Can Privacy Nudges be Tailored to Individuals' Decision Making and Personality Traits?

Logan Warberg
Carnegie Mellon University
warberg@cmu.edu

Alessandro Acquisti
Carnegie Mellon University
acquisti@andrew.cmu.edu

Douglas Sicker
Carnegie Mellon University
sicker@cmu.edu

ABSTRACT

While the effectiveness of nudges in influencing user behavior has been documented within the literature, most prior work in the privacy field has focused on 'one-size-fits-all' interventions. Recent behavioral research has identified the potential of tailoring nudges to users by leveraging individual differences in decision making and personality. We present the results of three online experiments aimed at investigating whether nudges tailored to various psychometric scales can influence participants' disclosure choices. Each study adopted a difference-in-differences design, testing whether differences in disclosure rates for participants presented with a nudge were affected by differences along various psychometric variables. Study 1 used a hypothetical disclosure scenario to measure participants' responses to a single nudge. Study 2 and its replication (Study 3) tested responses in real disclosure scenarios to two nudges. Across all studies, we failed to find significant effects robustly linking any of the measured psychometric variables to differences in disclosure rates. We describe our study design and results along with a discussion of the practicality of using decision making and personality traits to tailor privacy nudges.

CCS CONCEPTS

• **Security and privacy → Economics of security and privacy**; *Usability in security and privacy*; Privacy protections.

ACM Reference Format:
Logan Warberg, Alessandro Acquisti, and Douglas Sicker. 2019. Can Privacy Nudges be Tailored to Individuals' Decision Making and Personality Traits?. In *18th Workshop on Privacy in the Electronic Society (WPES'19), November 11, 2019, London, United Kingdom.* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3338498.3358656

1 INTRODUCTION

Nudges have emerged within the privacy and security literature as an effective means of affecting, and possibly assisting, user behavior [1]. Nudges work by modifying the structure of choices to encourage certain behaviors without altering economic incentives [40]. Recent research has applied nudges to diverse privacy and security scenarios where users face hurdles in the decision making process, and where those hurdles can result in negative outcomes. Within the security literature, nudges have been used to steer users towards better security behaviors in areas including password creation [42], heeding browser warnings [26], and selecting wireless networks [41]. In a similar vein, nudges within the privacy literature have been used to guide users to better privacy choices, often by providing them with salient information to aid the decision making process regarding online disclosures. Privacy nudges have been used to help add context to decisions such as setting mobile app permissions [5] and posting information on social networks [44]. Other nudges have focused on changes in the presentation of choices to influence user behavior; framing effects have been shown to encourage users to allow or prohibit information disclosures [34].

While effective, many of the applications of nudges within the privacy literature have focused on 'one-size-fits-all' approaches, where a certain behavioral intervention is applied to a diverse set of individuals, with no tailoring of the nudge based on characteristics unique to each individual. This is starting to change, where recent studies have examined nudges applied to privacy behaviors that are tailored to traits such as demographics [23]. Recent behavioral work has highlighted the possibility of making nudges more effective by tailoring them based on individual traits [15]. Given the prominent role that differences in individual traits can play in terms of privacy decision making, one could expect the privacy field to be one where the tailoring of nudges could prove particularly effective. Previous work has identified significant diversity in the privacy attitudes and behaviors across cultures and among different individuals within the same culture [27, 43, 46]. Existing research has also examined individual psychometric differences and found them to be predictors of differences in privacy attitudes [15]. This diversity suggests that people may approach privacy decisions differently. Consequently, behavioral interventions that capture such diversity may be more effective in changing user behavior. Specifically, differences in decision making and personality may have applications in creating personalized or 'tailored' nudges. In the security field, recent work has started investigating the effectiveness of 'tailored' nudges in influencing security behavior [26, 31].

We investigate psychometrically tailored nudges in the context of privacy behavior. We conducted three online experiments that attempted to identify effects for tailored nudges on data disclosure choices. Across the three studies, participants completed surveys in which they were presented with a hypothetical (Study 1) or real (Studies 2 and 3) disclosure choice. Each disclosure choice was paired with a nudge designed to encourage participants to either allow or prohibit the disclosure of potentially sensitive information. After recording participants' disclosure choices, we measured participants along a variety of psychometric scales designed to capture

ACM ISBN 978-1-4503-6830-8/19/11…$15.00
https://doi.org/10.1145/3338498.3358656

individual differences in decision making and personality. We examined whether changes in the differences in disclosure rates could be predicted by the measured psychometric variables. We modeled this relationship using logistic regression where the disclosure choice made by participants was our dependent variable.

In Study 1, we presented participants with a hypothetical disclosure choice and a single framing nudge with 'Opt-In' and 'Opt-Out' conditions. We examined the interaction between these conditions and two psychometric variables. Our analysis found a significant main effect for the nudge, but failed to identify significant effects for the *interaction* between the nudge and participants' psychometric traits. In other words: the nudge was effective in influencing participants' disclosure choice, but its effectiveness did not vary with participants' psychometric traits.

Lessons from Study 1 informed the design of Study 2. This study explored a wider scope of nudges and psychometric variables (15 psychometric variables across 2 nudges, across 4 experimental conditions) in a real disclosure setting. In our analysis of Study 2, we found - again - a main effect for the nudges. We additionally identified three potentially significant effects for the interaction between nudge conditions and psychometric variables. Thus, Study 2 suggested that the effectiveness of some nudges could in fact vary with participants' psychometric traits.

Building on the results of Study 2, we conducted a replication study (Study 3) to test the robustness of the effects found in Study 2. Study 3 used the same design and experimental conditions as Study 2. However, our analysis of Study 3 failed to replicate the effects we identified in Study 2.

We interpret the null findings (and in particular the failure to replicate the initial significant findings) as a cautionary tale for both the practical effectiveness of tailored nudges and for future research conducted in this area. The results suggest first, that while one cannot exclude on statistical grounds the theoretical possibility of tailoring nudges to individual differences in decision making and personality, the effects of some of these interventions may be fragile, and potentially impractical for many applications. Second, given the risk of spurious correlations emerging as significant from the interaction of multiple nudges with multiple psychometric variables, research in this area should pay particular attention to replicating and validating results of such interactions.

2 RELATED WORK

Our work builds upon the existing body of research on nudging [40]. As defined by Thaler and Sunstein, nudges are changes in the design or structure of a choice which predictably alter behavior without altering economic incentives. Building upon this idea, studies within the psychology and behavioral economics literature have explored and expanded upon nudge interventions, becoming a popular form of behavioral intervention [18].

2.1 Nudges in Privacy and Security

Within privacy and security, nudges have grown in popularity as a tool through which to help address a domain of problems where users may encounter difficulties in decision making processes. A variety of studies have shown nudges to be effective in changing user behavior [1]. In the security field, nudges have been used to help users make more secure choices. Ur et al. assess the effectiveness of visual password meters in encouraging users to create stronger passwords [42]. Turland et al. evaluate the effects of minor user interface changes on Android devices to encourage users to select more secure wireless networks [41]. In the privacy field, nudges have been used to guide users towards better privacy outcomes. Almuhimedi et al. examine the use of nudges on location disclosure decisions for mobile devices [5]. In a related study, Balebako et al. survey the extent to which nudge interventions are currently employed online and on mobile devices [6]. Each of these studies provide instances of cases where nudge interventions have been applied to guide users to better privacy and security outcomes. These studies inform the design of our nudges throughout our experiments.

While the application of nudges within privacy and security has often been 'one-size-fits-all', recent work has helped classify a variety of nudge types. Acquisti et al. review the application of nudges to problems in privacy and security and describe six categories of nudge interventions: information, presentation, defaults, incentives, reversibility, and timing [1]. This differentiation between nudge types raises the intriguing possibility that nudges may be personalized to users.

2.2 Individual Differences in Decision Making and Personality

Within the psychology literature, inventories have been developed to capture and quantify measures of individual differences in decision making and personality. It stands to reason that, if individuals differ based on personality and decision making traits, such differences may also translate to differential reactions to behavioral interventions across different domains of decision making - including privacy.

Personality inventories such as the well-known "Big Five" scale measure subjective personality traits which have been shown to influence behavior such as job performance [7]. In contrast, decision inventories attempt to measure aspects of users' decision making ability. We utilize both types of measures as a basis for drawing inferences regarding data disclosure decisions.

Multiple inventories have been developed within the psychology literature to measure both decision styles and decision skills. The Need for Cognition (NFC) and General Decision Making Style (GDMS) scales each measure aspects of individual decision making style [11, 35]. While NFC measures users along a single dimension, GDMS captures multiple attributes that describe decision making style.

Decision skill inventories vary from those which measure decision style in that they attempt to capture skills that individuals possess relevant to decision making. One measure of decision skills is the Adult Decision-Making Competence (A-DMC) inventory created by Bruine de Bruin et al. [10]. This inventory measures competence across seven decision skills, assessing competence in a way which the authors show is predictive of decision outcomes. Another common decision skill measure is numeracy [24]. This skill has been shown within the risk communication literature to be associated with risk perception [22, 33]. Users' perceptions of risk may have direct implications for how they approach privacy decisions.

Several of these inventories have been applied to privacy research, where Egelman et al. draw a link between privacy preferences and individual differences in decision making and personality [15]. In addition to identifying decision traits that are predictive of privacy preferences, the authors note the potential of individual differences as a tool to create personalized behavioral interventions for privacy and security decisions.

2.3 Tailored Nudges

In the past few years, studies have explored differential effects of nudges applied to decision problems in a number of domains. Allcot et al. assess the welfare impacts of nudges used within home energy conservation reports [4]. By capturing individual willingness to pay for the home energy reports, along with the magnitude of potential welfare gains, the authors design an algorithm to best target the nudges. In a similar study, Beshears et al. examine nudges for retirement savings plans and find the nudges to be more effective for certain demographics [8].

Within security, several studies have explored applications of tailored nudges that incorporate information about users' psychometric traits. Malkin et al. tests nudges for browser warnings that are personalized to the traits of the GDMS [26]. Although the authors identified several significant correlations between their nudges and the GDMS, they failed to find significant evidence of tailoring.

A later study by Peer et al. examined tailoring for password nudges [31]. In this study, the authors measure participants along multiple psychometric scales and were successful in identifying effects of tailoring for a subset of these. We employ several of the same measures used by the authors within our studies.

Existing work within the privacy literature has examined tailored nudges applied to disclosure decisions. Knijnenburg et al. identified differential effects for different 'justifications' within the context of disclosure rates for data on a hypothetical mobile app [23]. In practice, several of these justifications acted in a similar fashion to nudges. The authors found that characteristics such as gender moderated the effectiveness of some justifications. These findings further point to the potential for privacy nudges to be tailored to the individual traits of users. In a separate study, Coventry et al. examined the effects of a single nudge on cookie acceptance for web browsers [12]. The authors measure several personality traits, but fail to find evidence that the traits moderate the effectiveness of the nudge. We expand upon these works by focusing on multiple different psychometric traits as a basis for tailoring privacy nudges. Within the context of privacy, we are one of the first to explore the application of nudges tailored to individual differences in decision making and personality.

3 STUDY DESIGN OVERVIEW

Our examination of tailored privacy nudges takes place within the context of data disclosure decisions. To facilitate this, each of our studies begins by presenting participants with a hypothetical or real disclosure choice and asking them to make a 'Yes' or 'No' decision on whether or not they wish to disclose their data. We insert nudges into these scenarios by modifying the choice text to elicit differences in disclosure rates between groups of participants. We employ two types of nudges across our three studies: 'framing' nudges and 'social norms' nudges. The first of these modifies the choice text to leverage framing effects, while the second introduces additional information to the choice to establish social norms which might affect participants' disclosure behavior. In Study 1, we test only the framing nudge. In Study 2 and Study 3, we test both the framing and social norms nudges.

For each nudge, we create two different wordings or 'variants' of the nudge text. Relative to each other, these variants create the desired psychological effects we wish to leverage within our nudges. For example, we create the framing nudge for Study 1 by asking one group of participants whether they wish to 'Opt-In' to a service while asking the other group whether they wish to 'Opt-Out'. The framing effect exists only in the contrast between the two wordings. We apply a similar method to create the social norms nudge used in Study 2 and Study 3. We treat each nudge 'variant' as an experimental condition within our study. Participants are randomly assigned to a single nudge variant in a between-subjects design. For Study 1, this translates into two experimental conditions across one nudge. For Study 2 and Study 3, this translates into four experimental conditions across two nudges. Across the three studies, participants are only assigned to one disclosure decision that contains a single nudge variant. We provide additional details on the construction of the nudges and the nudge conditions in the study descriptions below.

Following the decision task, we ask all participants questions from psychometric inventories designed to measure their decision making and personality traits. We selected inventories for our studies that either relate to the cognitive effect being leveraged in the nudge or have been examined in other studies exploring psychometrically targeted nudges. In Study 1, we measure two scales from the A-DMC: Resistance to Framing and Applying Decision Rules [10]. For Study 2 and Study 3, we measure Resistance to Framing and Recognition of Social Norms (from the A-DMC) along with scales to capture Scientific Reasoning [14], Need for Cognition [11], Numeracy [24], General Decision Making Style [36], and the Big Five personality traits [7]. Additional details on each of these measures are provided in the study descriptions below. After answering questions from the psychometric inventories, participants conclude the study by answering questions on demographics and privacy attitudes.

4 STUDY 1

With Study 1, our goal was to identify whether psychometric variables could be used to infer differences in disclosure rates for a single nudge. Because tailoring privacy nudges implies the ability to select one nudge from multiple options, we planned to conduct further experimentation with multiple nudges in later studies. We do this in Study 2 and Study 3.

4.1 Study Design

4.1.1 Disclosure Scenario. We built the disclosure scenario for Study 1 around a hypothetical IoT service called 'Auto-Checkout'. This hypothetical IoT scenario was developed over several iterations and consisted of an automated checkout service that would allow users to bypass checkout lines in exchange for their indoor location information. Participants were given details of the potential benefits and drawbacks of the service, along with the choice to enroll or not with 'Yes' and 'No' response options. The details of the scenario along with the choice were presented to users on an interactive phone screen mock-up.[1] While the service would allow for faster trips to the grocery store, participants were informed that their data would be used to enable personalized advertising. [2]

4.1.2 Nudge Design. Within the choice text of the scenario described above, we embedded a nudge designed to leverage framing effects [21]. These manipulations are frequently used within the social science literature and typically yield strong effects [13]. Because a disclosure choice does not contain a natural 'baseline' condition against which to measure the effectiveness of our nudge, we create two 'variants' of our nudge ('Opt-In' and 'Opt-Out') with which we could measure differences in disclosure rates. We use the 'Opt-In' and 'Opt-Out' variants of the nudge as our experimental conditions. When participants were presented with the scenario, they were randomly assigned to receive either the 'Opt-In' or 'Opt-Out' variant of the choice text. These nudges manipulate the text of the disclosure question to change the 'default' state of disclosure from the perspective of the user. The text for both the 'Opt-In' and 'Opt-Out' conditions is given below.

"Do you wish to **OPT-IN** to Auto-Checkout?"

"Do you wish to **OPT-OUT** of Auto-Checkout?"

4.1.3 Psychometric Variable Selection. Following the disclosure choice, participants were presented with questions designed to measure different psychometric traits. We selected the Resistance to Framing and Applying Decision Rules scales from the A-DMC for the purposes of Study 1 [9, 10]. While the full A-DMC inventory contains questions to measure individual decision making ability across seven psychometric variables, the two we selected most closely measure participants' ability to recognize the cognitive bias leveraged by our framing nudge. These scales were also shown by the authors to be highly intercorrelated. When applied to a framing nudge, we might expect the nudge to have less of an effect on participants with higher Resistance to Framing and Applying Decision Rules scores compared to participants with lower scores. For the Resistance to Framing scale, the questions are separated into two identical sets which vary only in the framing of the text. These sets must be separated by an intermediary task or time interval. Because of this, in addition to a desire to reduce potential participant fatigue, we split our survey into two sections.[3]

To protect against participants randomly selecting answers for the survey questions, we designed attention check questions which we placed within the resistance to framing questions in both the initial and followup surveys. These checks are written in the style of the surrounding questions, but contain a strictly dominated choice. Participants that are paying attention should select this answer. By using these attention checks, we can filter out participants that do not honestly complete the survey (amounting to random noise within the data) while checking for comprehension of the task.

We collected basic demographic information from participants to serve as additional controls in our regressions analysis. Of potential interest are age and level of education. Either of these variables may capture a higher level of technological literacy which may help explain disclosure choice [30]. To capture level of education, participants are asked to list their highest attained degree. These responses are translated into years of education within the analysis to better facilitate regression. This translation is conducted by assigning year equivalents to each level of education between 12 and 20 years.

We additionally measured participants' scores along the Concern for Information Privacy (CFIP) scale to capture general attitudes regarding data collection and storage [37]. This variable is used as a control within our regression models.

4.1.4 Survey Structure. The survey for Study 1 was administered in two parts. In the initial survey, participants were presented with a hypothetical IoT disclosure choice followed by the psychometric variable scales for Resistance to Framing and Applying Decision Rules. The initial survey concluded with demographic questions.

After several days, participants were invited back to complete the followup survey in which participants answered additional questions from the Resistance to Framing scale and questions from the CFIP scale. We placed the IoT disclosure question at the front of the initial survey and the CFIP questions at the end of the followup. By presenting the questions in this order, we avoided priming participants with context which may have influenced their disclosure choice.

4.2 Results

4.2.1 Sample Demographics. We recruited a sample of 200 participants from Mechanical Turk to complete the initial survey. Although samples drawn from Mechanical Turk are not representative, numerous behavioral effects (including framing effects) have been replicated on the platform [29]. Of the participants recruited into the study, 176 completed the initial survey without incorrectly answering any of the attention check questions. Several days after the initial survey, this subset of participants was invited back to complete the followup survey. The response rate to this second stage survey was 81%, yielding a sample of 145 participants. Of these, only 2 participants failed to answer the attention check questions correctly. After excluding these participants, our analysis was conducted on a final sample of 143 participants.[4] Of this final sample, 60% (86) were male and the median age was 34. All participants had at least a high school degree, with 48% (69) having achieved a bachelors degree or higher.

4.2.2 Summary Statistics. Across the two experimental conditions, 55% (78) of participants chose to allow the data disclosure within the hypothetical IoT scenario. We measured the size of the framing effect by comparing the likelihood of participants to allow the data disclosure between nudge conditions. Of the 143 participants considered in our analysis, 64% (46) of those assigned to the 'Opt-In'

[1]Survey participants were provided with instructions on how to use the phone screen mock-ups prior to completing the disclosure task. These instructions and question format were tested in a pilot study prior to the experiment.

[2]The mock-ups, along with the full text of the disclosure scenario, are available in Appendix A.1.

[3]Sample questions from the Resistance to Framing scale are available in Appendix B.1.

[4]Our analysis for each study was conducted in R using publicly available packages for regression analysis and data presentation [16, 17, 20, 45, 47].

condition chose to allow the disclosure while only 44% (32) assigned to the 'Opt-Out' condition chose to do the same. The difference in disclosure rates between these conditions is significant with a p-value of $p = 0.02291$ when tested using Chi-squared, suggesting a significant main effect of the nudge.

We additionally measured the means and standard deviations for the two psychometric variables and the privacy preferences scale. The mean score for Resistance to Framing within the sample was 3.865 with a standard deviation of 0.403. The possible range of scores on this scale is 1 to 6. This result suggests a sample which is slightly resistant to framing effects overall. For the Applying Decision Rules scores, the sample mean was 0.615 with a standard deviation of 0.246. The range of possible scores for this variable is 0 to 1 where higher scores indicate greater ability to apply rules to decision problems. The sample mean for privacy attitudes as measured by the CFIP scale was 5.887 with a standard deviation of 0.984. The range of possible scores for this scale is 1 to 7. These scores potentially indicate a higher privacy sensitivity within the sample.

4.2.3 Logistic Regression. We used logistic regression models to conduct a difference-in-differences analysis of the response data. In each regression, the coefficient of interest to our research question was the interaction between the psychometric variable scores and the nudge condition assigned to the participant. This allowed us to see whether the effect of the assigned nudge on disclosure likelihood varied significantly with the psychometric variables.

In total, we examined two sets of models – one for each psychometric variable (Resistance to Framing and Applying Decision Rules). For each of the psychometric variables, we examined four regressions models, varying the number of control variables in each one. In the first model, we assessed the interaction between the framing nudge and psychometric variable score alone. In the subsequent models, we incrementally introduced controls for privacy attitudes (CFIP) and demographics (age, gender, and education). We represent the assigned experimental condition (whether participants were shown the 'Opt-In' or 'Opt-Out' variant of the nudge) in our models as a binary coded dummy variable (Opt-In Condition).

We used Generalized Additive Models (GAMs) to test for omitted transformations on the continuous regressors. This analysis revealed a non-linear transformation on the ADR psychometric variable. To best satisfy the 'linearity in parameters' assumption of the logistic regression model, this score variable was log transformed. A subsequent test showed the Log ADR score to be more linear. We use the Log ADR score throughout the remainder of our analysis.

Table 1 presents the regression coefficients expressed in log-odds along with standard errors for the models containing all controls for each psychometric variable.[5] Each column of the table presents the regression model for a separate psychometric variable. Overall, we did not identify significant effects for either of the interaction terms in our models.

While none of the interaction coefficients we examined were significant, the direction of the coefficients make intuitive sense. For both the Resistance to Framing and Log-ADR variable scores, the sign of the interaction coefficient is negative, indicating that a higher decision competence for these variables translates into a decrease in the effectiveness of the framing nudge on disclosure behavior.

[5]Regression models with incremental levels of controls for Study 1 are available in Appendix D.1. These models did not yield significant effects for the interaction terms.

Table 1: Logistic Regression Coefficients for Study 1

	Dependent variable:	
	disclosure	
	(1)	(2)
Opt-In Condition (OC)	3.909	0.745
	(3.496)	(0.560)
Framing Score (FR)	0.266	
	(0.607)	
Log-ADR Score (ADR)		1.001*
		(0.565)
CFIP Score	−0.238	−0.143
	(0.192)	(0.190)
Age	0.007	0.006
	(0.019)	(0.019)
Female	0.391	0.493
	(0.378)	(0.383)
Years of Education	0.141	0.076
	(0.093)	(0.096)
AC*FR	−0.768	
	(0.898)	
AC*ADR		−0.250
		(0.787)
Constant	−2.352	−0.381
	(2.835)	(1.818)
Observations	143	143
Log Likelihood	−92.626	−90.389
Akaike Inf. Crit.	201.252	196.778
Note:	*p<0.1; **p<0.05; ***p<0.01	

Although we did not identify significant effects for the interactions between the two psychometric variables and nudge conditions, our results may have been limited by factors including our sample size and the hypothetical nature of the disclosure scenario. Additionally, limitations within our nudge design and selection of psychometric variables may have impacted our results. We attempt to address these potential limitations and build upon the design of Study 1 within Study 2 and Study 3.

5 STUDY 2

Although Study 1 did not yield significant findings, the direction of our regression coefficients showed promise. In Study 2, we sought to determine whether the null result observed in Study 1 was robust by focusing on the potential limitations of Study 1 that may have contributed to our null results. These potential limitations include the design of our nudges and disclosure choices along with our selection of psychometric variables. If we selected the wrong nudges or psychometric variables, this may have lead to a false negative. In Study 2, we attempt to minimize these factors.

We addressed these issues in part by switching from hypothetical to real disclosure choices and including an additional nudge within our study design. We additionally expanded upon the selection of psychometric variables to capture a broader range of psychometric traits. Because of the high risk of identifying spurious correlations in our analysis within Study 2, we made the decision prior to collecting data to follow up Study 2 with a subsequent replication study (Study 3), which would attempt to confirm any effects possibly identified in Study 2.

5.1 Study Design

5.1.1 Disclosure Scenario. Our use of a hypothetical disclosure scenario in Study 1 potentially limited the size of the effect of our framing nudge. Although we iterated on the design of the scenario to make it more realistic, the hypothetical scenario lacked the perception of risk that users might face when making real disclosure decisions. Research in the privacy literature indicates that this may impact how users make privacy choices [28][3]. Users may respond differently to a real disclosure decision than they would to a hypothetical one. In the context of our study design, asking participants to make a real disclosure decision may yield stronger main effects for our nudges compared to a hypothetical decision. For Study 2, we employ deception to construct a scenario and disclosure question with real perceived risks by participants.

The design of our disclosure question and experimental manipulation for Study 2 is based off the design used by Samat et al. in a study of framing effects under different levels of risk [34]. In that study, the authors asked participants to answer a series of 'ethical behavior questions' drawn from Acquisti et al. [2]. Those questions asked participants how frequently they engaged in a variety of behaviors of varying degrees of sensitivity such as "Have you ever had a one-night stand". Before answering the questions, participants were presented with the ostensible choice to decide whether or not they wished to share their responses to the ethical behavior questions with a third-party. Samat et al. manipulated the identity of the third-party to raise or lower the perceived risk of the disclosure to participants. In the high-risk condition, participants were told that their responses would be shared with other users on Mechanical Turk community forums. Using this design, the authors found that the high-risk condition yielded the largest framing effect. No matter participants' answers to the disclosure choice, their responses were not actually shared.

In an early iteration of our design for Study 2, we combined the high-risk disclosure condition and framing nudge of Samat et al. with an expanded selection of psychometric variables. A pilot of this design on a small sample revealed that participants were skeptical of the deception. Several participants questioned why researchers would want to share their responses to ethical behavior questions with other users on Mechanical Turk community forums. From this feedback, we changed the third-party within the scenario from users of Mechanical Turk community forums to researches at other universities. A test of this deception yielded better results. Our final version of the disclosure scenario for Study 2 includes this change along with additional text designed to increase the perceived sense of risk for disclosure.

5.1.2 Nudge Design. We created two nudges to embed within the disclosure choice of our deceptive scenario. As in Study 1, these nudges each contain two variants that enable us to measure a relative difference in disclosure rates. This yielded a total of four experimental conditions spread across two nudges in a between-subjects design. Because of its strong effects within the literature and in Study 1, we again used a framing nudge within Study 2. The design of our framing nudge mimics that used by Samat et al. and is shown below.

> "**Allow** your responses to the ethical behavior questions to be shared with researchers outside of our study team, along with your Mechanical Turk ID?"

> "**Prohibit** your responses to the ethical behavior questions from being shared with researchers outside of our study team, along with your Mechanical Turk ID?"

The above 'Allow' and 'Prohibit' variants differ only in their presentation of the disclosure choice. Participants assigned to one of these conditions were then presented with the response options 'Yes' and 'No'. We randomized the order of these response options to avoid any potential ordering effects.

Research from the psychology literature on social norms has shown them to have some of the strongest effects among different nudges [39]. We constructed a second nudge that leverages those effects to create differences in disclosure rates similar to the framing nudge. This nudge consists of 'High Norms' and 'Low Norms' conditions that establish a social norm regarding the percentage of other participants that chose to disclose their responses. Unlike the framing nudge, the text for this nudge is appended to the end of the description of the disclosure scenario. We tested several versions of this nudge before arriving at a version that consistently yielded strong effects. The text of the experimental manipulations is shown below.

> "In our past studies, **73%** of participants chose to allow their responses to be shared with researchers outside of our study team."

> "In our past studies, **31%** of participants chose to allow their responses to be shared with researchers outside of our study team."

Because the disclosure question no longer contained the experimental manipulation, both the question and the response options for the social norms nudge conditions remained fixed. As for the conditions associated with the framing nudge, we randomized the order of the response options to avoid potential ordering effects. The full text of each experimental condition – including the text of the disclosure scenario and manipulations – is provided in Appendix A.2.

5.1.3 Psychometric Variable Selection. We expanded our selection of psychometric variables in Study 2 to cover a broader range of cognitive traits than in Study 1. In total, we selected 15 psychometric variables from the decision science and psychology literatures that measure a variety of decision making and personality characteristics. This blend of variables used scales that capture both objective skill measures (such as the scales of the A-DMC) and self-reported personality measures.

Due to its direct applicability to the framing nudge, we retained the Resistance to Framing scale that we used in Study 1. To this we added scales to measure Recognition of Social Norms (SN) [10], Scientific Reasoning (SRS) [14], Need for Cognition (NFC) [11], Numeracy (NUM) [24], General Decision Making Style (GDMS) [36], and the Big Five personality traits (B5) [7]. These additional scales were included due to either their relevance to the nudge being tested or their previous use in related studies on personalized nudges in security [26][31]. In the case of the Recognition of Social Norms scale, we expected these scores to have a strong relationship with the social norms nudge – where higher scores translate to greater effectiveness of the nudge on impacting disclosure rates.

Each of the scales with the exception of Numeracy consisted of multiple choice questions. The Big Five measures five personality traits that we consider independently. These are extraversion, agreeability, conscientiousness, neuroticism, and openness. Likewise, the GDMS identifies five decision making styles. These are rational, intuitive, dependent, avoidant, and spontaneous. We also consider these independently within our analysis. Both the resistance to framing and recognition of social norms scales are split into two sets of related questions that were administered to participants at different times.[6] In addition to the psychometric variables, we collect demographic information and participants' scores along the Internet Users' Information Privacy Concerns (IUIPC) scale [25] to use as controls within our regression models. We selected the IUIPC scale over the CFIP scale from Study 1 due to its wider use within the privacy literature.

5.1.4 Survey Structure. As in Study 1, the survey was administered in two parts to reduce the cognitive burden on participants and to allow for time separation as required by some of the psychometric variable measures. In the initial survey, participants completed the disclosure choice, ethical behavior questions, need for cognition scale, numeracy scale, and scientific reasoning scale. Participants also completed the initial sections of the scales designed to measure resistance to framing effects and recognition of social norms. The following day, participants were invited back to complete a followup survey which contained the second halves of the resistance to framing and recognition of social norms scales. In addition, the followup survey contained the measures for GDMS, Big Five, and IUIPC.

5.2 Results

5.2.1 Sample Demographics. We recruited 1,200 participants to complete Study 2 from Mechanical Turk. Of those recruited, 1,198 completed the initial survey and 966 completed the followup survey and passed the minimum threshold for attention check questions.[7] Of the 1,198 that completed the initial survey, 50% were male and the median age was 35. 58% (700) had completed a Bachelor's degree or higher. To facilitate our exploratory analysis, we split the sample by survey. Because the psychometric variable scales were split between the two surveys (and because the experimental manipulation was contained in the initial survey), several psychometric variables can be examined without needing data from the followup survey. By splitting our sample into a 'complete' sample (all 1,198 participants) and a 'partial' sample (the 966 participants that completed both surveys), we can take advantage of the larger sample size for a subset of the psychometric variables.

5.2.2 Summary Statistics. For both the framing and social norms nudges, we observed significant differences in the disclosure rates between conditions on the complete sample. Of the 597 participants that were assigned to the framing nudge, 66% (197) of those assigned to the 'allow' condition chose to disclose their responses while 39% (115) of those assigned to the 'prohibit' condition chose to disclose their responses. For the 601 assigned to the social norms nudge, 65% (197) of those assigned to the 'high norms' condition chose to disclose their responses while 47% (140) of those assigned to the 'low norms' condition chose to disclose their responses. Both differences in disclosure rates are significant when tested using Chi-squared. The differences between the conditions of the framing and social norms nudges are significant with p-values of $p = 7.584e-11$ and $p = 5.199e-06$ respectively. We observed similar difference in disclosure rates and significance levels for the partial sample.

Table 2 shows summary statistics for the psychometric variables and privacy concerns measured in Study 2 – including the mean, standard deviations, and ranges of possible scores for each variable. Higher scores for these variables translate to higher competence or greater intensity for the corresponding decision making and personality traits.

Table 2: Psychometric Variable Summary Statistics for Study 2

Variable	Mean	St. Dev.	Range
Scientific Reasoning	6.381	2.718	0 – 11
Need for Cognition	4.611	1.250	1 – 7
Numeracy	8.171	2.033	0 – 11
Resistance to Framing	4.924	0.486	1 – 6
Recognition of Social Norms	0.471	0.271	-1 – 1
GDMS Rational	3.888	0.674	1 – 5
GDMS Intuitive	3.293	0.904	1 – 5
GDMS Dependent	3.192	0.855	1 – 5
GDMS Avoidant	2.479	1.060	1 – 5
GDMS Spontaneous	2.510	0.890	1 – 5
Big 5 Extraversion	2.893	0.946	1 – 5
Big 5 Agreeableness	3.710	0.755	1 – 5
Big 5 Conscientiousness	3.918	0.771	1 – 5
Big 5 Neuroticism	2.668	0.970	1 – 5
Big 5 Openness	3.717	0.725	1 – 5
IUIPC	5.879	0.911	1 – 7

[6]Sample questions from the Recognition of Social Norms scale are available in Appendix B.2

[7]We developed 8 attention check questions for Study 2 that we placed throughout our survey. These attention check questions were designed to measure task comprehension and presented participants with a task in the style of surrounding questions that contained a strictly dominated answer. We used our attention check questions as a robustness check on the results of our analysis. While we present the results of participants that answered at least 1 attention check question correctly (which excludes the fewest participants), the results are consistent for higher numbers of attention check questions.

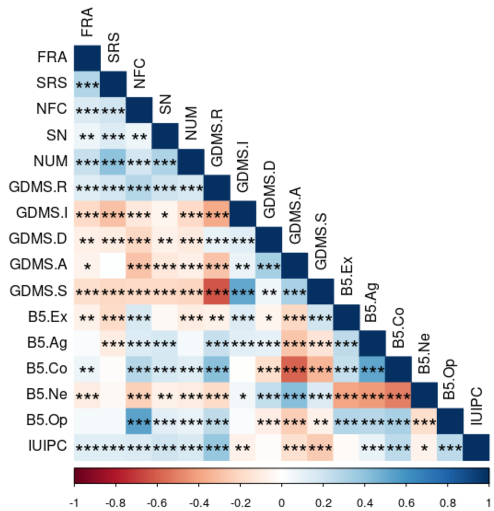

Figure 1: Correlation Strength and Significance for Psychometric Variables
Note: ***p<0.1; **p<0.05; ***p<0.01**

We computed the correlation strength and significance for each pairing of psychometric variables to assess the coverage of cognitive traits within our study. This correlation table is displayed in Figure 1. With the exception of two psychometric variables within the GDMS and Big Five, most of the variables are only weakly correlated with each other. This suggests that the selected scales are measuring a broad range of cognitive traits, with little overlap.

5.2.3 Logistic Regression. We tested multiple logistic regression models as part of our exploratory analysis to examine the relationship between disclosure choice, our nudges, and the measured psychometric variables. For each model, our dependent variable was participants' disclosure choice. Our explanatory variables were the assigned nudge condition and the relevant psychometric variable score. Across the two nudges and 15 psychometric variables, we constructed 30 sets of regression models. We used the IUIPC score and demographic variables as controls within our regressions. The interaction term between the assigned nudge condition and psychometric variable score is the primary coefficient of interest to our research question. As in Study 1, we represent the assigned experimental condition as a binary coded dummy variable ('Allow Condition' for models with the framing nudge and 'High Norms Condition' for models with the social norms nudge).

Out of the regression models that we created, we identified 3 pairs of psychometric variables and nudge conditions with potentially significant interaction terms. These included the Recognizing Social Norms score and social norms nudge, the Big Five Extraversion score and framing nudge, and the Big Five Conscientiousness score and framing nudge. The coefficients for these regression models expressed in log-odds along with the standard errors for each are presented in Table 3.[8]

Column 1 of Table 3 shows the regression model with control variables for the pair of the social norms nudge and Recognizing Social Norms score. Columns 2 and 3 show the regression models with controls for the framing nudge paired with the Big 5 Extraversion and Conscientiousness scores respectively.

Of the potentially significant effects, the interaction between the Recognizing Social Norms score and the social norms nudge is the most intuitive and has the strongest effect. The direction of the coefficient implies that as a participant is more likely to recognize social norms, they may be more likely to be influenced by a social norms nudge. While less intuitive, the effects between the two Big Five traits and the framing nudge can also be interpreted. The direction of the regression coefficient for Big Five Extraversion score and the framing nudge condition suggests that those who are more extroverted may be more likely to be influenced by a framing nudge. Likewise, the direction of the coefficient for the Big Five Conscientiousness score and framing nudge condition suggest that participants who are less conscientious may be more likely to be influenced by a framing nudge.

We employed the same statistical methods as in Study 1 to test for misspecification of the logistic regression models. We used generalized additive models to detect omitted transformations of the independent variables. This analysis indicated the each of the independent variables was roughly linear.

6 STUDY 3

Our goal in Study 2 was to determine whether the null effects we observed in Study 1 were robust or rather due to limitations in our study design. While we cannot prove the absence of an effect, we sought to minimize the likelihood of a false negative in Study 2 by expanding our selection of nudge types and psychometric variables. We found three potentially significant interactions between nudge conditions and psychometric variables. However, by testing additional nudges and psychometric scales, we risk identifying spurious correlations. Study 3 builds upon the results of Study 2 by attempting to replicate the three potentially significant effects we identified using a separate sample.

6.1 Study Design

For this study, we made only minimal changes to the study design and survey structure used in Study 2. Our primary change was to condense the survey questions into a single instrument by removing the scales for the psychometric variables we did not wish to test. Condensing the survey removed the need for a separate followup survey and allowed participants to complete the study in one session. In total, the survey used for Study 3 kept the same disclosure choice and nudge manipulations while measuring recognition of social norms, the Big 5 personality traits, and the IUIPC. Because many of our attention check questions from Study 2 were embedded within the scales for the other psychometric variables not considered in Study 3, we kept some of the questions from the other scales to preserve these attention checks. We additionally expanded our sample size to 2,000 participants (1,000 per nudge). This decision was informed in part by a post-hoc power analysis

[8]Results for the regression models using incremental levels of controls for the three potentially significant psychometric variables are available in Appendix D.2. Regression results for the non-significant pairs of psychometric variables and nudge conditions are available from the authors upon request.

Table 3: Logistic Regression Coefficients for Study 2

	Dependent variable:		
	disc		
	(1)	(2)	(3)
Recognizing Social Norms	-1.232^{**}		
	(0.574)		
Big 5 Extraversion		-0.056	
		(0.140)	
Big 5 Conscientiousness			0.258
			(0.193)
High Norms Condition	-0.322		
	(0.420)		
Allow Condition		-0.398	2.849^{***}
		(0.616)	(1.061)
IUIPC	-0.527^{***}	-0.485^{***}	-0.485^{***}
	(0.128)	(0.122)	(0.123)
Age	-0.021^{**}	-0.008	-0.007
	(0.010)	(0.009)	(0.009)
Female	-0.083	0.200	0.157
	(0.208)	(0.203)	(0.201)
African American	-0.012	0.162	0.171
	(0.363)	(0.369)	(0.368)
Hispanic	0.222	0.027	-0.052
	(0.405)	(0.365)	(0.366)
Asian	0.489	0.451	0.355
	(0.493)	(0.472)	(0.469)
Other Race	0.147	0.130	0.094
	(0.621)	(0.648)	(0.647)
High School	-12.697		
	(535.411)		
Associate Degree	-12.650	0.166	0.116
	(535.411)	(0.327)	(0.327)
Bachelor's Degree	-12.768	0.274	0.283
	(535.411)	(0.302)	(0.302)
Advanced Degree	-13.118	-0.086	-0.013
	(535.411)	(0.391)	(0.387)
Other Education		13.218	13.081
		(535.411)	(535.411)
Recognizing Social Norms*High Norms Condition	2.432^{***}		
	(0.784)		
Big 5 Extraversion*Allow Condition		0.487^{**}	
		(0.204)	
Big 5 Conscientiousness*Allow Condition			-0.464^{*}
			(0.262)
Constant	16.968	2.615^{***}	1.425
	(535.411)	(0.864)	(0.963)
Observations	453	482	482
Log Likelihood	-281.000	-301.875	-304.746
Akaike Inf. Crit.	591.999	633.749	639.492
Note:	*p<0.1; **p<0.05; ***p<0.01		

conducted using the observed effects sizes from Study 2.[9] We preregistered this study design along with our analysis plan on Open Science Framework prior to collecting data.[10]

6.2 Results

6.2.1 Sample Demographics. We recruited 2,000 participants to complete our replication study using Mechanical Turk. During recruitment, we excluded participants that had previously participated in Study 1 or Study 2. Out of this total, 1,996 participants provided complete responses and passed the minimum threshold for attention check questions. The makeup of this sample was 50% male with a median age of 36. 61% (1,222) had completed a Bachelor's degree or higher. Because we were able to condense our psychometric variable scales into a single survey, we did not split our sample as in Study 2.

6.2.2 Summary Statistics. We observed similar effect sizes on both the framing and social norms nudges compared to Study 2. For the framing nudge, 59% (293) of participants assigned to the 'allow' condition chose to share their responses to the ethical behavior questions while only 33% (165) of those assigned to the 'prohibit' condition chose to do the same. For the social norms nudge, 65% (321) of those assigned to the 'high norms' condition chose to disclose their responses while only 47% (239) of those in the 'low norms' condition shared their responses. Using a Chi-Squared test, both differences were, as in Study 2, significant at $p = 6.051e - 16$ and $p = 6.432e - 8$, respectively.

We additionally examined the means and standard deviations for the Recognizing Social Norms, Big Five Extraversion, and Big Five Conscientiousness psychometric variables. For the Recognizing Social Norms variable, we observed a sample mean of 0.474 and a standard deviation of 0.294. On the Big 5 Extraversion and Conscientiousness scales, we observed sample means of 2.874 and 3.829 with standard deviations of 0.864 and 0.773 respectively. For the IUIPC scale, we observed a mean privacy concern score of 5.767 with a standard deviation of 1.033. Many of these statistics were similar to the values observed in Study 2. Using a two-sample t-test, we did not find significant differences in the means of the three psychometric variables between Study 2 and Study 3. This suggests that the scores for these three variables observed in Study 2 are consistent with those in Study 3.

6.2.3 Logistic Regression. We tested the same logistic regression models that we used in Study 2 for the three potentially significant pairs of psychometric variables and nudge conditions (Recognizing Social Norms score and social norms nudge, the Big Five Extraversion score and framing nudge, and the Big Five Conscientiousness score and framing nudge). Overall, we failed to replicate the effects from Study 2. For the three logistic models, the interaction terms were either not significant, weakly significant (at the $p < 0.1$ level), or lost significance when controls were added. The coefficients for these regression models expressed in log-odds along with the standard errors for each are presented in Table 4.[11]

Whereas the interaction term for the Recognizing Social Norms score and social norms nudge pair was the strongest of the three effects in Study 2, the same effect was only significant at the $p < 0.1$ level when control variables were excluded from the regression in Study 3. When the control variables were added, the coefficient for the interaction term was no longer significant. For the Big Five Conscientiousness score and framing nudge pair, the direction of the interaction coefficient changed from negative to positive. Again, the interaction term for this variable and condition pair was only weakly significant for the regression model without control variables. When control variables were added, this term lost significance.

Only the significance of the interaction term for the Big Five Extraversion score and framing nudge pair remained comparable from Study 2 to Study 3. However, it did so only at the $p < 0.1$ level. While the interaction term was significant at the $p < 0.05$ level when no control variables were included, the significance level diminished when they were reintroduced.

To explore these effects in more detail, we conducted a secondary exploratory analysis where we split each of the three psychometric variables into tertiles. This allowed us to compare participants who scored in the bottom third for a particular variable to those in the middle and top thirds within our regression models. By splitting the data, we can potentially detect more detailed effects within the interaction. This analysis revealed significant effects for the interaction term between the middle and bottom tertiles for the Big Five Extraversion score and framing nudge pair at the $p < 0.01$ level. Although this is potentially highly significant, it only encompasses a portion of the range of potential Big Five Extraversion scores and indicates that the effect may be fragile. [12]

6.2.4 Johnson-Neyman Analysis. In addition to our regression analysis, we examined our data for Study 2 and Study 3 using the Johnson-Neyman technique. This technique has been previously used in the literature to examine the effectiveness of tailored nudges on security behavior [31]. This method splits apart interaction terms for linear models to help identify "regions of significance" within moderator variables [19]. Applied to our study, the Johnson-Neyman technique can help identify ranges within the psychometric variable scores where the interaction with the nudge condition is significant at the $p < 0.05$ level. While this method is potentially useful for identifying effects overlooked by traditional regression models, splitting apart the interaction term in this way increases the risk of identifying spurious correlations. When applied to the three potentially significant effects identified in Study 2, the Johnson-Neyman technique did identify regions of significance for all three. However, and importantly, we were unable to replicate these results for Study 3. In this replication, the regions identified by the Johnson-Neyman technique in Study 2 were either different from or non-existent in Study 3. These results are consistent with the results of our logistic regression analysis and indicate that, if present, the effects of psychometrically tailored privacy nudges are fragile.

[9] The results of this power analysis are available in Appendix C.
[10] https://osf.io/nfyma
[11] Results for the regression models using incremental levels of controls are available in Appendix D.3.
[12] Regression results for the tertiles analysis are available from the authors upon request.

Table 4: Logistic Regression Coefficients for Study 3

	Dependent variable:		
	disc		
	(1)	(2)	(3)
Recognizing Social Norms	-1.054^{***}		
	(0.377)		
Big 5 Extraversion		-0.062	
		(0.117)	
Big 5 Conscientiousness			-0.141
			(0.130)
High Norms Condition	0.363		
	(0.300)		
Allow Condition		0.283	0.364
		(0.475)	(0.692)
IUIPC	-0.953^{***}	-0.500^{***}	-0.496^{***}
	(0.101)	(0.071)	(0.072)
Age	0.011^{*}	-0.010	-0.009
	(0.006)	(0.006)	(0.006)
Female	-0.303^{**}	-0.016	-0.015
	(0.148)	(0.140)	(0.140)
Non-Binary	0.859	1.206	1.273
	(1.250)	(1.223)	(1.202)
African American	-0.414	-0.350	-0.279
	(0.265)	(0.237)	(0.233)
Hispanic	0.096	-0.107	-0.091
	(0.320)	(0.285)	(0.284)
Asian	-0.825^{**}	-0.088	-0.084
	(0.325)	(0.271)	(0.271)
Other Race	0.144	-0.476	-0.472
	(0.490)	(0.669)	(0.671)
High School	0.590	1.509	1.396
	(1.038)	(0.985)	(0.982)
Associate Degree	0.513	1.826^{*}	1.707^{*}
	(1.020)	(0.974)	(0.971)
Bachelor's Degree	0.374	1.354	1.267
	(1.016)	(0.969)	(0.966)
Advanced Degree	0.518	1.063	0.967
	(1.029)	(0.979)	(0.976)
Recognizing Social Norms*High Norms Condition	0.818		
	(0.534)		
Big 5 Extraversion*Allow Condition		0.287^{*}	
		(0.159)	
Big 5 Conscientiousness*Allow Condition			0.194
			(0.177)
Constant	5.367^{***}	1.331	1.702
	(1.160)	(1.093)	(1.118)
Observations	936	992	992
Log Likelihood	-551.202	-607.651	-609.227
Akaike Inf. Crit.	1,134.404	1,247.302	1,250.455
Note:	*p<0.1; **p<0.05; ***p<0.01		

7 DISCUSSION

Overall, the results of our three studies indicate that effects for tailored privacy nudges are difficult to identify with consistency. Although Study 2 identified three potentially significant effects, our replication of these effects in Study 3 found them to be either fragile or non-existent. This result suggests that tailored privacy nudges at the scale of our studies may not be practical in application.

We should note that the results observed in our three studies do not prove the lack of an effect for tailored privacy nudges. It is possible that a study with stronger main effects for the nudges, different psychometric variables, or larger sample sizes may find evidence for effects of psychometric tailoring on privacy nudges. Each of our study designs attempted to optimize these three dimensions to minimize the likelihood of a false negative. We constructed our nudges based on two of the strongest cognitive biases identified within the nudging literature [39]. Likewise, our selection of psychometric variables reflected a broad scope of cognitive traits and reflected the variables used by similar studies on tailored security nudges [26][31]. Our selected sample sizes were informed by power analyses using estimated effects sizes.

It is also possible that different results may be observed with a different sample population. Stewart et al. find the population of participants on Mechanical Turk available for behavioral research at any time to be relatively small (less than 10,000) [38]. Combined with a slow rate of turnover, this suggests that participants may become habituated to behavioral research over time. Participants that are repeatedly exposed to questions from the same psychometric scales may come to learn the 'correct' answers (particularly for scales that measure decision making ability). Peer et al. examine several alternative crowdsourcing platforms to Mechanical Turk, finding them to be less conditioned to behavioral research [32]. For Study 1, this conditioning effect may have contributed to the low variance we observed in scores for psychometric variables such as resistance to framing. In this way, samples drawn from crowdsourcing platforms may be different from representative samples of the general population.

While our studies employed large sample sizes, it is still possible that effects for tailored nudging could be identified using larger samples. However, as the sample sizes required to detect the effects from psychometric variables on nudges increase, tailored nudges become impractical from the perspective of many applications. The combination of large sample requirements and potentially fragile effects may make tailored nudges feasible only to organizations with vast amounts of user data such as Facebook or Google.

The nature of our results further emphasizes the importance of replication studies in behavioral research. When testing for interaction effects between multiple moderating variables, the risk increases of identifying spurious correlations with high significance. Future research should take care to validate such results.

8 CONCLUSION

We conducted three online studies with the goal of identifying effects that would indicate whether privacy nudges could be tailored according to individual differences in decision making and personality. In Study 1, we tested a hypothetical disclosure choice embedded with a framing nudge. Study 2 cast a wider net by testing more types of nudges and a wider variety of psychometric variables. This study yielded three potentially significant effects which we attempted to replicate in Study 3. Our replication did not confirm these effects, finding them to be fragile at best. Together, the results of these studies suggest that tailored privacy nudges are likely not feasible for many small-scale applications of nudges - such as those seen in the privacy literature - and reaffirms the importance of replicating potentially significant results. While on statistical grounds a null effect cannot be used to rule out their existence, the potential sample sizes required to identify robust means of tailoring privacy nudges likely make them impractical to all but the largest organizations.

ACKNOWLEDGMENTS

This research has been supported in part by DARPA and the Air Force Research Laboratory, as part of the DARPA Brandeis program, under grant number FA8750-15-2-0277. The US Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, the Air Force Research Laboratory, or the US Government. Our research was also supported by the Department of Engineering and Public Policy at Carnegie Mellon University. Acquisti gratefully acknowledges funding from the National Science Foundation under award 1514192.

REFERENCES

[1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2016. Nudges for Privacy and Security : Understanding and Assisting Users ' Choices Online. (2016), 1–40.
[2] A Acquisti, LK John, and G Loewenstein. 2012. The impact of relative standards on the propensity to disclose. *Journal of Marketing Research* 49, 2 (2012), 160–174. https://doi.org/10.1509/jmr.09.0215
[3] Idris Adjerid, Eyal Peer, and Alessandro Acquisti. 2016. Beyond the Privacy Paradox: Objective versus Relative Risk in Privacy Decision Making. (2016). https://doi.org/10.2139/ssrn.2765097
[4] Hunt Allcott and Judd Kessler. 2015. The Welfare Effects of Nudges : A case study of energy use social comparisons. (2015).
[5] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Cranor, and Yuvraj Agarwal. 2014. Your Location has been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. (2014).
[6] Rebecca Balebako, Pedro G Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2011. Nudging Users Towards Privacy on Mobile Devices. *CEUR Workshop Proceedings* 722 (2011), 23–26.
[7] Murray R Barrick and Michael K Mount. 1991. The Big Five Personality Dimensions and Job Performance : A Meta Analysis. *Personnel Psychology* 44, 1 (1991), 1–26. https://doi.org/10.1111/j.1744-6570.1991.tb00688.x
[8] John Beshears, James J. Choi, David Laibson, Brigitte C. Madrian, and Sean Wang. 2015. Who Is Easier to Nudge? *NBER Working Paper* 401 (2015). https://doi.org/10.1007/s13398-014-0173-7.2
[9] Wändi Bruine de Bruin, Andrew M Parker, and Baruch Fischhoff. 2007. Adult Decision-Making Competence Inventory. (2007), 938–956.
[10] Wändi Bruine de Bruin, Andrew M Parker, and Baruch Fischhoff. 2007. Individual Differences in Adult Decision-Making Competence. *Journal of Personality and Social Psychology* 92, 5 (2007), 938–956. https://doi.org/10.1037/0022-3514.92.5.938
[11] J T Cacioppo, R E Petty, and C F Kao. 1984. The efficient assessment of need for cognition. , 306–307 pages. https://doi.org/10.1207/s15327752jpa4803_13
[12] Lynne M. Coventry, Debora Jeske, John M. Blythe, James Turland, and Pam Briggs. 2016. Personality and social framing in privacy decision-making: A

study on cookie acceptance. *Frontiers in Psychology* 7, SEP (2016), 1–12. https://doi.org/10.3389/fpsyg.2016.01341
[13] James N Druckman, I Thank, Larry Bartels, Adam Berinsky, Gregory Bovitz, Gary Cox, Daniel Druckman, Nicole Druckman, Elisabeth Gerber, Zoltan Hajnal, David King, James Kuklinski, Arthur Lupia, Mathew Mccubbins, Thomas Nelson, Robert Oldendick, Samuel Popkin, and Michael Schudson. 2001. *On the Limits of Framing Effects: Who Can Frame?* Technical Report 4. 18 pages. http://www.journals.uchicago.edu/t-and-c
[14] Caitlin Drummond and Baruch Fischhoff. 2017. Development and Validation of the Scientific Reasoning Scale. *Journal of Behavioral Decision Making* 30, 1 (2017), 26–38. https://doi.org/10.1002/bdm.1906
[15] Serge Egelman and Eyal Peer. 2015. The Myth of the Average User. *NSPW 2015 Proceedings* (2015), 16–28.
[16] Alan Fernihough. 2014. mfx: Marginal Effects, Odds Ratios and Incidence Rate Ratios for GLMs. https://cran.r-project.org/package=mfx
[17] John Fox and Sanford Weisberg. 2011. *An {R} Companion to Applied Regression* (second ed.). Sage, Thousand Oaks {CA}. http://socserv.socsci.mcmaster.ca/jfox/Books/Companion
[18] Daniel G. Goldstein, Elke U. Weber, Craig Fox, Gerald Häubl, Ellen Peters, Richard P. Larrick, Brian Wansink, John W. Payne, David Schkade, Eric J. Johnson, Benedict G. C. Dellaert, and Suzanne B. Shu. 2012. Beyond nudges: Tools of a choice architecture. *Marketing Letters* 23, 2 (2012), 487–504. https://doi.org/10.1007/s11002-012-9186-1
[19] Andrew F. Hayes and Jörg Matthes. 2009. Computational procedures for probing interactions in OLS and logistic regression: SPSS and SAS implementations. *Behavior Research Methods* 41, 3 (2009), 924–936. https://doi.org/10.3758/BRM.41.3.924
[20] Marek Hlavac. 2015. stargazer: Well-Formatted Regression and Summary Statistics Tables. http://cran.r-project.org/package=stargazer
[21] Daniel Kahneman and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Source: Econometrica* 47, 2 (1979), 263–292. http://www.jstor.org/stable/1914185http://www.jstor.org/stable/1914185?seq=1{&}cid=pdf-reference{#}references{_}tab{_}contentshttp://about.jstor.org/terms
[22] Carmen Keller and Michael Siegrist. 2009. Effect of risk communication formats on risk perception depending on numeracy. *Medical Decision Making* 29, 4 (2009), 483–490. https://doi.org/10.1177/0272989X09333122
[23] Bart P. Knijnenburg and Alfred Kobsa. 2013. Helping users with information disclosure decisions. Association for Computing Machinery (ACM), 407. https://doi.org/10.1145/2449396.2449448
[24] Isaac M Lipkus, Greg Samsa, and Barbara K Rimer. 2001. General Performance on a Numeracy Scale among Highly Educated Samples. 21 (2001), 37–44. https://doi.org/10.1177/0272989X0102100105
[25] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information systems research* 15, 4 (2004), 336–355.
[26] Nathan Malkin, Arunesh Mathur, Marian Harbach, and Serge Egelman. 2017. Personalized Security Messaging: Nudges for Compliance with Browser Warnings. *EuroUSEC 17* April (2017). https://doi.org/10.14722/eurousec.2017.23008
[27] Barrington Moore. 1984. *Privacy: Studies in Social and Cultural History.*
[28] P A Norberg and D R Horne. 2007. The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors - NORBERG - 2007 - Journal of Consumer Affairs - Wiley Online Library. *Journal of Consumer ...* 41, 1 (2007), 100–126. https://doi.org/10.1111/j.1745-6606.2006.00070.x
[29] Gabriele Paolacci and Jesse Chandler. 2014. Inside the Turk: Understanding Mechanical Turk as a Participant Pool. *Current Directions in Psychological Science* 23, 3 (2014), 184–188. https://doi.org/10.1177/0963721414531598
[30] Yong Jin Park. 2013. Digital Literacy and Privacy Behavior Online. *Communication Research* 40, 2 (apr 2013), 215–236. https://doi.org/10.1177/0093650211418338
[31] Eyal Peer, Serge Egelman, Marian Harbach, Nathan Malkin, Arunesh Mathur, and Alisa Frik. 2019. Nudge Me Right: Personalizing Online Nudges to People's Decision-Making Styles. *SSRN Electronic Journal* (2019), 1–23. https://doi.org/10.2139/ssrn.3324907
[32] Eyal Peer, Sonam Samat, Laura Brandimarte, Alessandro Acquisti, and Corresponding Author Graduate. [n.d.]. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. ([n. d.]).
[33] Ellen Peters. 2008. Numeracy and the perception and communication of risk. *Annals of the New York Academy of Sciences* 1128 (2008), 1–7. https://doi.org/10.1196/annals.1399.001
[34] Sonam Samat, Alessandro Acquisti, Santa Clara, and Alessandro Acquisti. 2017. Format vs. Content: The Impact of Risk and Presentation on Disclosure Decisions. *Symposium on Usable Privacy and Security (SOUPS) 2017* Soups (2017), 9. https://www.usenix.org/conference/soups2017/technical-sessions/presentation/samat-disclosure
[35] Susanne Scott and Bruce Reginald. 1995. Decision-Making Style: The Development and Assessment of a New Measure. *Educational and psychological measurement* (1995), 818–831.
[36] Susanne G. Scott and Reginald A. Bruce. 1995. Decision-Making Style: The Development and Assessment of a New Measure. *Educational and Psychological Measurement* 55, 5 (1995), 818–831. https://doi.org/10.1177/0013164495055005017 arXiv:0803973233
[37] J.H Smith, S.J Milberg, and S.J Burke. 1996. Information privacy: measuring individuals concerns about organizational practices. *MIS quaterly* 20, 2 (1996), 167–196.
[38] Neil Stewart, Adam J L Harris, Daniel M Bartels, Ben R Newell, Gabriele Paolacci, and Jesse Chandler. 2015. The average laboratory samples a population of 7,300 Amazon Mechanical Turk workers. *Judgment and Decision Making* 10, 5 (2015), 479–491. https://doi.org/10.1017/CBO9781107415324.004 arXiv:arXiv:1011.1669v3
[39] Cass R. Sunstein. 2014. Nudging: A Very Short Guide. *Journal of Consumer Policy* 37, 4 (2014), 583–588. https://doi.org/10.1007/s10603-014-9273-1 arXiv:arXiv:1011.1669v3
[40] Richard Thaler and Cass Sunstein. 2009. *Nudge: Improving Decisions About health, Wealth, and Happiness.* Penguin Books.
[41] James Turland, Lynne Coventry, Debora Jeske, Pam Briggs, and Aad van Moorsel. 2015. Nudging Towards security: Developing an Application for Wireless Network Selection for Android Phones. *Proceedings of the 2015 British HCI Conference on - British HCI '15* (2015), 193–201. https://doi.org/10.1145/2783446.2783588
[42] Blase Ur, PG Kelley, Saranga Komanduri, and Joel Lee. 2012. How does your password measure up? The effect of strength meters on password creation. *Security'12 Proceedings of the 21st USENIX conference on Security symposium* (2012), 5–5. https://www.usenix.org/system/files/conference/usenixsecurity12/sec12-final209.pdf
[43] Paul Veyne and Arthur Goldhammer. 1987. *From pagan Rome to Byzantium. Vol. 1 of A history of private life.*
[44] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, and Norman Sadeh. 2014. A field trial of privacy nudges for facebook. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (2014), 2367–2376. https://doi.org/10.1145/2556288.2557413
[45] Taiyun Wei and Viliam Simko. 2017. R package "corrplot": Visualization of a Correlation Matrix. https://github.com/taiyun/corrplot
[46] Alan Westin and Oscar Ruebhausen. 1967. *Privacy and freedom.*
[47] Hadley Wickham and Romain Francois. 2016. dplyr: A Grammar of Data Manipulation. https://cran.r-project.org/package=dplyr

A TEXT OF EXPERIMENTAL MANIPULATIONS

A.1 Study 1: Auto Checkout Disclosure

A.1.1 'Opt-In' Condition. Imagine a new supermarket in your town has implemented sensors which allow it to determine your location within the store based on the position of your smartphone. This technology is used by the supermarket to enable a new service called "Auto-Checkout", which lets you skip checkout lines by using your location to identify which products you intend to purchase and charging your credit card. The supermarket may also use your location data to provide you with ads for products based on your previous purchases. These functionalities are implemented through a smartphone app that you have installed on your phone. The "Auto-Checkout" system has been tested by the supermarket to ensure accuracy, reliability, and security. As you enter the supermarket, the following message is displayed to you on your smartphone screen.

See Figure 2

A.1.2 'Opt-Out' Condition. Imagine a new supermarket in your town has implemented sensors which allow it to determine your location within the store based on the position of your smartphone. This technology is used by the supermarket to enable a new service called "Auto-Checkout", which lets you skip checkout lines by using your location to identify which products you intend to purchase and charging your credit card. The supermarket may also use your location data to provide you with ads for products based on your previous purchases. These functionalities are implemented through a smartphone app that you have installed on your phone. The "Auto-Checkout" system has been tested by the supermarket to ensure

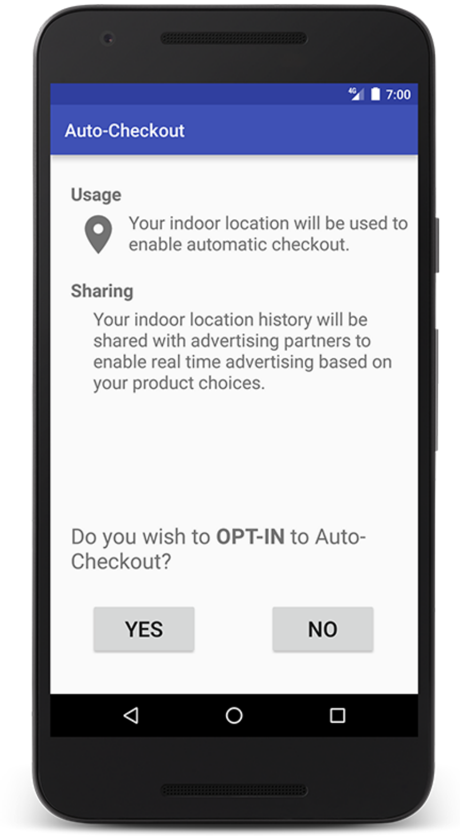

Figure 2: Interactive phone screen for 'Opt-In' Condition in Study 1

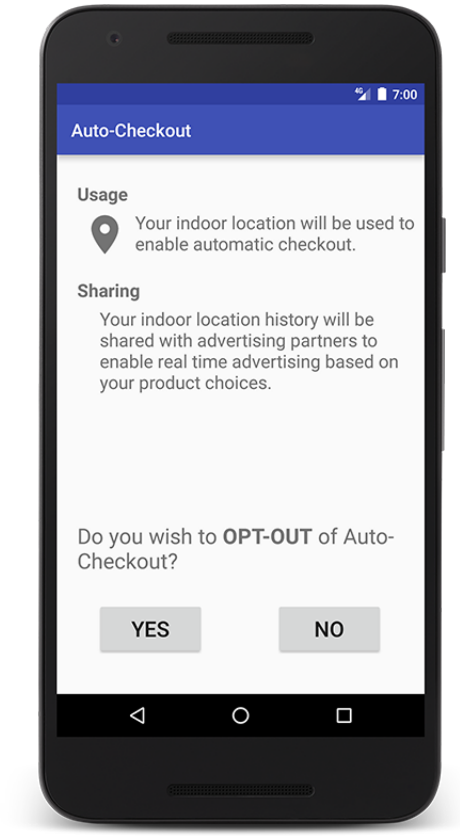

Figure 3: Interactive phone screen for 'Opt-Out' Condition in Study 1

accuracy, reliability, and security. As you enter the supermarket, the following message is displayed to you on your smartphone screen.

See Figure 3

A.2 Study 2 & 3: Ethical Behavior Questions Disclosure

A.2.1 Framing Nudge Allow Condition. Before we ask you the ethical behavior questions, we want to determine your preferences over the sharing of your responses. **You can have your responses to the ethical behavior questions, along with your Mechanical Turk ID, shared with researchers outside of our study team. If you consent to sharing, we will make this data available to researchers outside of our study team.** These researchers will use your responses as part of future published research studies. We show your Mechanical Turk ID as Participant's Mechanical Turk ID.

Please note that your Mechanical Turk ID may not be anonymous. Recent information has shown that your Mechanical Turk ID may be linked to information that can identify you such as your full name and Amazon purchase history. While we do not access or use this information in our study, other researchers may use this information in future published research studies.

Allow your responses to the ethical behavior questions to be shared with **researchers outside of our study team, along with your Mechanical Turk ID?**

A.2.2 Framing Nudge Prohibit Condition. Before we ask you the ethical behavior questions, we want to determine your preferences over the sharing of your responses. **You can have your responses to the ethical behavior questions, along with your Mechanical Turk ID, shared with researchers outside of our study team. If you consent to sharing, we will make this data available to researchers outside of our study team.** These researchers will use your responses as part of future published research studies. We show your Mechanical Turk ID as Participant's Mechanical Turk ID.

Please note that your Mechanical Turk ID may not be anonymous. Recent information has shown that your Mechanical Turk ID may be linked to information that can identify you such as your full name and Amazon purchase history. While we do not access or use this information in our study, other researchers may use this information in future published research studies.

Prohibit your responses to the ethical behavior questions from being shared with **researchers outside of our study team, along with your Mechanical Turk ID?**

A.2.3 Social Norms Nudge 'High Norms' Condition. Before we ask you the ethical behavior questions, we want to determine your preferences over the sharing of your responses. **You can have your responses to the ethical behavior questions, along with your Mechanical Turk ID, shared with researchers outside of our study team. If you consent to sharing, we will make this data available to researchers outside of our study team.** These researchers will use your responses as part of future published research studies. We show your Mechanical Turk ID as Participant's Mechanical Turk ID.

Please note that your Mechanical Turk ID may not be anonymous. Recent information has shown that your Mechanical Turk ID may be linked to information that can identify you such as your full name and Amazon purchase history. While we do not access or use this information in our study, other researchers may use this information in future published research studies. **In our past studies, 73% of participants chose to allow their responses to be shared with researchers outside of our study team.**

Please select a sharing preference below:

- **Allow** my responses and Mechanical Turk ID to be shared with researchers outside of the study team
- **Prohibit** my responses and Mechanical Turk ID from being shared with researchers outside of the study team

A.2.4 Social Norms Nudge 'Low Norms' Condition. Before we ask you the ethical behavior questions, we want to determine your preferences over the sharing of your responses. **You can have your responses to the ethical behavior questions, along with your Mechanical Turk ID, shared with researchers outside of our study team. If you consent to sharing, we will make this data available to researchers outside of our study team.** These researchers will use your responses as part of future published research studies. We show your Mechanical Turk ID as Participant's Mechanical Turk ID.

Please note that your Mechanical Turk ID may not be anonymous. Recent information has shown that your Mechanical Turk ID may be linked to information that can identify you such as your full name and Amazon purchase history. While we do not access or use this information in our study, other researchers may use this information in future published research studies. **In our past studies, 31% of participants chose to allow their responses to be shared with researchers outside of our study team.**

Please select a sharing preference below:

- **Allow** my responses and Mechanical Turk ID to be shared with researchers outside of the study team
- **Prohibit** my responses and Mechanical Turk ID from being shared with researchers outside of the study team

B PSYCHOMETRIC VARIABLE SCALES

B.1 Resistance to Framing

Questions are from the Adult Decision-Making Competence inventory. The complete set of questions used in our studies can be found at: http://www.sjdm.org/dmidi/Adult_-_Decision_Making_Competence.html.

B.1.1 Part 1 Sample Questions. Each of the following problems presents a choice between two options. Each problem is presented with a scale ranging from 1 (representing one option) through 6 (representing the other option). For each item, please select the number on the scale that best reflects your relative preference between the two options.

Q1: Imagine that recent evidence has shown that a pesticide is threatening the lives of 1,200 endangered animals. Two response options have been suggested:

- If Option A is used, 600 animals will be saved for sure.
- If Option B is used, there is a 75% chance that 800 animals will be saved, and a 25% chance that no animals will be saved.

Which option do you recommend to use?

Q2: Because of changes in tax laws, you may get back as much as $1200 in income tax. Your accountant has been exploring alternative ways to take advantage of this situation. He has developed two plans:

- If Plan A is adopted, you will get back $400 of the possible $1200.
- If Plan B is adopted, you have a 33% chance of getting back all $1200, and a 67% chance of getting back no money.

Which plan would you use?

B.1.2 Part 2 Sample Questions. Each of the following problems presents a choice between two options. Each problem is presented with a scale ranging from 1 (representing one option) through 6 (representing the other option). For each item, please select the number on the scale that best reflects your relative preference between the two options.

Q1: Imagine that recent evidence has shown that a pesticide is threatening the lives of 1,200 endangered animals. Two response options have been suggested:

- If Option A is used, 600 animals will be lost for sure.
- If Option B is used, there is a 75% chance that 400 animals will be lost, and a 25% chance that 1,200 animals will be lost.

Which option do you recommend to use?

Q2: Because of changes in tax laws, you may get back as much as $1200 in income tax. Your accountant has been exploring alternative ways to take advantage of this situation. He has developed two plans:

- If Plan A is adopted, you will lose $800 of the possible $1200.
- If Plan B is adopted, you have a 33% chance of losing none of the money, and a 67% chance of losing all $1200.

Which plan would you use?

B.2 Recognizing Social Norms

Questions are from the Adult Decision-Making Competence inventory. The complete set of questions used in our studies can be found at: http://www.sjdm.org/dmidi/Adult_-_Decision_Making_Competence.html.

B.2.1 Part 1 Sample Questions. The following problems ask whether it is sometimes OK to do different things. For each question, please indicate whether *in your opinion* the answer is yes or no.

Do you think it is sometimes OK...

- to steal under certain circumstances?
- to smoke cigarettes?
- to commit a crime which could put you in jail?
- to keep things you find in the street?
- to experiment with marijuana?

B.2.2 Part 2 Sample Questions. The following problems ask out of *100 people your age*, how many would say that it is sometimes OK to do different things. For each question, please select a number between 0 (meaning *no one* thinks that it is sometimes OK) and 100 (meaning *everyone* thinks that it is sometimes OK).

Out of 100 people your age, how many would say it is sometimes OK ...

- to steal under certain circumstances?
- to smoke cigarettes?
- to commit a crime which could put you in jail?
- to keep things you find in the street?
- to experiment with marijuana?

C POST-HOC POWER ANALYSIS RESULTS

C.1 Social Norms Nudge and Recognizing Social Norms Score

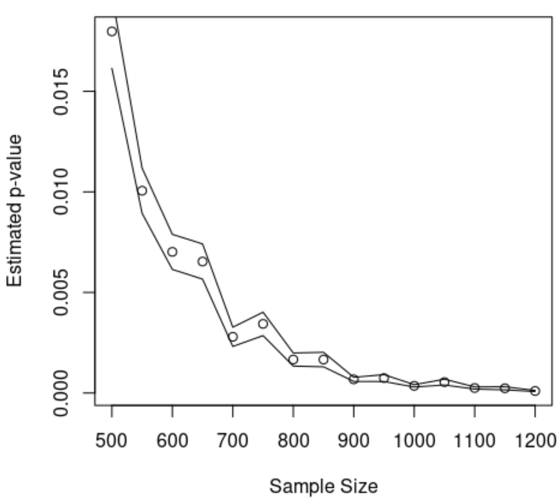

Figure 4: Estimated p-values for different sample sizes for the Social Norms Nudge and Recognizing Social Norms score with 95% confidence interval

C.2 Framing Nudge and Big 5 Extraversion Score

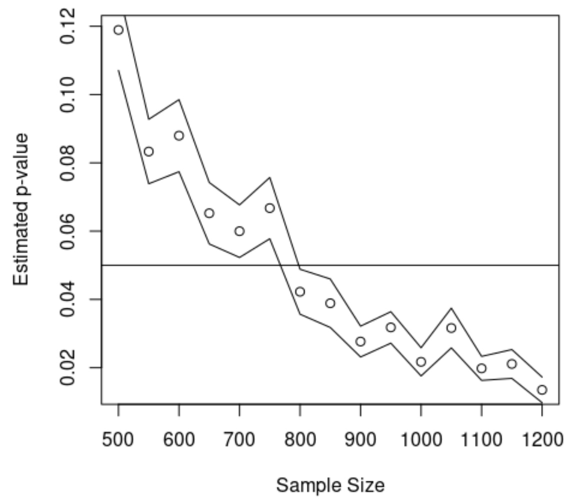

Figure 5: Estimated p-values for different sample sizes for the Framing Nudge and Big 5 Extraversion score with 95% confidence interval

C.3 Framing Nudge and Big 5 Conscientiousness Score

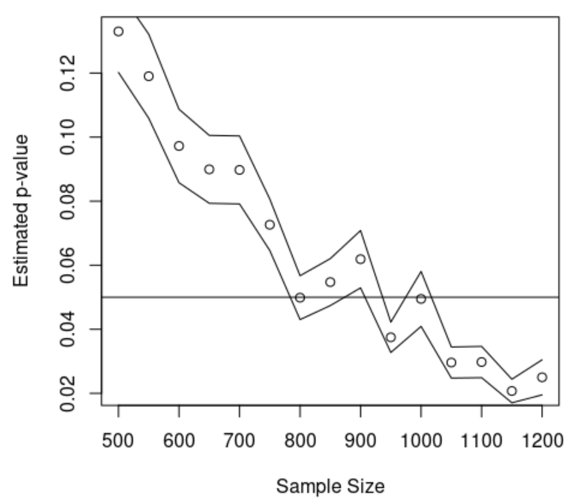

Figure 6: Estimated p-values for different sample sizes for the Framing Nudge and Big 5 Extraversion score with 95% confidence interval

D ADDITIONAL REGRESSION TABLES

D.1 Study 1 Regression Tables

Table 5: Framing Score and Log-ADR Score Logit Regressions for Study 1

	Dependent variable:							
	disclosure							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Allow Frame	3.780	3.959	3.884	3.909	0.639	0.632	0.683	0.745
	(3.396)	(3.436)	(3.447)	(3.496)	(0.544)	(0.544)	(0.554)	(0.560)
Framing Score	0.382	0.358	0.352	0.266				
	(0.592)	(0.594)	(0.598)	(0.607)				
CFIP Score		−0.179	−0.220	−0.238		−0.097	−0.130	−0.143
		(0.181)	(0.188)	(0.192)		(0.184)	(0.189)	(0.190)
Age			0.009	0.007			0.008	0.006
			(0.018)	(0.019)			(0.019)	(0.019)
Female			0.414	0.391			0.506	0.493
			(0.374)	(0.378)			(0.381)	(0.383)
Years of Education				0.141				0.076
				(0.093)				(0.096)
Allow Frame*Framing Score	−0.761	−0.814	−0.773	−0.768				
	(0.871)	(0.882)	(0.885)	(0.898)				
Log-ADR Score					1.064^{**}	1.024^{*}	1.112^{**}	1.001^{*}
					(0.526)	(0.525)	(0.550)	(0.565)
Allow Frame*Log-ADR Score					−0.232	−0.221	−0.311	−0.250
					(0.763)	(0.760)	(0.783)	(0.787)
Constant	−1.703	−0.540	−0.797	−2.352	0.408	0.967	0.682	−0.381
	(2.306)	(2.589)	(2.612)	(2.835)	(0.380)	(1.123)	(1.229)	(1.818)
Observations	143	143	143	143	143	143	143	143
Log Likelihood	−95.133	−94.630	−93.792	−92.626	−91.910	−91.768	−90.703	−90.389
Akaike Inf. Crit.	198.265	199.260	201.585	201.252	191.820	193.536	195.406	196.778
Note:								*p<0.1; **p<0.05; ***p<0.01

D.2 Study 2 Regression Tables

Table 6: Social Norms Nudge and Recognizing Social Norms Regressions for Study 2

	Dependent variable:			
	disc			
	(1)	(2)	(3)	(4)
Recognizing Social Norms	-1.744^{***}	-1.318^{**}	-1.194^{**}	-1.232^{**}
	(0.538)	(0.564)	(0.573)	(0.574)
High Norms Condition	−0.429	−0.281	−0.293	−0.322
	(0.395)	(0.412)	(0.419)	(0.420)
IUIPC		-0.591^{***}	-0.536^{***}	-0.527^{***}
		(0.125)	(0.127)	(0.128)
Age			-0.020^{**}	-0.021^{**}
			(0.010)	(0.010)
Female			−0.086	−0.083
			(0.207)	(0.208)
African American			−0.023	−0.012
			(0.360)	(0.363)
Hispanic			0.205	0.222
			(0.404)	(0.405)
Asian			0.491	0.489
			(0.489)	(0.493)
Other Race			0.197	0.147
			(0.620)	(0.621)
High School				−12.697
				(535.411)
Associate Degree				−12.650
				(535.411)
Bachelor's Degree				−12.768
				(535.411)
Advanced Degree				−13.118
				(535.411)
Recognizing Social Norms*High Norms Condition	2.542^{***}	2.294^{***}	2.400^{***}	2.432^{***}
	(0.743)	(0.772)	(0.783)	(0.784)
Constant	0.572^{**}	3.870^{***}	4.225^{***}	16.968
	(0.283)	(0.764)	(0.808)	(535.411)
Observations	453	453	453	453
Log Likelihood	−298.305	−285.683	−282.242	−281.000
Akaike Inf. Crit.	604.611	581.366	586.485	591.999
Note:		*p<0.1; **p<0.05; ***p<0.01		

Table 7: Framing Nudge and Big 5 Extraversion Regressions for Study 2

	Dependent variable:			
	disc			
	(1)	(2)	(3)	(4)
Big 5 Extraversion	−0.022	−0.045	−0.058	−0.056
	(0.133)	(0.137)	(0.139)	(0.140)
Allow Condition	−0.260	−0.336	−0.400	−0.398
	(0.595)	(0.605)	(0.612)	(0.616)
IUIPC		-0.497^{***}	-0.480^{***}	-0.485^{***}
		(0.118)	(0.121)	(0.122)
Age			−0.009	−0.008
			(0.009)	(0.009)
Female			0.205	0.200
			(0.202)	(0.203)
African American			0.170	0.162
			(0.367)	(0.369)
Hispanic			0.056	0.027
			(0.360)	(0.365)
Asian			0.417	0.451
			(0.464)	(0.472)
Other Race			0.136	0.130
			(0.650)	(0.648)
Associate Degree				0.166
				(0.327)
Bachelor's Degree				0.274
				(0.302)
Advanced Degree				−0.086
				(0.391)
Other Education				13.218
				(535.411)
Big 5 Extraversion*Allow Condition	0.433^{**}	0.456^{**}	0.486^{**}	0.487^{**}
	(0.197)	(0.200)	(0.203)	(0.204)
Constant	−0.336	2.657^{***}	2.796^{***}	2.615^{***}
	(0.404)	(0.824)	(0.846)	(0.864)
Observations	482	482	482	482
Log Likelihood	−314.760	−304.871	−303.238	−301.875
Akaike Inf. Crit.	637.520	619.742	628.477	633.749
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 8: Framing Nudge and Big 5 Conscientiousness Regressions for Study 2

	Dependent variable:			
	disc			
	(1)	(2)	(3)	(4)
Big 5 Conscientiousness	0.137	0.280	0.282	0.258
	(0.177)	(0.188)	(0.192)	(0.193)
Allow Condition	3.041^{***}	2.992^{***}	2.915^{***}	2.849^{***}
	(1.023)	(1.053)	(1.058)	(1.061)
IUIPC		-0.494^{***}	-0.484^{***}	-0.485^{***}
		(0.120)	(0.122)	(0.123)
Age			−0.008	−0.007
			(0.009)	(0.009)
Female			0.159	0.157
			(0.200)	(0.201)
African American			0.185	0.171
			(0.366)	(0.368)
Hispanic			−0.022	−0.052
			(0.361)	(0.366)
Asian			0.343	0.355
			(0.460)	(0.469)
Other Race			0.114	0.094
			(0.649)	(0.647)
Associate Degree				0.116
				(0.327)
Bachelor's Degree				0.283
				(0.302)
Advanced Degree				−0.013
				(0.387)
Other Education				13.081
				(535.411)
Big 5 Conscientiousness*Allow Condition	-0.514^{**}	-0.506^{*}	-0.482^{*}	-0.464^{*}
	(0.253)	(0.260)	(0.262)	(0.262)
Constant	−0.933	1.420	1.530	1.425
	(0.706)	(0.930)	(0.949)	(0.963)
Observations	482	482	482	482
Log Likelihood	−316.376	−307.109	−305.940	−304.746
Akaike Inf. Crit.	640.752	624.218	633.881	639.492
Note:	*p<0.1; **p<0.05; ***p<0.01			

D.3 Study 3 Regression Tables

Table 9: Social Norms Nudge and Recognizing Social Norms Regressions for Study 3

	Dependent variable:			
	disc			
	(1)	(2)	(3)	(4)
Recognizing Social Norms	-1.233^{***}	-0.985^{***}	-1.031^{***}	-1.054^{***}
	(0.338)	(0.364)	(0.374)	(0.377)
High Norms Condition	0.438	0.309	0.374	0.363
	(0.269)	(0.291)	(0.299)	(0.300)
IUIPC		-0.924^{***}	-0.952^{***}	-0.953^{***}
		(0.097)	(0.100)	(0.101)
Age			0.012^{*}	0.011^{*}
			(0.006)	(0.006)
Female			-0.291^{**}	-0.303^{**}
			(0.147)	(0.148)
Non-Binary			0.911	0.859
			(1.252)	(1.250)
African American			−0.408	−0.414
			(0.263)	(0.265)
Hispanic			0.110	0.096
			(0.320)	(0.320)
Asian			-0.837^{***}	-0.825^{**}
			(0.324)	(0.325)
Other Race			0.168	0.144
			(0.488)	(0.490)
High School				0.590
				(1.038)
Associate Degree				0.513
				(1.020)
Bachelor's Degree				0.374
				(1.016)
Advanced Degree				0.518
				(1.029)
Recognizing Social Norms*High Norms Condition	0.529	0.916^{*}	0.800	0.818
	(0.481)	(0.520)	(0.532)	(0.534)
Constant	0.501^{***}	5.829^{***}	5.781^{***}	5.367^{***}
	(0.189)	(0.611)	(0.642)	(1.160)
Observations	936	936	936	936
Log Likelihood	−618.414	−560.791	−551.897	−551.202
Akaike Inf. Crit.	1,244.828	1,131.581	1,127.794	1,134.404
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 10: Framing Nudge and Big 5 Extraversion Regressions for Study 3

	Dependent variable:			
	disc			
	(1)	(2)	(3)	(4)
Big 5 Extraversion	−0.096	−0.125	−0.099	−0.062
	(0.110)	(0.115)	(0.116)	(0.117)
Allow Condition	0.195	0.256	0.246	0.283
	(0.452)	(0.468)	(0.471)	(0.475)
IUIPC		-0.508^{***}	-0.483^{***}	-0.500^{***}
		(0.069)	(0.070)	(0.071)
Age			-0.010^{*}	−0.010
			(0.006)	(0.006)
Female			−0.007	−0.016
			(0.139)	(0.140)
Non-Binary			0.984	1.206
			(1.230)	(1.223)
African American			−0.295	−0.350
			(0.235)	(0.237)
Hispanic			−0.094	−0.107
			(0.280)	(0.285)
Asian			−0.169	−0.088
			(0.266)	(0.271)
Other Race			−0.368	−0.476
			(0.667)	(0.669)
Associate Degree				1.509
				(0.985)
Bachelor's Degree				1.826^{*}
				(0.974)
Advanced Degree				1.354
				(0.969)
Other Education				1.063
				(0.979)
Big 5 Extraversion*Allow Condition	0.302^{**}	0.281^{*}	0.289^{*}	0.287^{*}
	(0.151)	(0.157)	(0.158)	(0.159)
Constant	−0.421	2.580^{***}	2.796^{***}	1.331
	(0.329)	(0.532)	(0.554)	(1.093)
Observations	992	992	992	992
Log Likelihood	−648.681	−618.239	−615.658	−607.651
Akaike Inf. Crit.	1,305.363	1,246.478	1,255.315	1,247.302
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 11: Framing Nudge and Big 5 Conscientiousness Regressions for Study 3

	Dependent variable:			
	disc			
	(1)	(2)	(3)	(4)
Big 5 Conscientiousness	-0.380^{***}	-0.198	-0.176	-0.141
	(0.121)	(0.127)	(0.128)	(0.130)
Allow Condition	-0.014	0.238	0.230	0.364
	(0.659)	(0.685)	(0.686)	(0.692)
IUIPC		-0.491^{***}	-0.473^{***}	-0.496^{***}
		(0.070)	(0.071)	(0.072)
Age			-0.009	-0.009
			(0.006)	(0.006)
Female			-0.005	-0.015
			(0.139)	(0.140)
Non-Binary			1.078	1.273
			(1.214)	(1.202)
African American			-0.235	-0.279
			(0.231)	(0.233)
Hispanic			-0.080	-0.091
			(0.280)	(0.284)
Asian			-0.162	-0.084
			(0.266)	(0.271)
Other Race			-0.370	-0.472
			(0.669)	(0.671)
Associate Degree				1.396
				(0.982)
Bachelor's Degree				1.707^{*}
				(0.971)
Advanced Degree				1.267
				(0.966)
Other Education				0.967
				(0.976)
Big 5 Conscientiousness*Allow Condition	0.284^{*}	0.216	0.221	0.194
	(0.169)	(0.176)	(0.176)	(0.177)
Constant	0.747	2.877^{***}	3.065^{***}	1.702
	(0.466)	(0.575)	(0.591)	(1.118)
Observations	992	992	992	992
Log Likelihood	-645.753	-618.697	-616.582	-609.227
Akaike Inf. Crit.	1,299.507	1,247.394	1,257.163	1,250.455
Note:		*p<0.1; **p<0.05; ***p<0.01		