solutions found using PCA exposes interesting demographic differences in our responses. In particular, Factor 1 solution correlated more with the strength of the reported results, our second factor solution, Factor 2, better separated the respondent's attitude with respect to their age and willingness to provide a pattern or statistics about those patterns for the experiment.

To summarize, the paper makes the following contributions and findings:

- An analysis of the efficacy of a short privacy/security questionnaire regarding threats to mobile devices.
- Exposing latent factors via PCA that describe clear difference in participants between present and external threats to mobile device security.
- Demonstrating correlations between latent factors of the survey with demographic information, particularly age and location.
- Identifying the connection of individual questions to unlock authentication strength, but inconclusive results for the entirety of the prompts.

Based on these results, we conclude that while we suspect that a relationship between authentication choice and security/privacy attitudes of mobile devices seems probable, we were not able to find enough evidence here to draw strong conclusions. Further study is needed, and we show a number of promising directions regarding measurement methods that could be applied to expose this link. In particular, there are encouraging findings with respect to specific prompts, as well as when applying factor analysis to the entire set of privacy prompts. This suggests that using a short survey to better understand user perceptions regarding privacy on mobile devices could be possible and would be a useful tool in directing users in a personalized way, based upon their own security and privacy perceptions, towards more secure choices.

## 2 Related Work

Researchers have aimed to optimize both security and usability for mobile unlocking mechanisms, through extensive studies examining both attitudes and usage behaviors relating to securing mobile technologies. In terms of attitudes, mobile device users have been found to express concerns about security [27] and make clear demands for protection against unauthorized use [13]. However, difficulties have been experienced when complying with general security guidance, such as updating passcodes [13], even if users are ‘tech savvy’ [27]. Research suggests that a number of mobile device users prefer not to lock their devices [14]. Reasons for this include concerns about emergency personnel not being able to identify them, not having their devices returned if lost, and not believing they had any data worth protecting. Lack of motivation and awareness were also found to be reasons behind lack of compliance in a study by Qiu et al. [36]. The usability challenges (i.e., the time taken to use locking screens) were identified by Malkin et al. [33] as deterrents for locking. When examining perceptions in more detail, Egelman et al. [14] identified a strong correlation between use of security features and risk perceptions, which indicates rational behavior. The researchers also observed that most users likely underestimate the extent to which data stored on their smartphones pervades their identities, online and offline. Hayashi et al. [20] found that most of the participants in their study wanted at least half of their applications to be accessible without requiring an unlock code, which according to Egelman et al. [14], suggests opportunities for improving current locking mechanisms.

Studies have examined a range of factors which may play a role in impacting privacy and security attitudes. For example, Benensen et al. [7] conducted a study to identify whether mobile device users utilizing different operating systems had varying perceptions of security. Findings from their study showed that 20% of Android users and 5% of iOS users were privacy aware. Mare et al. [34] examined the impact of age, among other factors, on attitudes to privacy and security. The researchers identified that teenage users and not just adults were interested in these areas. However, teenage users were thought to have less useful understandings of how to achieve it.

Studies have been conducted to investigate locking behaviors by users from eight countries (e.g. [18, 33]). Findings revealed that level of protection varies by location. Participants in most non-U.S. countries were between 31% and 76% more likely than Americans to have a secure lock screen. Age and sex were also thought to play a role in determining locking behavior. Reasons for lack of compliance appeared to be similar across participants from most countries examined [18]. Egelman and Peer [15] examined the impact of individual differences on privacy and security attitudes. The five factor model, often used to examine the impact of these differences on decision making, was found to be a weak predictor of privacy preferences and behaviors. The researchers identified that decision-making style and risk-taking attitudes are strong predictors of privacy attitude.

Qiu et al. [36] have focused on recording and analyzing locking behaviors among mobile device users. These include the length of unlocking and interaction sessions with the devices, outcomes (failed/succeeded) of the unlocking attempts, type of locking mechanism (e.g. manually or through auto-lock), participants’ attitudes of sharing their devices with other people, device locations and information gathered from ac-

<i>n</i> = 750	<i>Reported Statistics</i>	<i>Reported Pattern</i>	<i>Urban</i>	<i>Suburban</i>	<i>Rural</i>	<i>Male</i>	<i>Female</i>
Participants	307	443	278	330	142	407	343
18-24	58	148	83	90	33	123	83
25-34	134	228	143	154	65	199	163
35-44	85	52	41	66	30	67	70
45-54	24	13	8	18	11	17	20
55-64	5	2	3	2	2	1	6
65+	1	0	0	0	1	0	1
Urban	119	159				170	108
Suburban	133	197				173	157
Rural	55	87				64	78
Male	153	254	170	173	64		
Female	154	189	108	157	78		

Table 1: Participant demographics for those that only reported statistics on their authentication pattern and those that reported their authentication pattern.

celerometers/gyroscopes, the model of the device and its screen size. Although the study has yet to be completed, the findings would offer interesting insights to interface designers aiming to develop security solutions to better meet the needs of mobile device users. Hintze et al. [21] also examined usage behavior when using mobile devices. Findings showed that a noteworthy amount of users do not lock their devices after usage, increasing the possibility of malicious interaction with the device when considering unlocked usage, as devices stay unlocked for a short period afterwards.

In terms of challenging perceptions of risk, concerns and attitudes to screen locking, Albayram et al. [1] presented a video to participants based on 'fear appeal' highlighting the potential consequences associated with poor locking compliance. The video was found to significantly affect levels of perceived severity, vulnerability, and response efficacy. The researchers suggest that risk communication can effectively change risk perception, which can be a key to promoting secure behavior such as the use of a secure screen lock mechanism on smartphones.

### 3 Methodology

The goal of this research was to determine if security and privacy attitudes towards data stored on mobile devices impact the strength of chosen passcodes. Our primary result involved the analysis of participants' agreement/disagreement with statements regarding information security/privacy of mobile phones (on a Likert scale). This type of approach has been used by other researchers in prior work [11, 28, 31, 8, 12]). Some of those participants also self-reported their mobile device unlock passcode (an Android graphical pattern) in addition to reporting their agreement/disagreement with the privacy statements, allowing us to compare their privacy sentiment with regards to their mobile device and the relative strength of their passcodes.

To measure the strength of the reported passcodes, we use password guessability

metrics [9, 35, 23] and particularly the guessing methods developed in related work [4, 39] trained using non-overlapping data of user-generated passcodes obtained from the same related work [39, 30, 40]. All the data collection methods were reviewed and approved by our institutions' IRBs.

### 3.1 Self Reporting Passcodes

This study builds upon the analysis and data collection conducted previously by Aviv et al. [4, 5]. The study was conducted on Amazon Mechanical Turk and participants who self-identified as Android mobile device users were recruited with the following prompt:

You will be asked to (optionally) self-report what your Android unlock pattern is, or report statistics about your pattern. If you do not use the Android unlock pattern, please do not complete the HIT.

Participants entering the survey were aware of the nature of the study and what was going to be asked of them. This may have provided a bias to the results as those *opting-in* may have been less privacy conscious than those *opting-out*; however, the option to *not* report a pattern and still complete the survey allowed for more privacy conscious users to participate. The statistics that participants reported about their patterns included the start contact point and whether any of a pre-selected set of tri-grams were present, i.e., a connection of three contact points, but were not required to reveal their precise pattern. Both those who reported their actual pattern or statistics were compensated equally, \$0.75, for their participation and fully completing the survey.

Once the HIT was accepted, the structure of the survey was as follows:

1. Informed Consent: Participants were made aware of the intent of the survey, the data collection and protection procedures, and the kinds of information that would be collected.
2. Demographic Information: Participants answered background questions about their age, gender, and place of residence, either rural, suburban, or urban. Participants were limited to U.S. areas using Amazon Mechanical Turk's settings.
3. Pattern/Stats: Participants were given the option to either report statistics about their Android unlock pattern or report their actual pattern.
4. Privacy Prompts: Participants were asked to rate their agreement/disagreement with the privacy prompts (see below for specific prompts).
5. Confirm/Attention Tests: At the end of the survey, participants were asked to re-enter their pattern (or stats) and answer a question about providing truthful responses. Participant data was removed if the participant failed to confirm the same pattern/stats or answered negatively to the truthfulness question. Seven participants were excluded for this reason.

Further, participants were required to complete the entirety of the survey on an Android device. The survey, including pattern entry screens were written using a

combination of JavaScript, PHP, and HTML, and run completely within a browser on a mobile device. Once the HIT was accepted, the participants were provided with a link to visit on their Android device's browser window. To ensure that participants used an Android device, the `user-agent` header was checked. While the `user-agent` header can be spoofed, it is likely more burden than it is worth given the compensation for the survey of \$0.75. Participants who did not own an Android device should have returned the HIT in order to complete other tasks.

The demographics of the study are presented in Table 1, and in total 750 individuals participated in the study, 443 of which also self-reported their authentication passcode for their Android unlock pattern.

### 3.2 Security/Privacy Prompts

Each participant, regardless of if they self-reported a unlock pattern, were provided with the following prompts/questions which they responded disagreement/agreement with on the following scale: "Strongly Disagree," "Disagree," "Neither Agree nor Disagree," "Agree," or "Strongly Agree."

- **Q1:** I am concerned about the privacy of the data stored on my mobile device(s).
- **Q2:** I am more concerned about the privacy of the information stored on my mobile device as compared to my laptop or desktop.
- **Q3:** I am more concerned about the privacy of the information stored on my mobile device as compared to information stored online.
- **Q4:** I am concerned about identity theft.
- **Q5:** I am concerned about my privacy online.
- **Q6:** I am concerned about unauthorized access to my device by known actors (people I know, for example, friends, family, and my peers).
- **Q7:** I am concerned about unauthorized access to my device by unknown actors (people I don't know, for example, police, government, and corporations).

Based on the responses, a security/privacy attitude average can be calculated by assigning a value of 1 to "Strongly Disagree," 2 to "Disagree," 3 to "Neither Agree nor Disagree", 4 to "Agree", and 5 to "Strongly Agree" that related to the general privacy attitude of a participant as it relates to information stored on mobile devices. For those participants that also self-reported their authentication passcode, we next analyzed the strength of that passcode as it related to their security/privacy attitudes.

While other published scales have been used by researchers to determine perceptions/attitudes towards security/privacy, such as the ones described by Buchanan et al. [10], Malhotra et al. [32], Kumaraguru and Cranor [25]-(Westin Primary Segmentation index), and others, some parts of these scales are out of date, not focusing on mobile devices (and aspects of unlock authentication) or were deemed to be prohibitively long for a participant to take on a mobile device. Instead, we developed our prompts from prior published privacy scales for computing, generally [11, 28]). We do not imply






Figure 1: Magnitude of responses to each question prompt for those that reported statistics and those that reported patterns, where: SD is “Strongly Disagree,” D is “Disagree,” N is “Neither Disagree nor Agree,” A is “Agree,” and SA is “Strongly Agree.”

that these prompts provide complete clarity of security/privacy attitudes. However, they do expose some key aspects of the mental models of users. For example, the difference between security/privacy of data stored on the device vs. security/privacy of data stored online, as well as the difference between unauthorized access by unknown actors and known actors. Further, as we will discuss in the next section, we found acceptable internal consistency with two different factors, similar to that of other published security/privacy questionnaires.

## 4 Identifying Latent Factors

As a first step in the analysis, consider the results in Figure 1 for a general overview of the responses to each of the privacy prompts. Clearly observed is the fact that there are varied distributions in the responses, by which some of the participants were strongly in agreement for most of the questions, while other questions provided more varied responses. The variance of the responses is not ideal as some questions have mostly agreement or strong-agreement (Q1, Q4, Q5, and Q7), while others show more spread variance (Q2, Q3, and Q6). To better identify the impact of all the prompts, i.e., to find

	F1	F2
Q1	-0.411	-0.245
Q2	-0.435	
Q3	-0.401	-0.114
Q4	-0.290	-0.411
Q5	-0.282	-0.255
Q6	-0.458	0.812
Q7	-0.328	-0.176

Table 2: Weighting for the two factors found via PCA where 1 is “Strongly Disagree” and 5 is “Strongly Agree”. The Cronbach  $\alpha > 0.7$  (acceptable) for these values

an appropriate weighting, we can apply a naive method or a more advanced factor-based analysis.

The naive method for computing a quantitative privacy response into a continuous variable for each participant is to simply take the average of each of the responses to each of the questions, on a scale from 1 to 5, which we refer to as the *privacy average*. A privacy average across the prompts of 5 would indicate strong agreement with all the prompts, and a 1 would indicate strong disagreement. Values in-between would indicate a more nuanced perspective, but likely shifted towards more agreement or less agreement with the prompts. We apply this quantization in much of analysis and justify it through principal component analysis to follow.

As a more advanced technique for determining appropriate weighting of the questions in computing a continuous value, we use principal component analysis (PCA) to identify latent factors in the responses and the associated question weights (loading) that lead to high internal consistency in responses between questions. This is a common technique for verifying privacy scales, and we follow the verification technique in prior published work, particularly that of Buchanan et al. [10].

Just as in Buchanan et al. [10], we required a loading greater than 0.3 for a Varimax analysis, which left us with a three-factor solution. Further we used Saucier’s criteria and then Cronbach’s  $\alpha > 0.7$  to determine factor purity, after which we were left with two viable factors. The loadings of each of the factors are presented in Table 2. Factor 1 (F1) had a Cronbach alpha of 0.74 and Factor 2 (F2) had a Cronbach alpha of 0.71. Comparing the two-factor solutions, we find that they are weakly correlated. Figure 2 plots the factors in 2D space, and shows the weak negative correlation (Spearman rank of  $\rho = -0.09$ ,  $p < 0.05$ ), suggesting that each factor is exposing different latent factors of the. As  $\rho$  is quite small, these two factors are exposing different aspects of the responses, and overall, 55% of the variance is explained by the two factor solution.

Observing the loadings, F1 is very similar to the privacy average metric described previously because each of the questions are loaded at nearly equal weight, but the loading is all negative. Since the contribution of each of the questions is significant and none are significantly greater than any other, it suggests that using an average of the results provides a reasonable and straightforward quantitative measure of privacy attitudes based on the question prompts.

The F2 solution provides a very different loading. Question 2, which compares the privacy of information stored on the mobile device compared to a laptop or desktop,




Figure 2: A scatter plot of two factor solution correlation. The two factors are weakly negatively correlated based on Spearman’s rank correlation coefficient ( $\rho = -0.088$ ,  $p < 0.05$ ).

does not contribute to the solution based on the minimization criteria requiring high internal consistency. While Q2 contributes highly in F1, it does not provide any loading in F2. Instead, Q6 and Q4 have the largest impact, one positive and one negative, respectively.

These two questions, Q6 and Q4, highlight one of the big dividers we found in this survey: *present vs. distant threats*. Q4 refers to the risk of identity threat, which mostly occurs from unknown actors, while Q6 refers to the threat of unauthorized access by familiar, known actors, like friends and family. When the external threat factor is combined with the other negatively loaded factors, which all deal with external threats to data security and privacy, it is clear that F2 is picking up on differences between participants fears of known vs. unknown actors. As we will show, F2 provides significant differences in results when considering the willingness of demographic factors, particularly surrounding age.

## 5 Guessability Metric

For the strength metric of Android unlock patterns, we use guessability [9, 35, 23]. Guessability is a metric that describes how many guesses would an informed attacker require to guess a user’s password in an offline setting.

The guessability metric requires the creation of a guessing algorithm which outputs

a series of guesses in order of most likely used passwords to least likely used passwords. Prior work in this space has provided a framework for designing these algorithms [4, 39] which require training data. One advantage present for studying Android unlock patterns is that, unlike text-based passwords where there are essentially an unbound number of passwords, for Android unlock patterns there are finite number of patterns, exactly 389,112 [6]. Thus, it is possible to develop a complete ranking of patterns using a guesser.

The best known guessing algorithms [4, 39] leverage the properties of commonly selected patterns, such as the **Z**- and **M**-shaped patterns (and rotations/flips thereof), to construct a Markov model based on transitions within the pattern grid space. Recall that Android unlock patterns require connecting a set of 3x3 grid points in a continuous gesture without repetition or avoidance. The order in which those grid points connect is well modeled using a bi-gram transition matrix of likelihoods, where the likelihoods are measured based on training data. Additional probabilities considering the start point, end point and the length of the pattern can also be modeled using similar measures. Once the model is constructed, all 389,112 possible patterns (enumerated using brute force, as in Aviv et. al [6]) are analyzed using the Markov model, producing a likelihood score. Ranking the patterns based on this likelihood score provides a guessability metric where lower values indicate patterns that are more likely to be guessed and higher values indicate patterns that are less likely to be guessed.

We followed the Markov model construction from prior work by Aviv et al. [4] (particularly the pen-and-paper data). In addition to training the model using the Android unlock patterns used in prior collections, we also trained the Markov model on patterns provided via Uellenbeck et al. [39]. As a limitation, these models change over time as more training data is available; however, repeated observations of user choice in Android unlock patterns find that the set of common patterns appears to be consistent [2, 37, 41, 30, 40].

## 6 Results

### 6.1 Demographic Differences in Responses

Figure 1 provides the general overview of the magnitude of responses for participants to the security/privacy prompts. The figure separates by two groups: those willing to self-report their authentication passcode (or pattern) and those unwilling. The willingness to reveal this information could also be considered a sensitivity regarding security and privacy, and we present the statistical analysis in Table 3.

An omnibus MANOVA test showed no significant differences between these groups. Additionally, the Factor 1 (F1) solution also showed no statistical differences between the responses of participants who were or were not willing to provide a pattern for analysis. Interestingly, the Factor 2 (F2) solution does show a significant difference that is not present in prior tests. This result suggests that F2 is capturing some more difference between the willingness of those to share information about their authenticator than what is simply captured under an average of the responses.

While we found no significant differences with respect to gender (see Table 4), we

	Reported Statics	Reported Pattern	<i>t</i> -test
Q1	$\mu = 3.90, \sigma = 0.97$	$\mu = 4.01, \sigma = 0.94$	$p = 0.135$
Q2	$\mu = 3.25, \sigma = 1.22$	$\mu = 3.34, \sigma = 1.14$	$p = 0.319$
Q3	$\mu = 2.95, \sigma = 1.14$	$\mu = 3.04, \sigma = 1.11$	$p = 0.284$
Q4	$\mu = 3.90, \sigma = 0.98$	$\mu = 4.00, \sigma = 0.87$	$p = 0.135$
Q5	$\mu = 4.05, \sigma = 0.83$	$\mu = 4.12, \sigma = 0.76$	$p = 0.266$
Q6	$\mu = 3.33, \sigma = 1.27$	$\mu = 3.22, \sigma = 1.18$	$p = 0.229$
Q7	$\mu = 4.14, \sigma = 0.90$	$\mu = 4.04, \sigma = 0.99$	$p = 0.137$
avg	$\mu = 3.65, \sigma = 0.65$	$\mu = 3.68, \sigma = 0.64$	$p = 0.474$
F1	$\mu = -9.36, \sigma = 1.75$	$\mu = -9.44, \sigma = 1.70$	$p = 0.521$
F2	$\mu = -1.95, \sigma = 1.08$	$\mu = -2.12, \sigma = 1.00$	$*p = 0.033$

Table 3: Security/Privacy average per question/prompt comparing those that reported statistics about their authentication to those that self-reported their actual authentication pattern. Data was normal ( $p < 0.01$ ) according to Anderson-Darling test. Performed MANOVA tests, differences across responses were *not* significant ( $Z^* = 14.07, \alpha = 0.056$ ).

	Male	Female	<i>t</i> -test
Q1	$\mu = 3.97, \sigma = 1.00$	$\mu = 3.92, \sigma = 0.91$	$p = 0.460$
Q2	$\mu = 3.23, \sigma = 1.20$	$\mu = 3.35, \sigma = 1.17$	$p = 0.163$
Q3	$\mu = 2.98, \sigma = 1.16$	$\mu = 2.99, \sigma = 1.10$	$p = 0.951$
Q4	$\mu = 3.90, \sigma = 0.97$	$\mu = 3.98, \sigma = 0.88$	$p = 0.239$
Q5	$\mu = 4.08, \sigma = 0.82$	$\mu = 4.08, \sigma = 0.78$	$p = 0.993$
Q6	$\mu = 3.31, \sigma = 1.22$	$\mu = 3.26, \sigma = 1.25$	$p = 0.540$
Q7	$\mu = 4.15, \sigma = 0.90$	$\mu = 4.03, \sigma = 0.98$	$p = 0.069$
avg	$\mu = 3.66, \sigma = 0.63$	$\mu = 3.66, \sigma = 0.66$	$p = 0.940$
F1	$\mu = -9.39, \sigma = 1.69$	$\mu = -9.39, \sigma = 1.77$	$p = 0.941$
F2	$\mu = -2.00, \sigma = 1.07$	$\mu = -2.04, \sigma = 1.04$	$p = 0.569$

Table 4: Security/Privacy average per question/prompt comparing male and female responses. Data was normal ( $p < 0.01$ ) according to Anderson-Darling test. Performed MANOVA tests, differences across responses were *not* significant ( $Z^* = 14.07, \alpha = 0.093$ ).

did observe differences within other demographic groupings. The reported location of participants (Urban, Sub-Urban, vs. Rural) showed significant differences for the factor results and under MANOVA tests ( $p = 0.012$ ); post-hoc, one-vs-all analysis suggests that Urban respondents differed significantly compared to other groupings (see Table 5). This is also captured using the F1 weighting and using the average of responses, as expected given the generally smooth weighting of F1. In analyzing the questions, using a Bonferoni correction, Q2 (more concern about security/privacy on mobile vs. desktop/laptop) and Q6 (more concern about known actors vs. unknown actors) showed the most differences for these groups. While we exclude Q3 (more concern about privacy on mobile vs. online) due to the correction, it is worth noting that it does seem to have some impact on differentiating respondents and is highly weighted under F1, so is considered well in that comparison.

	Urban	Sub-Urban	Rural	ANOVA
Q1	$\mu = 4.04, \sigma = 1.00$	$\mu = 3.92, \sigma = 0.91$	$\mu = 3.85, \sigma = 0.99$	$p = 0.114$
Q2	$\mu = 3.37, \sigma = 1.18$	$\mu = 3.15, \sigma = 1.18$	$\mu = 3.42, \sigma = 1.20$	* $p = 0.024$
Q3	$\mu = 3.07, \sigma = 1.20$	$\mu = 2.86, \sigma = 1.05$	$\mu = 3.09, \sigma = 1.15$	$p = 0.032$
Q4	$\mu = 3.94, \sigma = 1.01$	$\mu = 3.94, \sigma = 0.85$	$\mu = 3.94, \sigma = 0.96$	$p = 1.000$
Q5	$\mu = 4.13, \sigma = 0.84$	$\mu = 4.04, \sigma = 0.74$	$\mu = 4.09, \sigma = 0.86$	$p = 0.411$
Q6	$\mu = 3.43, \sigma = 1.23$	$\mu = 3.15, \sigma = 1.22$	$\mu = 3.32, \sigma = 1.23$	* $p = 0.017$
Q7	$\mu = 4.13, \sigma = 0.97$	$\mu = 4.12, \sigma = 0.86$	$\mu = 3.98, \sigma = 1.05$	$p = 0.249$
avg	$\mu = 3.73, \sigma = 0.70$	$\mu = 3.60, \sigma = 0.56$	$\mu = 3.67, \sigma = 0.71$	* $p = 0.04$
F1	$\mu = -9.59, \sigma = 1.86$	$\mu = -9.20, \sigma = 1.51$	$\mu = -9.43, \sigma = 1.89$	* $p = 0.02$
F2	$\mu = -1.95, \sigma = 1.07$	$\mu = -2.10, \sigma = 1.04$	$\mu = -1.96, \sigma = 1.05$	$p = 0.09$

Table 5: Security/Privacy average per question/prompt comparing locales. Data was normal ( $p < 0.01$ ) according to Anderson-Darling test. Using MANOVA indicates that there are significant differences across responses ( $Z^* = 23.68, p = 0.012$ ), and post-hoc analysis (one vs. all) suggests that responses from Urban participants were significantly different. Using a Bonferoni correction, Q2 and Q6 show significant different responses within questions. F1 comparisons also suggest significant differences.

	18-24	25-34	35-44	45-54	ANOVA
Q1	$\mu = 3.89, \sigma = 1.02$	$\mu = 3.91, \sigma = 0.99$	$\mu = 4.05, \sigma = 0.83$	$\mu = 4.19, \sigma = 0.65$	$p = 0.407$
Q2	$\mu = 3.21, \sigma = 1.25$	$\mu = 3.30, \sigma = 1.18$	$\mu = 3.45, \sigma = 1.13$	$\mu = 3.08, \sigma = 1.07$	$p = 0.302$
Q3	$\mu = 2.88, \sigma = 1.19$	$\mu = 2.92, \sigma = 1.13$	$\mu = 3.26, \sigma = 1.01$	$\mu = 3.24, \sigma = 1.02$	* $p = 0.017$
Q4	$\mu = 3.87, \sigma = 1.03$	$\mu = 3.96, \sigma = 0.94$	$\mu = 3.97, \sigma = 0.78$	$\mu = 4.03, \sigma = 0.75$	$p = 0.814$
Q5	$\mu = 4.01, \sigma = 0.90$	$\mu = 4.09, \sigma = 0.80$	$\mu = 4.15, \sigma = 0.65$	$\mu = 4.14, \sigma = 0.58$	$p = 0.439$
Q6	$\mu = 3.51, \sigma = 1.24$	$\mu = 3.25, \sigma = 1.25$	$\mu = 3.23, \sigma = 1.15$	$\mu = 2.84, \sigma = 1.05$	* $p = 0.003$
Q7	$\mu = 4.22, \sigma = 0.87$	$\mu = 4.02, \sigma = 1.00$	$\mu = 4.14, \sigma = 0.86$	$\mu = 4.11, \sigma = 0.86$	$p = 0.080$
avg	$\mu = 3.66, \sigma = 0.69$	$\mu = 3.64, \sigma = 0.65$	$\mu = 3.75, \sigma = 0.56$	$\mu = 3.66, \sigma = 0.50$	$p = 0.385$
F1	$\mu = -9.40, \sigma = 1.85$	$\mu = -9.33, \sigma = 1.74$	$\mu = -9.63, \sigma = 1.50$	$\mu = -9.34, \sigma = 1.37$	$p = 0.477$
F2	$\mu = -1.79, \sigma = 0.96$	$\mu = -2.03, \sigma = 1.13$	$\mu = -2.16, \sigma = 0.97$	$\mu = -2.52, \sigma = 0.85$	* $p < 0.001$

Table 6: Security/Privacy average per question/prompt comparing Age. Age ranges for 55+ were excluded to limited sample size. Data was normal ( $p < 0.01$ ) according to Anderson-Darling test. Using MANOVA indicates that there are significant differences across responses ( $Z^* = 23.68, p = 0.012$ ), post-hoc analysis (one vs. all) suggest that responses from 18-24 and 45-54 were significantly different. Using a Bonferoni correction, Q3 and Q6 show significant difference in responses. Additionally, F2 captures differences in age well, again, 18-24 and 45-54 showing the most difference.

When comparing age groups (see Table 6), again we see that particular questions expose significant differences between the groups, supported under MANOVA analysis ( $p = 0.012$ ). Post-hoc MANOVA analysis, using a one-vs.all method, suggest that 18-24 and 45-54 age groups were significant different in responses to other respondents. In particular, for older participants (45-54) there is more agreement with Q3 regarding more concern for the security/privacy of data stored on mobile devices as compared to data stored online, but younger participants (18-24), however, show more disagreement, suggesting that they may be more concerned with their online footprint. Further, there is a divide regarding Q6, known vs. unknown actors. Younger participants show more security/privacy concern regarding unauthorized access from both known and unknown

threats, and are particularly sensitive to known actors as compared to other groups.

The second factor solution (F2) also exposes an overall difference in the security/privacy sensitivity for age. Recall that more negative responses correlated with more security/privacy sensitivity, and we see an increase in that sensitivity (more negative) for increased age. Looking at the weightings to questions within F2, the response to Q1 (concern about security/privacy), Q4 (identify threat), and Q5 (online privacy) are the most strongly negative within the F2 solution. The result may indicate that these security/privacy threats resonate more strongly with older participants and that F2 is capturing that.

As a limitation to these results from age and location, there is overlap between individual groupings in these statistical tests, e.g., more young people live in urban settings, so results may be oversampling a sub-group of the population. Even correcting for that, across the hypothesis (using smallest p-values), it would suggest that Q2, Q3, and Q6 play the most significant individual roles; this is further supported in the PCA weightings for both F1 and F2 solutions.

## 6.2 Guessability Strength

In Figure 3, a scatter plot is presented which depicts the PCA two factor solutions for each participant who self-reported a pattern compared to the guessing number (log scale). Using Pearson Correlation (on the non-log guess values), we found that there exists a small correlation between the average security/privacy response and the guessing number ( $r = -0.08, p < 0.1$ ) for F1. This suggests that as individuals express stronger security/privacy preferences (more negative results in F1 indicates more agreement with the security/privacy prompts), there is a small increase in self-reported passcode strength, as represented in the guessing number. *We acknowledge that this is a limited and disheartening small effect and possibly inconclusive or a negative result, but this and other results do suggest that there likely exists some correlation between security/privacy attitudes and security of authentication choice under the right measurement circumstances.* Our intuition for there likely being a connection is due to the responses to individual questions with regard to the security of authentication choices.

Breaking down the results further per question to analyze the attitudes towards specific security/privacy prompts that have the largest impact. Figure 4 displays a box-and-whisker plot of the guessing numbers for self-reported patterns broken down by response within each question prompt. To measure significance, we performed a Mann-Whitney U-test (guessing numbers are not necessarily normally distributed) within group analysis comparing the guessing number of participants who expressed agreement (“Agree” or “Strongly Agree”) with those that express disagreement (“Disagree” or “Strongly Disagree”). We do not include the neutral (“Neither Agree nor Disagree”) group in this analysis. In total, we make 7 hypothesis tests over the same overlapping data (namely guessing numbers) that examine whether agreement/disagreement of the prompt impacts guessing number, and as such, we consider significance with a Bonferroni correction. Table 7 displays the primary results.

We find that in the case for question prompt Q1 and Q6 there is a significant difference between individuals who agree/disagree and the security of self-reported




Figure 3: Scatter plot of factor solutions to the guessing number of reported patterns. There is a negative correlation of  $r = -0.08$  using Pearson rank for Factor 1 with small a effect with a small, an inconclusive effect size  $p < 0.1$ , leading to a potential negative result. No potentially significant correlation for Factor 2 was found.

Android unlock patterns; however, with Bonferroni correction with  $m = 7$ , a weak significance of  $\alpha = 0.1$  occurs when  $p < 0.014$  so the effect may be limited. Recall that Q1 refers to a general concern with the security/privacy of data stored on mobile devices, and Q2 refers to concern with unauthorized access by known actors (e.g., friends, family, coworkers). In analysis with Q1, although the Mann-Whitney  $U$ -test is compatible with unequal group sizes, it is important to note that there were only 56 participants in the Disagree/Strongly Disagree category as compared to 351 in the Agree/Strongly Agree category, limiting the power of the test. However, for Q6, there is a more balanced sample (albeit still unequal). While the power of the test for Q6 is reduced by Bonferroni corrections, it does suggest that there are factors in the mental model of the threat—is it someone I know?—that could impact the authentication choices and requires further investigation.

## 7 Limitations

There are a number of limitations in this study that are worth noting. Foremost is that the data relies on both self-reported attitudes towards privacy prompts as well as self-reported authentication choices. While it could be the case that individuals are not forthcoming, the reported authentication choices are in line with other published results [4, 39, 38, 3].

Additionally there may be selection bias in that participants that are more willing to participate in the survey may already be less privacy conscious than those who do not.






Figure 4: Box-and-Whisker plot for guessing number strength (log-scale) compared to security/privacy response where: SD is “Strongly Disagree,” D is “Disagree,” N is “Neither Disagree nor Agree,” A is “Agree,” and SA is “Strongly Agree.”

	Disagree/Strongly Disagree	Agree/Strongly Agree	<i>U</i> -Test	Cohen- <i>d</i>
Q1	$\mu = 2568.30, \sigma = 8935.48, N = 56$	$\mu = 7111.18, \sigma = 32861.65, N = 351$	$*p = 0.011$	$d=0.148$
Q2	$\mu = 3293.11, \sigma = 8799.23, N = 150$	$\mu = 4609.32, \sigma = 20723.04, N = 201$	$p = 0.392$	$d=0.079$
Q3	$\mu = 4218.82, \sigma = 20383.17, N = 184$	$\mu = 6801.41, \sigma = 33295.58, N = 166$	$p = 0.188$	$d=0.094$
Q4	$\mu = 1496.86, \sigma = 2688.57, N = 51$	$\mu = 7028.29, \sigma = 33289.51, N = 339$	$p = 0.238$	$d=0.178$
Q5	$\mu = 5974.97, \sigma = 15860.97, N = 32$	$\mu = 6285.81, \sigma = 31447.69, N = 378$	$p = 0.418$	$d=0.010$
Q6	$\mu = 3450.97, \sigma = 10158.65, N = 146$	$\mu = 8277.26, \sigma = 39266.21, N = 235$	$*p = 0.010$	$d=0.153$
Q7	$\mu = 3499.11, \sigma = 10800.16, N = 35$	$\mu = 6629.40, \sigma = 31722.95, N = 377$	$p = 0.296$	$d=0.102$

Table 7: Significance testing for self-reported pattern strength based on agreement/disagreement with question prompt. Significance testing performed using Mann-Whitney *U*-test as guessing numbers are not normally distributed. The effect size was calculated considering groups of unequal size.

We attempted to control for this effect by allowing participants to opt-out of reporting their pattern directly, and in comparing those two groups (reported patterns vs. stats) we did not see any differences in the responses to the privacy prompts, suggesting that within the groups there is a diversity of privacy attitudes. This control does not account for measurements of individuals who chose not to participate at all and avoided the survey due to privacy concerns.

It is acknowledged that this study focuses solely on graphical password choices, rather than the range of unlock authentication mechanisms used on mobile devices (i.e., text-based password and PIN/passcodes). The focus on graphical passwords was thought to be a reasonable choice given prior-published guessing/security metrics exist for Android’s unlock pattern that would allow for ready-made comparisons. We do expect that similar results would likely hold for other unlock authentication options.

For purposes of analysis, we examined participant agreement/disagreement with statements regarding the information privacy of mobile phones, similar to that of a Westin Scale. The Westin Scale has been examined extensively by researchers in order to determine the merits of comparing self-report vs observed behavior data regarding privacy practices [22], enabling researchers to categorize individuals into privacy groups [42]. While there have been some challenges identified using such scales (e.g., between the correlation between categories and behavioral intent, and between the Westin categories and consequences), the same researchers have found that segmentations derived from the results may help to reveal deeper insights into views about privacy [42]. Other researchers have questioned the applicability of the scale (e.g., [24]), and the limitations imposed by its unidimensional nature compared with other scales [16]. However, according to [26], described by [16], it is still a useful tool that researchers use to examine how privacy attitudes evolve over time. Further work would be undertaken to better understand the limitations of the scale.

We also acknowledge throughout that the effect size and power of many of the statistical tests are limited. We believe that this work shows that there is a potential connection between privacy attitude and unlock authentication strength. There are many factors, both in the choice of the privacy scale and the choice (and construction/training) of the guessability metric that limit the scope of the results. Further analysis of this connection would need to be performed to make stronger conclusions.

Finally, there is a limitation in that this data was collected in the Fall of 2015, before the FBI and Apple incident involving the unlocking of the San Bernardino terrorist's iPhone [29]. It may be the case that individuals are more privacy sensitive if the study were re-run today given the publicity of government surveillance, particularly to Q7. It would be interesting future work to compare these results to a similar study conducted in the current climate regarding the privacy of mobile devices.

## 8 Conclusion

To conclude, we analyzed a dataset collected via an online survey on Amazon Mechanical Turk of participants agreement/disagreement with a set of seven security/privacy prompts related to information security on mobile devices. Analyzing the responses using principal component analysis (PCA) revealed two factor-solutions. We further analyzed the solutions' weighting and responses to individual questions based on demographic breakdowns, willingness to reveal or not reveal an unlock pattern, and the strength of the unlock pattern if revealed.

Originally, we hypothesized that (H1) there is a relationship between the responses to the security/privacy prompts and the strength of the reported authentication stimuli, and (H2) that there is a relationship between the participants' demographic information (e.g., age, location of residence) and the responses to the security/privacy prompts.

In regards to H1, we were unable to find sufficient evidence to accept or reject H1 outright. We found weak evidence, potentially inconclusive evidence with regard to one of our factor solutions (F1), in support of H1, but the effect size and the correlation results are not significant. Our second factor solution (F2) showed no significant differences between those that revealed and those that did not reveal their unlock pattern,

but again, this does not allow us to draw conclusions regarding the relationship between pattern strength and security/privacy attitudes. Further research would be needed and perhaps a modification to either the security/privacy prompts or the security metric, or both.

However, while we were unable to identify larger trends that include the entire set of security/privacy prompts, we did find that individual questions provided significant differences in the security of reported unlock patterns. Both Q1 (concern for privacy of data on mobile devices) and Q6 (concern about unauthorized access by known actors) proved to divide the groups well where those who agreed more strongly selected stronger patterns. This suggests that modifications to the mobile device specific security/privacy prompts should include questions that explore this connection further.

While results for H1 are inconclusive, we are able to accept H2. There are demographic differences in the responses to the security/privacy prompts, particularly as it relates to both of our factor solutions (F1 and F2). With respect to place of residence, when applying F1, we find that both urban and rural participants have similar privacy responses (more privacy conscious) that differ from sub-urban participants (less privacy conscious). When considering age, F2 showed very strong differences in the responses as participants got older, and based on the weighting, we believe this relates to the responses to Q1 (concern about privacy), Q4 (identify threat), and Q5 (online privacy) which older participants had more agreement with.

Based on these findings, we believe that incorporating some or all of the security/privacy prompts as a short survey within the password creation process could help inform a system or policy makers on how to better personalize the selection process. For example, if Android used a similar prompt during initialization, the information from that response can nudge users who are more casual about information security and privacy towards stronger authenticator.

## Acknowledgments

This work was supported by the Office of Naval Research and the National Security Agency. We would like to thank Stephen Chan, Devon Budzitowski, Jeanne Lunning-Prak, and Justin Maguire for help supporting this research.

## References

- [1] Yusuf Albayram, Mohammad Maifi Hasan Khan, Theodore Jensen, and Nhan Nguyen. “...better to use a lock screen than to worry about saving a few seconds of time”: Effect of fear appeal in the context of smartphone locking behavior. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, pages 49–63, Santa Clara, CA, 2017. USENIX Association.
- [2] Panagiotis Andriotis, George Oikonomou, Alexios Mylonas, and Theo Tryfonas. A study on usability and security features of the android pattern lock screen. *Information & Computer Security*, 24(1):53–72, 2016.

- [3] Panagiotis Andriotis, Theo Tryfonas, George Oikonomou, and Can Yildiz. A pilot study on the security of pattern screen-lock methods and soft side channel attacks. In *Proceedings of the sixth ACM conference on Security and privacy in wireless and mobile networks*, WiSec'13, pages 1–6, 2013.
- [4] Adam J Aviv, Devon Budzitwoski, and Ravi Kuber. Is bigger better? comparing user-generated passwords on 3x3 vs. 4x4 grid sizes for android's pattern unlock. In *Proceedings of the 31th Annual Computer Security Applications Conference*, ACSAC'15, 2015.
- [5] Adam J Aviv and Dane Fichter. Understanding visual perceptions of usability and security of Android's graphical password pattern. In *Proceedings of the 30th Annual Computer Security Applications Conference*, ACSAC'14, pages 286–295, 2014.
- [6] Adam J Aviv, Katherine Gibson, Evan Mossop, Matt Blaze, and Jonathan M Smith. Smudge attacks on smartphone touch screens. In *Proceedings of the 2010 Workshop on Offensive Technology*, WOOT'10, 2010.
- [7] Zinaida Benenson, Freya Gassmann, and Lena Reinfelder. Android and ios users' differences concerning security and privacy. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, pages 817–822, New York, NY, USA, 2013. ACM.
- [8] Zinaida Benenson, Anna Girard, Ioannis Krontiris, Vassia Liagkou, Kai Rannenberg, and Yannis Stamatou. User acceptance of privacy-abcs: an exploratory study. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, pages 375–386. Springer, 2014.
- [9] Joseph Bonneau. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, SP'12, pages 538–552, 2012.
- [10] Tom Buchanan, Carina Paine, Adam N. Joinson, and Ulf-Dietrich Reips. Development of measures of online privacy concern and protection for use on the internet. *J. Am. Soc. Inf. Sci. Technol.*, 58(2):157–165, January 2007.
- [11] KE Caine. Understanding privacy behaviors and disclosures. *Unpublished Dissertation. Georgia Institute of Technology. Atlanta, GA*, 2009.
- [12] Farah Chanchary and Sonia Chiasson. User perceptions of sharing, advertising, and tracking. In *SOUPS*, pages 53–67, 2015.
- [13] Nathan L Clarke and Steven M Furnell. Authentication of users on mobile telephones—a survey of attitudes and practices. *Computers & Security*, 24(7):519–527, 2005.
- [14] Serge Egelman, Sakshi Jain, Rebecca S. Portnoff, Kerwell Liao, Sunny Consolvo, and David Wagner. Are you ready to lock? In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, pages 750–761, 2014.

- [15] Serge Egelman and Eyal Peer. Predicting privacy and security attitudes. *ACM SIGCAS Computers and Society*, 45(1):22–28, 2015.
- [16] Serge Egelman and Eyal Peer. Scaling the security wall: Developing a security behavior intentions scale (sebis). In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 2873–2882, New York, NY, USA, 2015. ACM.
- [17] Marian Harbach, Alexander De Luca, and Serge Egelman. The anatomy of smartphone unlocking: A field study of android lock screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 4806–4817, 2016.
- [18] Marian Harbach, Alexander De Luca, Nathan Malkin, and Serge Egelman. Keep on lockin' in the free world: A multi-national comparison of smartphone locking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 4823–4827, 2016.
- [19] Marian Harbach, Emanuel von Zezschwitz, Andreas Fichtner, Alexander De Luca, and Matthew Smith. It's a hard lock life: A field study of smartphone (un) locking behavior and risk perception. In *Symposium on Usable Privacy and Security (SOUPS)*, 2014.
- [20] Eiji Hayashi, Oriana Riva, Karin Strauss, AJ Brush, and Stuart Schechter. Goldilocks and the two mobile devices: going beyond all-or-nothing access to a device's applications. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 2012.
- [21] Daniel Hintze, Rainhard D. Findling, Muhammad Muaaz, Sebastian Scholz, and René Mayrhofer. Diversity in locked and unlocked mobile device usage. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, UbiComp '14 Adjunct, pages 379–384, New York, NY, USA, 2014. ACM.
- [22] Carlos Jensen, Colin Potts, and Christian Jensen. Privacy practices of internet users: self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63(1):203–227, 2005.
- [23] Patrick Gage Kelley, Saranga Komanduri, Michelle L Mazurek, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, and Julio Lopez. Guess again (and again and again): Measuring password strength by simulating password-cracking algorithms. In *Proceedings of the IEE Symposium on Security and Privacy*, SP'12, pages 523–537, 2012.
- [24] Jennifer King. Taken out of context: An empirical analysis of westin's privacy scale. In *Workshop on Privacy Personas and Segmentation*, page 2014, 2014.
- [25] Ponnurangam Kumaraguru and Lorrie Faith Cranor. Privacy indexes: a survey of westin's studies. 2005.

- [26] Ponnurangam Kumaraguru and Lorrie Faith Cranor. Privacy indexes: A survey of westin’s studies. 2005. Available as *ISRI Technical Report CMU-ISRI-05-138*, 2015.
- [27] Stan Kurkovsky and Ewa Syta. Digital natives and mobile phones: A survey of practices and attitudes about privacy and security. In *Proceedings of the IEEE International Symposium on Technology and Society (ISTAS)*, pages 441–449, 2010.
- [28] Michelle Kwasny, Kelly Caine, Wendy A Rogers, and Arthur D Fisk. Privacy and technology: folk definitions and perspectives. In *CHI’08 Extended Abstracts on Human Factors in Computing Systems*, pages 3291–3296. ACM, 2008.
- [29] Eric Lichtblau and Katie Benner. U.s. presses bid to force apple to unlock iphone in new york. 2016. <http://www.nytimes.com/2016/04/09/technology/us-presses-bid-to-force-apple-to-unlock-iphone-in-new-york.html>.
- [30] Marte Loge, Markus Duermuth, and Lillian Rostad. On user choice for android unlock patterns. In *1st European Workshop on Usable Security*, 2016.
- [31] Miguel Malheiros, Sören Preibusch, and M Angela Sasse. “fairly truthful”: The impact of perceived effort, fairness, relevance, and sensitivity on personal data disclosure. In *International Conference on Trust and Trustworthy Computing*, pages 250–266. Springer, 2013.
- [32] Naresh K Malhotra, Sung S Kim, and James Agarwal. Internet users’ information privacy concerns (iupc): The construct, the scale, and a causal model. *Information systems research*, 15(4):336–355, 2004.
- [33] Nathan Malkin, Marian Harbach, Alexander De Luca, and Serge Egelman. The anatomy of smartphone unlocking: Why and how android users around the world lock their phones. *GetMobile: Mobile Comp. and Comm.*, 20(3):42–46, January 2017.
- [34] Shrirang Mare, Mary Baker, and Jeremy Gummesson. A study of authentication in daily life. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 189–206, Denver, CO, 2016. USENIX Association.
- [35] Michelle L Mazurek, Saranga Komanduri, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Patrick Gage Kelley, Richard Shay, and Blase Ur. Measuring password guessability for an entire university. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & communications security, CCS’13*, pages 173–186, 2013.
- [36] Lina Qiu, Ildar Muslukhov, and Konstantin Beznosov. Advancing the understanding of android unlocking and usage. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 2016.

- [37] Hossein Siadati, Payas Gupta, Sarah Smith, Nasir Memon, and Mustaque Ahamad. Fortifying android patterns using persuasive security framework. *UBICOMM 2015*, page 81, 2015.
- [38] Youngbae Song, Geumhwan Cho, Seongyeol Oh, Hyoungshick Kim, and Jun Ho Huh. On the effectiveness of pattern lock strength meters: Measuring the strength of real world pattern locks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, 2015.
- [39] Sebastian Uellenbeck, Markus Dürmuth, Christopher Wolf, and Thorsten Holz. Quantifying the security of graphical passwords: The case of Android unlock patterns. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, CCS'13, 2013.
- [40] Emanuel von Zezschwitz, Malin Eiband, Daniel Buschek, Sascha Oberhuber, Alexander De Luca, Florian Alt, and Heinrich Hussmann. On quantifying the effective password space of grid-based unlock gestures. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia*, MUM '16, pages 201–212, New York, NY, USA, 2016. ACM.
- [41] Emanuel von Zezschwitz, Malin Eiband, Daniel Buschek, Sascha Oberhuber, Alexander De Luca, Florian Alt, and Heinrich Hussmann. On quantifying the effective password space of grid-based unlock gestures. In *Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia*, pages 201–212. ACM, 2016.
- [42] Allison Woodruff, Vasyli Pihur, Sunny Consolvo, Lauren Schmidt, Laura Brandimarte, and Alessandro Acquisti. Would a privacy fundamentalist sell their dna for 1000... if nothing bad happened as a result? the westin categories, behavioral intentions, and consequences. In *SOUPS*, pages 1–18, 2014.