Predicting Privacy and Security Attitudes

Serge Egelman
International Computer Science Institute
Berkeley, CA, USA
egelman@cs.berkeley.edu

Eyal Peer
Bar-Ilan University
Ramat Gan, Israel
eyal.peer@biu.ac.il

ABSTRACT

While individual differences in decision-making have been examined within the social sciences for several decades, this research has only recently begun to be applied by computer scientists to examine privacy and security attitudes (and ultimately behaviors). Specifically, several researchers have shown how different online privacy decisions are correlated with the "Big Five" personality traits. However, in our own research, we show that the five factor model is actually a weak predictor of privacy preferences and behaviors, and that other well-studied individual differences in the psychology literature are much stronger predictors. We describe the results of several experiments that showed how decision-making style and risk-taking attitudes are strong predictors of privacy attitudes, as well as a new scale that we developed to measure security behavior intentions. Finally, we show that privacy and security attitudes are correlated, but orthogonal.

Categories and Subject Descriptors

H.1.2. [**Models and Principles**]: User/Machine Systems; K.6.5. [**Management of Computing and Information Systems**]: Security and Protection

1. INTRODUCTION

Current systems, when they take a user-centric approach, are designed for the "average user," but no one person perfectly fits this definition. For instance, due to varying privacy preferences, no one set of default privacy settings will accommodate all users. As a result, current systems are often designed to satisfy majorities, pluralities, or the most vocal. As a result, systems can only reach local maxima when designed for human beings in the general case. Thus, we believe that we should stop satisficing and start optimizing by designing privacy and security mitigations that account for individual differences.

Psychology researchers have studied how individual differences impact decision-making [2]; literature has shown how particular behaviors are correlated with latent constructs (e.g., attitudes towards risk), and that various scales can be used to measure those constructs. If some of these constructs are also predictive of privacy preferences, then measurements of those latent constructs (e.g., using scales or observations of related behaviors) can be used to infer an individual's privacy preferences without directly asking her. Similarly, if other constructs are correlated with security decision-making, then measurements of those constructs can be used to tailor security messaging to result in better security outcomes (e.g., warning messages that appear more salient, and are therefore more likely to be obeyed [4]).

The goal of studying individual differences in decision-making is to deepen the understanding of a certain decision-making phenomena and explore whether a certain effect is more pronounced for individuals who exhibit a high or low degree in one or more individual trait measures. For example, several studies have shown that individuals with low numeracy are less likely to understand health risks that are presented to them, and that they are more susceptible to effects of mood and how the information is presented, framed or ordered [20]. Some preliminary evidence also exists for how individual differences predict privacy attitudes: for instance, Pedersen showed that individuals showing low self-esteem are more likely to seek solitude [19].

Our goal is to design systems that can make inferences about their users that can then be used to tailor privacy and security mitigations to an individual user's needs. As a first step, we have begun examining how individual differences can be used to predict privacy and security attitudes and behaviors. In this article, we describe an initial series of experiments that we performed to show how certain individual differences are predictive of privacy and security attitudes and behaviors. Contrary to the existing literature on predicting privacy preferences using personality traits, we show that the "five factor model" is a very weak predictor of privacy preferences, relative to other well-studied individual differences in the psychology literature.

2. RELATED WORK

Studying how individual differences affect people's decisions has only recently gained attention [2]. "Individual differences" covers any variable that differs between people, from demographics to abilities to personality. The Decision Making Individual Differences Inventory is an online collection, assembled by various decision-making researchers, and lists an extensive array of individual differences measures that relate to decision style or decision approach, measures of risk attitudes and behaviors, cognitive abilities, motivational measures, personality traits and more [2].

The Big 5 personality model, also known as the "five factor model," is one of the most widely used personality models in the field of psychology [7]. The five dimensions are:

- **Openness to new experiences**: the extent to which someone seeks intellectual stimulation.
- **Conscientiousness**: the extent to which someone is organized or self-disciplined.

- **Extraversion**: the extent to which someone is outgoing and enjoys socializing.
- **Agreeableness**: the extent to which someone is compassionate or empathetic.
- **Emotional Stability**: the extent to which someone is stable versus neurotic, insecure, or nervous.

Very recently, several researchers have begun to examine how privacy preferences may be predicted by using the Big Five model. For instance, Junglas *et al.* found that agreeableness, conscientiousness, and openness to new experiences all correlate with an individual's concern for using location-based services [13]. Korzaan and Boswell found that agreeableness correlated with "concern for information privacy" [14]. These correlations go beyond self-reported privacy concerns, and can also be observed with regard to behaviors: Gou *et al.* found that aspects of users' public Twitter tweets can be used to infer their Big 5 dimensions [10].

Although individual differences in privacy attitudes have been extensively investigated [1], there is little (if any) research about how other individual differences—beyond the Big 5 traits—predict people's privacy attitudes and behaviors. Exploring the effect of individual differences on self-disclosure behaviors and preferences may lead to systems that better empower users to act according to their stated privacy preferences. Similarly, since we are unaware of any previous research that has examined how security attitudes and behaviors may correlate with individual differences, we believe that exploring this may lead to higher compliance rates with security messaging.

In the privacy domain, we conducted two experiments to correlate psychometrics with observed privacy attitudes and behaviors. First, we examined whether personality traits correlate with privacy preferences and privacy-preserving behaviors. Subjects completed the Ten Item Personality Inventory (TIPI) [9], which we found to weakly correlate with privacy attitudes and privacy-preserving behaviors. Based on our results, we performed a followup experiment to show that constructs relating to risk-taking and decision-making style are much stronger predictors of privacy attitudes. In the security domain, we developed and validated a new scale to measure users' security behavior intentions, and then showed how it correlates with various individual differences.

3. EXPERIMENT 1: PERSONALITY

Our first experiment focused on personality traits (i.e., the Big 5 model [7]) as potential predictors of privacy preferences and behaviors. While we observed that personality traits correlate with preferences and behaviors, consistent with prior research, the overall predictive value is quite low. In this section, we describe our method and result.

3.1 Method

In this experiment, participants completed a personality test, as well as several different privacy metrics, which measured both stated preferences (i.e., privacy attitudes) and observed behaviors (i.e., participants' willingness to disclose private information about themselves). The order in which they completed each test was randomized, as was the question ordering within each test.

We measured participants' personality dimensions using the Ten Item Personality Index (TIPI) [9], a 10-question survey instrument featuring two questions per personality dimension. Each question is answered using a Likert scale. We used the five dimensions as independent variables in a regression, which also included demographic factors (i.e., gender, income, and education level). Thus, each regression model featured eight independent variables.

Our dependent variables consisted of various privacy attitude and behavior metrics. We examined privacy attitudes using the Privacy Concerns Scale (PCS) [5]. The PCS is a set of 16 Likert-scale questions used to evaluate privacy attitudes on a unidimensional scale with regard to how concerned Internet users are with various scenarios involving misuse of personal information. We also used both the Westin Index [15] and the Internet Users Information Privacy Concerns (IUIPC) scale [16]. The Westin Index measures consumers' general attitudes about privacy using 3 Likert-scale questions that segment the population into three categories: "Fundamentalists," "Pragmatists," and "Unconcerned." Despite being used for several decades, researchers have recently raised questions about its validity [23]. The IUIPC scale features 10 Likert-scale questions evaluated across three dimensions: control over personal information, awareness of privacy practices, and data collection concerns.

Finally, we measured privacy behaviors by examining participants' self-disclosure behaviors two different ways. First, we used the 10-item Strahan-Gerbasi version of the Marlowe-Crowne Social Desirability Scale (SDS) [22]. The SDS measures social desirability bias, which is the propensity for people to respond to questions in ways that make them appear more desirable to others. To that end, the scale features 10 true/false statements; half reflect socially desirable behaviors that are rare (e.g., "I am always a good listener"), whereas the other reflect socially undesirable behaviors that are common (e.g., "I sometimes try get even rather than forgive and forget"). We coded SDS responses by adding the number of "true" responses to socially undesirable traits with the number of responses of "false" to socially desirable traits. Thus, self-disclosure scores ranged from 0 to 10.

Our second metric for self-disclosure behaviors was an unethical behaviors scale developed by John *et al.* [11]. This scale included 14 items that pertained to unethical or immoral behaviors (e.g., "Have you ever stolen anything worth more than $25?") to which participants could indicate a response between 1 (never) to 5 (many times) or skip the item if they wished not to answer it. We followed John *et al.* and coded responses for all items as Affirmative Admissions Rates (AARs) [11], which represented the frequency with which participants reported engaging in the unethical behaviors (i.e., not selecting "never" or skipping the item).

To examine whether and how personality traits (measured by the TIPI) predict privacy attitudes and self-disclosure, we ran multiple regression analyses using TIPI, gender, education level and income level as predictors on the following dependent variables: PCS, IUIPC (both overall and the 3 sub-scales), Westin Index, disclosure of socially undesirable traits (SDS) and disclosure of unethical behavior (AARs).

To minimize the likelihood of participants selecting responses to questions at random, we included two attention-check questions. First, the beginning of the survey featured the following instructions and questions:

> *This study requires you to voice your opinion using the scales below. It is important that you take the time to read all instructions and that you read questions carefully before you answer them. Pre-*

> *vious research on preferences has found that some people do not take the time to read everything that is displayed in the questionnaire. The questions below serve to test whether you actually take the time to do so. Therefore, if you read this, please answer 'three' on the first question, add three to that number and use the result as the answer on the second question. Thank you for participating and taking the time to read all instructions.*
>
> *I would prefer to live in a large city rather than a small city.* [Strongly disagree (1), (2), (3), (4), (5), (6), Strongly agree (7)]
>
> *I would prefer to live in a city with many cultural opportunities, even if the cost of living was higher.* [Strongly disagree (1), (2), (3), (4), (5), (6), Strongly agree (7)]

If participants did not select "3" and "6," respectively, we gave them a second opportunity to answer these questions correctly. Upon answering them incorrectly a second time, we disqualified them from completing the survey. Additionally, we included an 11th item within the SDS questions: *I do not read the questions in surveys.* We filtered out participants who responded "true" to this question *post hoc.*

We recruited 500 participants from Amazon's Mechanical Turk (MTurk). We restricted participation to those over 18, based in the U.S., and with previous task completion rates exceeding 95%. After filtering out 43 responses (8.6% of 500) based on the second attention-check question (i.e., we did not receive results from participants who incorrectly answered the attention-check at the beginning of the survey), we were left with a sample of 457 valid responses. Of our sample, 58.2% were male, and had a mean age of 32.91 ($\sigma = 11.19$). Most participants had either completed high school (33%) or held a bachelor's degree (33.3%) or an associate's degree (15.5%). Median income category was $35K-$50K and the majority of participants (79%) reported an income lower than $75K per year.

3.2 Results

We first performed Principal Component Analysis (PCA) to verify each scale's dimensionality, as well as determined internal reliability using Cronbach's α. PCA with Varimax rotation on the PCS showed two components with eigenvalues greater than 1 that predicted 58.72% of the total variance. However, the second component only added 8.26% to the predicted variance and the reliability of the entire scale was high ($\alpha = .933$) so we treated it as measuring one factor, as prescribed by Buchanan *et al.* [5], and computed average PCS scores for all participants. Regarding the IUIPC, PCA showed the three original components predicted 75.18% of the total variance: Control ($\alpha = .792$), Awareness ($\alpha = .776$), and Collection ($\alpha = .908$). We noted that one item was cross-loaded on Collection (.495) and Awareness (.455). Based on the recommendations of Matsunaga [17], we made a judgment call to retain the original structure (keeping the item with its intended factor, Awareness). The Westin Index showed adequate reliability ($\alpha = .692$), and we did not assess the internal reliability of the TIPI, as its authors recommend against it [9].[1]

Table 1 summarizes our regression analyses. As can be seen, personality traits had a low predictive ability towards any of the privacy scales, and the total predicted variance was less than 1% for all dependent variables. Among the personality sub-scales, only Agreeableness predicted the PCS (along with income level); Openness to new experiences was the highest and most stable predictor of IUIPC (overall and all sub-scales), followed by Conscientiousness (which predicted awareness and collection but not control), Agreeableness (which predicted only awareness) and Extraversion (which predicted control). Agreeableness also predicted SDS, followed by Conscientiousness, which also predicted AARs, as did Openness and income level.

For the Westin Index, which classifies individuals into three groups, we performed a multinomial regression analysis with the same predictors. Only education level showed a significant result in predicting classification to the three groups ($\chi^2(14) = 23.95$, $p = .046$) while none of the other variables showed any significant prediction. This corroborates prior research showing that the Westin Index is a poor predictor of privacy preferences or behaviors [23], and therefore we decided to not consider it further.

4. EXPERIMENT 2: DECISION-MAKING

Our second experiment focused on individual differences in decision-making as potential predictors of privacy attitudes. Overall, we observed that decision-making style and risk-taking attitudes were much better predictors.

4.1 Method

We made three changes from our first experiment. First, we did not include the Westin Index, as it performed poorly compared to the other privacy attitudes scales. Second, we also chose not to include behavioral tendency measures (such as the SDS and admissions to unethical behaviors) and focused solely on privacy attitudes scales (i.e., the PCS and IUIPC). Finally, we decided to use decision-making psychometrics as our predictors: Need for Cognition (NFC) [6], the General Decision Making Style (GDMS) scale [21], and the Domain Specific Risk Attitude (DoSpeRT) scale [3].

NFC is a unidimensional scale that measures tendency to engage in "thoughtful endeavors" [6]. The GDMS scale measures decision-making style across five dimensions [21]:

- **Rational**: Using logic when making decisions.
- **Avoidant**: Delaying making decisions.
- **Dependent**: Making decisions by looking to others.
- **Intuitive**: Making "gut" decisions.
- **Spontaneous**: Making quick decisions.

The DoSpeRT measures attitudes towards engaging in risks across five dimensions [3]: financial, health/safety, recreational, ethical, and social. Just as before, the order of all questionnaires was randomized between participants.

We recruited a new cohort of 500 participants from Amazon's Mechanical Turk and required that they had not participated in the previous experiment. Using the same screening requirements as in our previous experiment, we filtered out 4 responses (0.8% of 500) and were left with a sample of 496. Half (51%) of participants were male, and the mean age was 35.33 ($\sigma = 11.6$). Most participants had either completed high school (33%) or held a bachelor's degree (33.5%)

[1] Reliability is established via test-retest [9], which we did not do due to the high internal reliability of the other scales.

	PCS	IUIPC				SDS	AARs
		Overall	Control	Awareness	Collection		
Extraversion	-0.002	-0.085	**-0.108**	-0.066	-0.051	-0.034	0.095
	(0.962)	(0.085)	(0.030)	(0.178)	(0.311)	(0.486)	(0.054)
Agreeableness	**0.129**	0.060	0.029	**0.126**	0.015	**0.218**	-0.063
	(0.011)	(0.233)	(0.562)	(0.011)	(0.770)	(<0.001)	(0.208)
Conscientiousness	0.092	**0.110**	0.083	**0.099**	**0.104**	**0.118**	**-0.190**
	(0.071)	(0.029)	(0.103)	(0.048)	(0.044)	(0.019)	(<0.001)
Emotional Stability	-0.102	**0.110**	-0.107	-0.092	-0.088	0.074	0.061
	(0.067)	(0.046)	(0.054)	(0.093)	(0.115)	(0.179)	(0.266)
Openness	0.074	**0.249**	**0.245**	**0.231**	**0.180**	0.005	**0.109**
	(0.141)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(0.916)	(0.028)
Income level	**-0.099**	0.059	**0.129**	0.042	-0.006	-0.010	**-0.183**
	(0.033)	(0.197)	(0.005)	(0.360)	(0.896)	(0.821)	(<0.01)
Education level	-0.003	0.050	0.051	0.010	0.062	0.062	-0.042
	(0.952)	(0.279)	(0.270)	(0.831)	(0.187)	(0.177)	(0.360)
Male	-0.091	-0.069	0.008	**-0.095**	-0.091	0.063	-0.024
	(0.064)	(0.155)	(0.863)	(0.050)	(0.065)	(0.190)	(0.617)
F	**3.665**	**5.247**	**4.413**	**6.062**	**5.546**	**6.136**	**6.003**
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
Adjusted R^2	.045	.069	.056	.082	.040	.083	.081

Table 1: Regression analysis with privacy preferences/behaviors as dependent variables and the Big Five personality traits as independent variables, controlling for demographic factors. Numbers in parentheses show the p-value; values in bold are statistically significant at the .05 level.

	PCS		IUIPC							
			Overall		Control		Awareness		Collection	
NFC	-0.038	(0.467)	-0.03	(0.54)	-0.041	(0.419)	-0.005	(0.919)	-0.027	(0.606)
GDMS-Intuitive	**0.175**	(0.001)	**0.149**	(0.002)	0.038	(0.436)	**0.144**	(0.002)	**0.173**	(0.001)
GDMS-Rational	**0.252**	(<0.001)	**0.315**	(<0.001)	**0.301**	(<0.001)	**0.283**	(<0.001)	**0.228**	(<0.001)
GDMS-Avoidant	0.076	(0.142)	0.001	(0.977)	-0.055	(0.272)	-0.038	(0.421)	0.064	(0.219)
GDMS-Dependent	**0.102**	(0.044)	-0.01	(0.831)	-0.06	(0.214)	0.0	(0.993)	0.022	(0.667)
GDMS-Spontaneous	0.000	(1.000)	0.008	(0.895)	-0.055	(0.358)	0.022	(0.699)	0.039	(0.53)
RT-Ethical	0.078	(0.215)	**-0.125**	(0.034)	-0.112	(0.067)	**-0.235**	(<0.001)	-0.02	(0.756)
RT-Health/Safety	**-0.213**	(0.001)	-0.105	(0.090)	-0.064	(0.317)	0.014	(0.816)	**-0.169**	(0.011)
RT-Recreational	**0.116**	(0.040)	-0.003	(0.949)	-0.027	(0.617)	-0.066	(0.199)	0.053	(0.347)
RT-Social	0.072	(0.164)	**0.259**	(<0.001)	**0.230**	(<0.001)	**0.241**	(<0.001)	**0.195**	(<0.001)
RT-Financial	0.032	(0.559)	-0.09	(0.082)	-0.071	(0.184)	**-0.102**	(0.043)	-0.063	(0.259)
Male	**-0.121**	(0.009)	-0.039	(0.363)	-0.024	(0.584)	-0.03	(0.472)	-0.04	(0.381)
Education level	0.005	(0.919)	-0.032	(0.454)	-0.043	(0.337)	-0.003	(0.937)	-0.031	(0.501)
Income level	0.02	(0.665)	-0.054	(0.215)	-0.032	(0.482)	-0.037	(0.387)	-0.06	(0.198)
F	**5.332**	(<0.001)	**11.120**	(<0.001)	**8.123**	(<0.001)	**13.58**	(<0.001)	**5.385**	(<0.001)
Adjusted R^2	.113		.230		.174		.270		.114	

Table 2: Regression analysis with privacy attitudes as dependent variables and decision-making psychometrics as independent variables, controlling for demographic factors. Numbers in parentheses show the p-value; values in bold are statistically significant at the .05 level.

or an associate's degree (17.3%). Median income category was \$25K-\$50K and the majority of participants (87%) reported an income lower than \$75K per year. These demographics are very similar to those of our initial experiment.

4.2 Results

We first analyzed our data in terms of scale reliability. As before, PCS showed high reliability ($\alpha = .936$). A confirmatory PCA on IUIPC showed the original three factors (this time all items loaded highest on their predicted factor), and the factors showed high reliability ($\alpha = .805$, .829, and .908 for Control, Awareness and Collection, respectively). NFC also showed high reliability ($\alpha = .952$). A confirmatory PCA on GDMS showed that the original five factors all had an eigenvalue larger than 1 and predicted a total of 68.18% of the variance. The factors included the different decision-making styles labeled Rational $\alpha = .787$), Avoidant ($\alpha = .918$), Dependent ($\alpha = .809$), Intuitive ($\alpha = .897$), and Spontaneous ($\alpha = .863$). For the DoSpeRT, a confirmatory PCA showed the original five factors which had an eigenvalue larger than 1 and predicted 53.44% of the total variance: Ethical risk-taking ($\alpha = .772$); Health/Safety risk-taking ($\alpha = .737$); Recreational risk-taking ($\alpha = .846$); Social risk-taking ($\alpha = .744$); Financial risk-taking ($\alpha = .831$). Thus, we concluded that our data were reliable, and we proceeded to build our regression models.

Table 2 summarizes the results of multiple regression analyses conducted on all dependent variables (PCS, IUIPC overall and sub-scales) with NFC, GDMS sub-scales, and DoSpeRt sub-scales as predictors, alongside gender, education and income level. While NFC did not show a significant correlation with any of the privacy attitudes scales, two GDMS styles significantly predicted privacy attitudes on (almost) all scales: Intuitive and Rational. Given the positive standardized coefficients, this suggests that stronger privacy attitudes are the result of rational decision-making, as well as people having "gut feelings" about not wanting to divulge information.

	Device Securement: The extent to which someone locks their computer screen or uses a PIN/password to lock a tablet/smartphone.
1	I set my computer screen to automatically lock if I don't use it for a prolonged period of time.
2	I use a password/passcode to unlock my laptop or tablet.
3	I manually lock my computer screen when I step away from it.
4	I use a PIN or passcode to unlock my mobile phone.
	Password Generation: The extent to which someone chooses strong passwords or does not reuse passwords between different accounts.
5	I do not change my passwords, unless I have to.^{r}
6	I use different passwords for different accounts that I have.
7	When I create a new online account, I try to use a password that goes beyond the site's minimum requirements.
8	I do not include special characters in my password if it's not required.^{r}
	Proactive Awareness: The extent to which someone pays attention to contextual cues, such as the URL bar or other browser indicators.
9	When someone sends me a link, I open it without first verifying where it goes.^{r}
10	I know what website I'm visiting based on its look and feel, rather than by looking at the URL bar.^{r}
11	I submit information to websites without first verifying that it will be sent securely (e.g., SSL, "https://", a lock icon).^{r}
12	When browsing websites, I mouseover links to see where they go, before clicking them.
13	If I discover a security problem, I continue what I was doing because I assume someone else will fix it.^{r}
	Updating: The extent to which someone applies security patches or otherwise keeps their software up to date.
14	When I'm prompted about a software update, I install it right away.
15	I try to make sure that the programs I use are up-to-date.
16	I verify that my anti-virus software has been regularly updating itself.

Table 3: The SeBIS items and sub-scales. Six questions are reverse-scored (denoted by ^{r}). Responses are reported on the following scale: *Never (1)*, *Rarely (2)*, *Sometimes (3)*, *Often (4)*, **and** *Always (5)*.

Among the risk-taking measures, social risk-taking significantly predicted almost all of the privacy attitudes scales, and health/safety risk-taking negatively predicted PCS and the collection sub-scale of IUIPC. That is, people who are more likely to challenge social norms are also more likely to question company policies about how personal information is handled. Similarly, those who take fewer safety risks are more likely to have concerns about their privacy.

Comparing the results of our two experiments, we can see that the second model has strong fit; averaged across all five dependent variables, the coefficient of determination was over three times as large in the second model, relative to the first. Thus, future research to predict privacy attitudes and behaviors should probably focus on decision-making psychometrics, rather than the five factor model.

5. EXPERIMENT 3: SECURITY

While several scales exist in the literature for measuring privacy attitudes, we are unaware of any similar scales for measuring security attitudes. Thus, we developed a new scale, which we are calling the Security Behavior Intentions Scale (SeBIS) [8]. SeBIS features 16 items, spread across 4 dimensions, and scored on a 5-point Likert scale ("never," "rarely," "sometimes," "often," and "always"). The sub-scales and associated items are depicted in Table 3. We performed an experiment to examine how various decision-making psychometrics predict security attitudes, as measured by SeBIS.

5.1 Method

We used our previous methodology to examine correlations between SeBIS and our previous decision-making psychometrics (i.e., GDMS, DoSpeRT, and NFC), but also added in the Barratt Impulsiveness Scale (BIS) and Consideration for Future Consequences (CFC). BIS measures impulsivity across three dimensions [18]: attention, motor, and nonplanning. CFC is a unidimensional scale that measures how much attention people pay to long-term consequences [12].

	Securement	Passwords	Awareness	Updating
RT_e		**-.201	**-.226	**-.201
RT_h			**-.204	**-.164
RT_r				
RT_s		*.141		
RT_f				
$GDMS_i$				
$GDMS_r$		**.145	**.224	**.229
$GDMS_a$	*-.133	**-.220	**-.230	**-.247
$GDMS_d$			**-.157	
$GDMS_s$				*-.129
NFC	**.164	**.290	**.231	**.253
BIS_a		**-.243	*-.140	**-.218
BIS_m			**-.147	**-.145
BIS_p		**-.235	**-.171	**-.247
CFC	**.184	**.317	**.307	**.303

Table 4: Correlations between SeBIS sub-scales and various psychometrics. $^{*}p < 0.005$, $^{**}p < 0.001$.

We recruited 500 participants from MTurk to complete SeBIS, as well as the 5 psychometric tests. Because this work was exploratory in nature, we opted to perform Pearson correlations,[2] rather than building a full regression model.

5.2 Results

Table 4 shows the resulting correlation matrix. To counteract effects from repeated testing, we only report correlations significant at the $p < 0.005$ level. Across all four sub-scales, participants who were more inquisitive (as determined by the NFC scale) were more likely to report engaging in better security practices. We also noted that the highest correlation across all four sub-scales was with consideration for future consequences (CFC), suggesting that good security behaviors are tied to long-term thinking.

[2] Q-Q plots indicated that average scale scores were normally distributed. We did not perform Shapiro-Wilk or Kolmogorov-Smirnov due to our large sample size.

Our initial hypothesis was that willingness to take risks involving health/safety would be inversely correlated with computer security behaviors; we observed this was true, but only for behaviors involving proactive awareness and keeping software updated. Similarly, across all sub-scales, people who engaged in "better" security behaviors were less likely to procrastinate (as determined by the GDMS avoidant subscale). People scoring low on GDMS dependence scored high on SeBIS awareness; being proactive about computer security means *not* relying on other people. Many of the security behaviors also correlated inversely with impulsivity: this suggests that adhering to security advice involves foresight.

6. CONCLUSION

Through our preliminary experiments, we show that privacy attitudes can be predicted by examining several well-studied psychometrics from the psychology literature. While previous research has shown that personality traits are predictive of privacy attitudes (e.g., the five factor model), we show that individual differences pertaining to decision-making and risk-taking are much stronger predictors.

Similarly, these individual differences are also predictive of security behavior intentions, though in different ways. For instance, while the "avoidant" dimension of the General Decision-Making Style (GDMS) scale had no predictive power over any of the privacy attitudinal metrics that we examined, it was correlated with all four dimensions of the Security Behavior Intentions Scale (SeBIS). Thus, privacy and security attitudes are different constructs, and are therefore correlated with different individual differences. (We corroborated this by showing that the Privacy Concerns Scale (PCS) is only weakly correlated with SeBIS [8].)

The purpose of this preliminary work is to illustrate how an individual's privacy and security attitudes can be predicted based on other individual differences. Our goal is to design systems that can infer these differences and then use these inferences to customize privacy and security mitigations on an individual basis. Thus, our next steps are to examine the myriad ways of inferring these differences based on observations of other types of behavior, as well as to explore the ways in which knowledge of users' privacy and security attitudes can be used to improve the user experience and nudge them towards making better security decisions.

7. ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under award CNS-1343433.

8. REFERENCES

[1] A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy Enhancing Technologies Workshop (PET '06)*, 2006.

[2] K. C. Appelt, K. F. Milch, M. Handgraaf, and E. U. Weber. The decision making individual differences inventory and guidelines for the study of individual differences in judgment and decision-making research. *Judgment and Decision Making*, 6(3):252–262, April 2011.

[3] A.-R. Blais and E. U. Weber. A domain-specific risk-taking (dospert) scale for adult populations. *Judgment and Decision Making*, 1(1):33–47, 2006.

[4] J. Blythe, J. Camp, and V. Garg. Targeted risk communication for computer security. In *Proceedings of the 16th International Conference on Intelligent User Interfaces*, IUI '11, pages 295–298, New York, NY, USA, 2011. ACM.

[5] T. Buchanan, C. Paine, A. N. Joinson, and U.-D. Reips. Development of measures of online privacy concern and protection for use on the internet. *Journal of the American Society for Information Science and Technology*, 58(2):157–165, 2007.

[6] J. T. Cacioppo, R. E. Petty, and C. Feng Kao. The efficient assessment of need for cognition. *Journal of personality assessment*, 48(3):306–307, 1984.

[7] P. T. Costa and R. R. McCrae. The revised neo personality inventory (neo-pi-r). *The SAGE handbook of personality theory and assessment*, 2:179–198, 2008.

[8] S. Egelman and E. Peer. Scaling the Security Wall: Developing a Security Behavior Intentions Scale (SeBIS). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '15, New York, NY, USA, 2015. ACM. To appear.

[9] S. D. Gosling, P. J. Rentfrow, and W. B. Swann Jr. A very brief measure of the big-five personality domains. *Journal of Research in personality*, 37(6):504–528, 2003.

[10] L. Gou, M. X. Zhou, and H. Yang. Knowme and shareme: Understanding automatically discovered personality traits from social media and user sharing preferences. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '14, pages 955–964, New York, NY, USA, 2014. ACM.

[11] L. K. John, A. Acquisti, and G. Loewenstein. Strangers on a plane: context-dependent willingness to divulge sensitive information. *Journal of Consumer Research*, 37(5):858–873, 2011.

[12] J. Joireman, M. J. Shaffer, D. Balliet, and A. Strathman. Promotion orientation explains why future-oriented people exercise and eat healthy evidence from the two-factor consideration of future consequences-14 scale. *Personality and Social Psychology Bulletin*, 38(10):1272–1287, 2012.

[13] I. A. Junglas, N. A. Johnson, and C. Spitzmuller. Personality traits and concern for privacy: an empirical study in the context of location-based services. *European Journal of Information Systems*, 17(4):387–402, print 2008.

[14] M. L. Korzaan and K. T. Boswell. The influence of personality traits and information privacy concerns on behavioral intentions. *Journal of Computer Information Systems*, 48(4):15–24, 2008.

[15] P. Kumaraguru and L. F. Cranor. Privacy Indexes: A Survey of Westin's Studies. Technical Report CMU-ISRI-5-138, Carnegie Mellon University, December, 2005. `http://reports-archive.adm.cs.cmu.edu/anon/isri2005/abstracts/05-138.html`.

[16] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet users' information privacy concerns (iuipc): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, December 2004.

[17] M. Matsunaga. How to factor-analyze your data right: Do's, don'ts, and how-to's. *International Journal of*

Psychological Research, 3(1):97–110, 2010.

[18] J. H. Patton, M. S. Stanford, et al. Factor structure of the barratt impulsiveness scale. *Journal of clinical psychology*, 51(6):768–774, 1995.

[19] D. M. Pedersen. Personality correlates of privacy. *The Journal of Psychology*, 112(1):11–14, 1982.

[20] V. F. Reyna, W. L. Nelson, P. K. Han, and N. F. Dieckmann. How numeracy influences risk comprehension and medical decision making. *Psychological bulletin*, 135(6):943, 2009.

[21] S. G. Scott and R. A. Bruce. Decision-making style: The development and assessment of a new measure. *Educational and psychological measurement*, 55(5):818–831, 1995.

[22] R. Strahan and K. C. Gerbasi. Short, homogeneous versions of the marlowe-crowne social desirability scale. *Journal of clinical psychology*, 1972.

[23] A. Woodruff, V. Pihur, S. Consolvo, L. Brandimarte, and A. 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 *Proceedings of the 2014 Symposium on Usable Privacy and Security*, pages 1–18. USENIX Association, 2014.