



To Self-Persuade or be Persuaded: Examining Interventions for Users' Privacy Setting Selection

Isadora Krsek*

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Kimi V. Wenzel*

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Sauvik Das

Georgia Institute of Technology
Atlanta, Georgia, USA

Jason I. Hong

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Laura A. Dabbish

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

ABSTRACT

User adoption of security and privacy (S&P) best practices remains low, despite sustained efforts by researchers and practitioners. Social influence is a proven method for guiding user S&P behavior, though most work has focused on studying peer influence, which is only possible with a known social graph. In a study of 104 Facebook users, we instead demonstrate that crowdsourced S&P suggestions are significantly influential. We also tested how reflective writing affected participants' S&P decisions, with and without suggestions. With reflective writing, participants were less likely to accept suggestions – both social and Facebook default suggestions. Of particular note, when reflective writing participants were shown the Facebook default suggestion, they not only rejected it but also (unknowingly) configured their settings in accordance with expert recommendations. Our work suggests that both non-personal social influence and reflective writing can positively influence users' S&P decisions, but have negative interactions.

CCS CONCEPTS

- Security and privacy → social aspects of security and privacy.

KEYWORDS

Social Cybersecurity; Social Influence; Social Proof; Authority; Reflective Writing; Surveys; Privacy Behaviors; Decision Making; Privacy; Individual Differences; Quantitative Methods; Qualitative Methods

ACM Reference Format:

Isadora Krsek, Kimi V. Wenzel, Sauvik Das, Jason I. Hong, and Laura A. Dabbish. 2022. To Self-Persuade or be Persuaded: Examining Interventions for Users' Privacy Setting Selection. In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3491102.3502009>

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution International 4.0 License.

1 INTRODUCTION

Roughly two-thirds of U.S. Internet users are estimated to have experienced data theft of some kind [65] and over half of U.S. Internet users surveyed have expressed wanting to do more to protect their online S&P [72]. Despite these sentiments, as well as attempts to raise the security sensitivity of the general populace, the advice of experts is largely ignored and general knowledge of S&P remains low [66, 90]. Moreover, this discrepancy between what S&P experts recommend and what end-users do is not limited to the U.S. – similar findings have been found in other parts of the world as well (e.g., [20]).

Research in behavioral economics highlights bounded rationality and cognitive biases as key hindrances to the capabilities of users in S&P decision making [1–3, 52]. Even when users are fully informed of S&P best practices however, they struggle to consider long-term consequences [1, 3], ultimately making decisions that leave themselves, and sometimes others, vulnerable and insecure [40, 63, 87, 89]. Information on privacy risks are difficult for people to find, understand, and implement [41, 42, 46, 56, 79] making users' existing penchant for defaults (status quo bias) all the more impactful. Prior research also found that individual differences such as end-user security attitudes, demographic characteristics, or personality can shape user's online S&P decision making behavior [23]. For example, users' social orientation (collective identity, individual identity, and segmented identity) [16, 47] seems to predict their behaviors and preferences for settings relating to interpersonal S&P [47].

In this paper, we investigate two novel interventions to help people configure safer online security and privacy (S&P) settings. The first intervention is non-personal social influence. Past work in usable S&P found that users tend to rely on the opinions of peers when uncertain or overwhelmed while configuring settings [30, 75]. Researchers and practitioners in HCI have also leveraged peer influences in encouraging positive behavior change [14, 19, 31, 50, 75]. However, while a growing number of studies acknowledge the impact of social influence on end-user S&P behavior, much of this prior work focuses on the influence of *personal* connections and/or expert recommendations. Yet, the social metadata necessary to show people personalized peer recommendations is rarely available, and it is expensive to solicit expert recommendations for all S&P settings. How effective is social influence from non-personal, non-expert referent groups (e.g., the public at large) at encouraging safer configuration of S&P settings? Collecting recommendation information from this referent group may be more scaleable and

cost-effective, so it is worth understanding its efficacy at encouraging end-users to configure safer online S&P settings relative to, e.g., expert and default recommendations.

The second intervention is reflective writing. Past work has found that, in the context of consumer behavior, reflecting and writing about future outcomes can positively aid decision making [86, 94]. Furthermore, writing self-persuasive content can aid in long-term risk perception [54] and affect decision-making behaviors more broadly [6, 11, 80]. Research on how reflective writing can affect S&P decisions, however, is limited. This lack of prior work may be because writing can be cognitively demanding and time-consuming for users. However, unlike other interventions which are externally motivated, reflective writing may provide users with more intrinsic motivation through deep personal thought [4]. As intrinsic motivators have a stronger influence on users' online self-disclosure behavior than extrinsic motivators [78], exploring how such motivations are formed and how they persist may provide important points of innovation for future interventions.

We evaluated the main and interaction effects of non-personal social influence and reflective writing on how people configured 14 S&P settings on Facebook. Social influence suggestions were concisely annotated beside the corresponding setting options. No explanations justifying the recommendations were given so as to minimize users' attention to the intervention. We chose Facebook because of its various categories of information disclosure, its granular S&P settings, its widespread use, and because of the wealth of social psychology literature focused on Facebook.

We recruited 104 Facebook users from Amazon Mechanical Turk to participate in a 2x4 mixed design, controlled survey experiment, with reflective writing serving as a within-subjects factor (reflective writing vs. no reflective writing) and non-personal social influence serving as a between-subjects factor (expert influence vs. public influence vs. Facebook default control vs. experimental control), (see Fig. 1). We tested two non-personal social influence referent groups: experts and the general public. Recommendations from these two groups were collected through a separate survey prior to the main study. We also included two control conditions: the Facebook default suggestion control, and the no suggestion control. Participants were asked to configure 14 Facebook S&P settings, one at a time. Prior to configuring, participants were asked to engage in reflective writing for one of two groups of seven of these settings (counterbalanced across participants). While making their selection, participants were shown a suggestion or not, depending on the non-personal social influence condition to which they were assigned. We also collected a number of individual psychometric measures for individual users – e.g., their security attitudes and their orientation towards a collectivist identity – to understand how individual personality traits might interact with social suggestions and reflective writing on configuration decisions.

We found that participants who were shown S&P setting suggestions from both experts and the public were significantly more likely to adhere to those suggested configurations than those who saw the Facebook default or no suggestion at all. Moreover, we did not observe a significant difference in the effectiveness of social suggestions from experts and the public, suggesting that non-personal social influence from the public may be a good proxy for the authoritative and personal social influence effects found in prior work.

We also found that reflective writing diminished the effects of the social suggestions, but also increased participants' likelihood of selecting more private configuration settings than the default settings suggested by Facebook. To better understand why, in an adhoc analysis of participants' reflective writing responses, we identified three potential motivations for setting configuration preferences that dampened user adherence to social suggestions. In short, our findings suggest that both social influence and reflective writing can encourage safer configuration of S&P settings, but have negative interaction effects and should be used independently. Finally, we found that two of the individual psychometric measures we collected also correlated with the effectiveness of social influence suggestions: participants with a high collective identity score were more likely to follow social suggestions, while those who used Facebook more passively were less likely to follow those suggestions.

This paper makes the following novel research contributions:

- We introduce two mechanisms for cybersecurity behavior persuasion: non-personal social influence (opinions sourced from the public and experts) and reflective writing.
- We examine how non-personal social influence and reflective writing influence S&P setting configuration.
- We examine how the effects of non-personal social influence and reflective writing vary based on collective identity and social media use patterns.
- We introduce a taxonomy of users' S&P setting configuration motivations.

2 RELATED WORK

We have organized related work into three categories: the current burden of managing S&P (both general and on Facebook), social influence in cybersecurity contexts, and the relationship between reflective writing, self-influence, and behavior change.

2.1 The Security & Privacy Burden

Challenges of Configuration. In describing the key challenges of S&P decision-making, Acquisti and Grossklags detailed several gaps in the assumption that perfect rationality dictates the nuances of an individual's privacy and security behavior, citing bounded rationality as a significant limitation to decision-making behavior [3]. Further research by Das *et al.* identified the concept of *security sensitivity* as a key factor hampering S&P decision making, and is defined by Das as "the awareness of, motivation to use, and knowledge of how to use security tools" and practices [29]. Many studies have explored individual aspects affecting factors of security sensitivity using Facebook as the foundation, due to its popularity and the granularity of its settings. For example Strater *et al.* found interface usability to affect and impair users' notion of S&P features or threats, subsequently hampering their ability to act within their desired boundaries of privacy [81]. When asking individuals to reassess their audience preferences for previous posts on Facebook, Mondal *et al.* observed that users' posts from the past have active privacy settings that no longer match user's consent – highlighting a lack of awareness posed by risks from users shifting privacy preferences [61]. Other research has observed how individual factors like personality, or heuristics and biases can affect their motivation to act securely on Facebook [82, 88]. Facebook has paid close

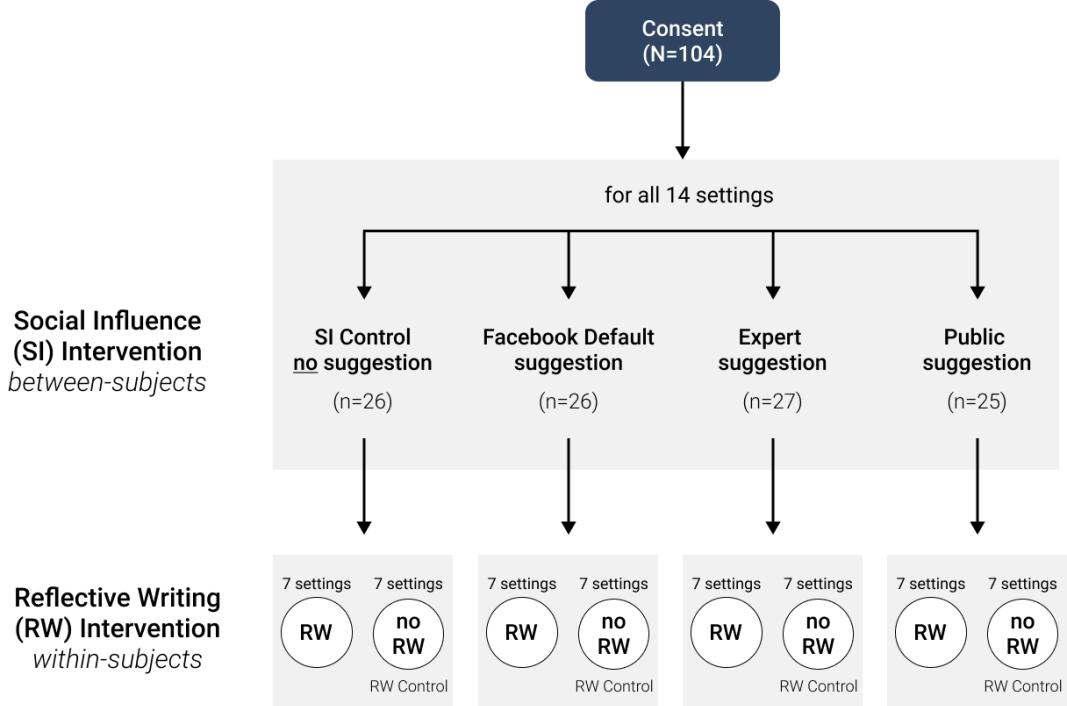


Figure 1: Flowchart depicting the breakdown of experimental conditions for participants. Participants were randomly sorted into one of four between-subjects non-personal social influence conditions. For the within-subjects reflective writing intervention, participants were further split into two groups to determine which settings they reflected on. Setting order was randomized within each reflective writing block, and reflective writing was counterbalanced such that an equal number of participants experienced reflective writing in the first and second half of the study, for each setting.

attention to research in the domain of S&P over the years, and has regularly been developing new features or settings to address user concerns resulting in an extremely complex and granular S&P setting interface [8, 13, 67, 76].

In observing the repercussions of the Facebook interface design, several researchers have identified discrepancies between users' perception of Facebook settings and their actual function or scope, and other gaps in expectations and instances of mishandled S&P. For example, Liu *et al.* found that Facebook privacy settings only matched users' expectations 37% of the time, and when incorrect, almost always expose content to a wider audience than users expected [53]. Later work showed that Facebook users frequently underestimate the audience, guessing that their audience is only 27% of its true size [12]. Such misalignments can be detrimental, as users' online presence play an important role in their reputations; intimate details posted about one's life intended to be seen by friends or family can easily be found inappropriate by other coworkers. The complexities of these settings thus compound existing issues of security sensitivity through worsening missteps for users motivated by S&P concerns or exacerbating a lack of action, and often require knowledge beyond what most end-users have [84].

Our research adds to existing literature by building on our understanding of what types of interventions can be leveraged to mitigate the burden of managing S&P. More specifically our work shows

support for the use of non-personal social influences as an effective intervention for S&P by demonstrating its effectiveness with different non-personal advice-sources (general public and cybersecurity experts). We also present a technique for behavioral persuasion that to our knowledge has yet to be studied in the domain of S&P, reflective writing, and highlight its ability to influence users to either use or shy away from recommended practices.

2.2 Social Influence in S&P

Prior work in psychology has well established social influence as a tool for behavioral persuasion [25, 77]. For example, Burke *et al.* found social learning to be pivotal in affecting how new Facebook users participate on the platform; new account holders who saw friends contributing on the platform would go on to share more content themselves [19]. Researchers have also observed the effects of social proof, a specific type of social influence describing individuals' inclination to look to others for cues on how to behave [26], on S&P behaviors of individuals. The work of Das *et al.* for example has indicated that social interventions are strong motivators in changing awareness, attitude, and the adoption of expert-recommended practices [30–32]; for example, after being exposed to information detailing the security behavior of their friends, Facebook users were more likely to adopt those same behaviors [30]. Relatedly, Rader *et al.* highlighted how colloquially shared

security stories from non-cybersecurity expert peers strongly influenced individuals' behaviors [70]. Redmiles *et al.* expanded on this idea, highlighting common reasons users accept or reject the online security advice shared with them. They found when accepting or rejecting advice, users heavily relied on their evaluation of the trustworthiness as well as technology expertise of their advice-source; some of the most common sources they trust included their workplace, IT professionals, family members, and friends [75].

Work has further explored the role of social influence through the lens of authority. Though not as common in the domain of S&P, the authority principle [26] supports the idea that users' decisions are significantly affected by advice from individuals with known expertise in a subject-matter [36, 59].

That being said, the influence of authority in S&P varies based on the source. Many Americans do not trust authoritative organizations like the federal government or social media companies to protect their data [65], with this sentiment surrounding Facebook in particular [17, 18]. Online behaviors endorsed by social proof (as opposed to organized entities), on the other hand, increase user trust and acceptance [49, 55]. It then follows that participants may be more likely to follow non-personal social influence recommendations rather than options supported by Facebook. However the social meta-data required to scale-up peer recommendations for end-users is rarely available, and for many websites and services is both difficult and expensive to assess since they won't know their users' first degree connections. Crowdsourced recommendations offer an alternative that may both be more cost-effective and easier to scale. As such, it's worth comparing the efficacy of social proof from non-personal sources in encouraging users to configure S&P settings more privately compared to that of experts and defaults.

Our study builds on existing literature on social influence S&P by exploring the effects of non-personal social influence, and whether there is a significant difference in users' adherence to non-personal social influence depending on the advice-source, in this case, cybersecurity experts and the public.

- **H1:** Participants will be more likely to adhere to the non-personal social influence suggestions (experts and the public) than to the Facebook default.

Security Advice Sources. A plethora of research has investigated how users source advice for defensive security configuration. Beaumont *et al.* coined the term "compliance budget" to describe user's limited capacity for compliance with S&P tools and best practices because of their finite amount of energy and time. With limited "compliance budgets" [10], users often rely on advice from a variety of sources such as websites, news media, and peers in an effort to mitigate the energy needed to upkeep their compliance with best practices for privacy and security online [60, 70, 73, 75]. Relatedly, prior work has shown that the likelihood of users' compliance with advice varies based on the source of that advice. Factors such as level of expertise, the strength of a social tie, gender, and more, can influence the acceptance of advice [15, 57, 59, 85]. While recent research in the domain of S&P has explored how the security advice of different social connections is accepted amongst individuals, these studies have explicitly focused on comparing the effect of figures of authority to the advice of friends (with varying degrees of relationship strength) [36, 59]. Work in social psychology has pointed

to crowd advice as a robust mechanism for influence [15, 57] (e.g. in simulated judgement environments estimating unemployment rate for the next quarter), though we are not aware of any work that has been done to compare the effectiveness of crowdsourced advice to that of experts in S&P. We therefore seek to investigate whether there is a difference in effect of expert advice vs. the advice of a non-specified crowd (the public) on users' behavior.

- **H2:** The effect of authority in the expert advice-source condition will differ from the effect of social proof in the public advice-source condition.

Security Attitudes & Uncertainty. End-user security attitudes have also commonly been explored as factors affecting the way that users behave in cybersecurity environments [71]. In the context of S&P, attitudes represent users' evaluation and awareness of, their motivation to use, and their knowledge of how to use S&P measures and best practices — namely, their *security sensitivity* [29, 30]. Faklaris *et al.* further showed that the attitude users have towards security measures correlate with their intentions and actual behaviors in implementing said measures [39]. In the context of social influence, Redmiles *et al.* found that individuals of low security sensitivity are less discerning of S&P advice based on their perceived trust in the source, and they lack feelings of self-efficacy in regards to their S&P compared to individuals of higher security sensitivity. While this study provides insight, the sample was limited to 14 participants, leaving room for further validation and investigation for how security sensitivity affects adherence to social influence overall [75]. What is known however, in both general psychology [25, 33] and S&P in particular [30, 75], is that social proof is especially effective when individuals are faced with uncertainty. Presuming that individuals with a lower awareness of recommended security protocols and tools are more likely to be uncertain about their S&P settings, our hypothesis then follows:

- **H3:** Participants with a low security attitude score (SA-6) [39] will be more likely to adhere to the non-personal social influence suggestions (experts and the public) compared to other participants.

2.3 Collectivism as an Individual Difference

Collectivism. Research into theory on self-disclosure on SNS offers insight into how various individual difference variables such as demographic characteristics, personality, or culture can shape users' disclosure of information online [23]. For example, an individual's orientation to their social world (collectivism vs. individualism) [44] has been established in prior research as correlating with different privacy attitudes [24, 44, 47]. Collectivism increases trust-propensity [91, 93], and individuals' likelihood of being influenced by social proof [51], even on a personal level [27]. Furthermore, collectivists are more likely to foster a sense of similarity with others, in spite of social differences [51], suggesting that collectivists can identify with individuals of varying expertise. However, to our knowledge, no work has yet been published specifically exploring the effects of varying expertise on collectivists' behavior, let alone in the domain of S&P. Altogether this leads to hypothesis H4.

- **H4:** Participants with a high collective identity [16] will be more likely to adhere to the suggestions of non-personal

social influence advice-sources (expert and the public) compared to other participants.

2.4 Reflective Writing

The process of reflecting on and imagining one's own course of action has been a prominent method in behavior change literature. One relevant form of reflection is mental simulation, "the cognitive construction of hypothetical scenarios or the reconstruction of real scenarios" [83]. It is an established methodology in health [64], sports [21], motor skill development [5], and marketing domains [22, 38, 92]. In a meta-analysis of over 100 studies, Cole *et al.* found mental simulation to have a moderate positive effect on behavior change [28].

In our study, we adapt mental simulation to an S&P context and deploy it as a written, as opposed to a purely thought-based, task. Zhao *et al.* showed that asking individuals to simulate and write about future outcomes can prime them to make decisions that are desirable but require more effort [94], and such findings have since been corroborated [86]. Other types of reflective writing exercises, namely self-persuasive writing, have repeatedly been found effective in encouraging positive decision making. For example, Bernritter *et al.* found that asking individuals to participate in a reflective writing exercise prior to paying for their bill (e.g. describing why they thought their waitress deserved a tip) was more effective in increasing consumers' generosity than direct persuasion attempts (e.g. being presented with reasons why their waitress deserves a tip) [11]. Other work has shown self-persuasion to be effective in increasing individuals perceived risk of alcohol consumption [54], increasing vegetable and fruit consumption [80], and increasing attempts at smoking cessation [6]. Self-persuasion provides individuals with intrinsic motivation, making it more powerful than other forms of persuasion in which individuals may be aware that someone is trying to enforce a perspective onto them [4]. Thus, writing about one's own S&P decisions may serve as a persuasive form of mental simulation that overrides the effect of social influence.

- **H5:** Engaging in reflective writing prior to configuring a setting will diminish the effect of social influence.

3 METHODOLOGY

We recruited Facebook users from Amazon Mechanical Turk to participate in a 2 (reflective writing) x 4 (non-personal social influence) mixed design, controlled, survey-based experiment. Participants were asked to configure 14 S&P related settings for their Facebook accounts, either with or without a reflective writing exercise (within-subjects) and given information about either (i) what "experts" recommend, (ii) what the "majority of people" recommend, (iii) the Facebook default setting, or (iv) no information at all (between-subjects). To measure the effect of non-personal social influence and reflective writing on participants' behavior, we recorded their configuration decisions, in addition to written reflections as to their motivations for preferred setting configurations on half of all the settings they saw. To choose the settings we presented to users, we began by cataloguing all of Facebook's S&P related settings, and then narrowed our scope to 14 settings (see Table 1) based on the expert and public suggestion data to which we had

access, and to reduce the time commitment for participants and the overall cost of running the study. Researchers ensured that the chosen settings ranged in diversity. For example, some settings dealt with interpersonal privacy, while others dealt with account security. Similarly, some settings affected people's everyday networking experiencing on Facebook (such settings #9 or #10) while others covered user data preferences (such setting #6). All settings were displayed to participants as screenshots which we took directly from the Facebook interface so as to simulate what they would see in the real-world as closely as possible. Our project was approved by the IRB at our Institute.

3.1 Recruitment (N=104)

We recruited a total of 134 participants through a call for participation in a survey on Facebook settings on MTurk. Participants were screened for eligibility at the beginning of the survey and disqualified before beginning the experiment. Requirements for eligibility included having an active Facebook account, being a U.S. citizen, and being aged 18 or older. We removed 30 participants responses from the analysis after quality control assessments. The results that follow are based off of the responses of 104 participants. Data collection occurred in August of 2020 and participants received USD \$7.50 for completing the survey.

3.2 Survey Experiment

Survey Content. Participants completed an online survey regarding how they would configure 14 different Facebook S&P settings (see Table 1 above for the full list). For half of the settings (counterbalanced across the sample), all participants responded to a reflective writing prompt (Fig. 2). For all settings, participants in non-personal social influence treatment conditions were presented with a setting configuration suggestion. We piloted the experiment with 12 participants whose data was excluded from the analysis.

Facebook setting:

Do you want search engines outside of Facebook to link to your profile?

Options:

- Yes
- No

Please tell us which configuration you would choose for this setting, then reflect on how your choice would affect your experience on Facebook. Be specific when writing your reflection (you may provide example scenarios, etc.).

Write a minimum of 50 characters.

Figure 2: Reflective writing prompt for privacy setting #6.

Reflective Writing. Reflective writing was manipulated within subjects with participants asked to reflect on the first or second half of the 14 total settings they viewed (Fig. 2). We divided the settings

#	Facebook Setting	Facebook Default	Expert	Public	Control
1	Use two-factor authentication	Off	On	On	–
2	Choose 3 to 5 friends to contact if you get locked out	Disable	Enable	Enable	–
3	Who can look you up using the email address you provided?	Everyone	Only me	Only me	–
4	Who can look you up using the phone number you provided?	Everyone	Only me	Only me	–
5	Do you want search engines outside of Facebook to link to your profile?	Yes	No	No	–
6	Do you want Facebook to be able to recognize you in photos and videos?	Yes	No	No	–
7	Encrypted notification emails	Disable	Enable	Enable	–
8	Who can see your friends list?	Public	Friends	Friends	–
9	Who can post on your timeline?	Friends	Friends	Friends	–
10	Allow others to share your posts to their stories?	Enable	Disable	Disable	–
11	Who can see posts you're tagged in on your timeline?	Friends of friends	Friends	Friends	–
12	Review posts you're tagged in before the post appears on your timeline?	Disable	Enable	Enable	–
13	Allow others to share your public stories to their own story?	Allow	Don't allow	Don't allow	–
14	Who can follow me	Friends	Friends	Friends	–

Table 1: Suggested setting configuration based on experimental condition. Non-personal social influence recommendations (cybersecurity experts and the public) were identical. There was no suggested configuration in the control condition.

into two groups in order to balance the number of participants reflecting on each setting and to account for any potential order effects of reflection or that setting group. This meant an equal number of participants reflected on or did not reflect on each setting in the first and second half of the session.

In the writing prompt, participants were asked to first consider what setting configuration they would choose, and then write about how that choice would affect their experience on Facebook (Figure 2). Participants were not given any instructions on how much time to spend for the reflective writing, but were required to write a minimum of 50 characters. The writing prompt was given to them directly prior to seeing a suggested configuration and making a final decision on configuring their own setting. We structured the flow of the experiment in this way to evaluate the effect reflective writing had on the non-personal social influence manipulations.

Social Influence. We deployed the social influence manipulation between-subjects for all 14 settings. Participants were randomly sorted into one of four conditions, which determined what source

they saw setting configuration advice from. The different sources are as follows:

- **Cybersecurity Experts** (social influence)
- **Public** (social influence), referred to as the “majority of people” in the study
- **Facebook Default Configuration** (real-world control)
- **Social Influence Control** (experimental control)

Suggestions for both the general public and cybersecurity expert groups were sourced from two prior surveys conducted in the fall of 2019, which gathered the Facebook setting configuration suggestions of cybersecurity experts (N=10) and crowd workers from MTurk (N=200). For each setting, the setting configuration supported by the majority of experts and crowd workers, respectively, were selected as the setting suggestions for those two conditions in the experiment. Interestingly, the majority suggestions from both the expert and the crowd worker (“public”) groups were identical. To identify the “Facebook default” configuration, researchers created a new account on Facebook and recorded the initial configuration for each setting. The Facebook default differed from the

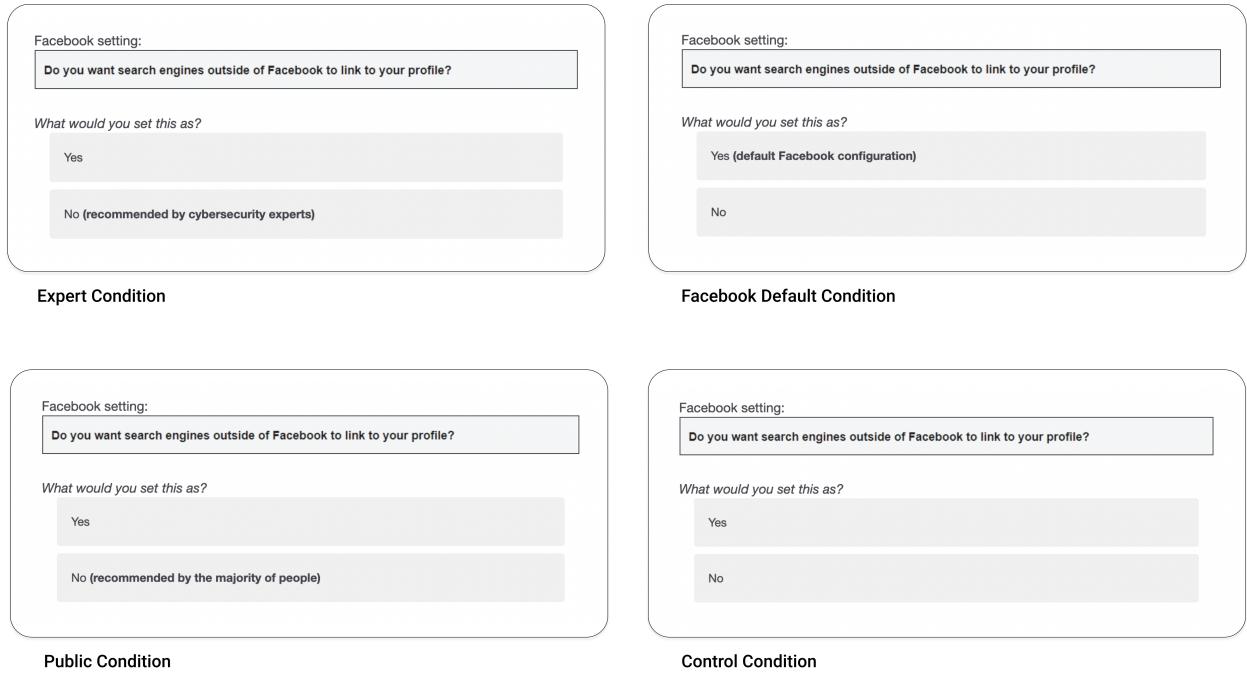


Figure 3: Setting configuration prompt for privacy setting #6, showing how each non-personal social influence condition appears.

suggestions of experts and the crowd workers for all instances. For each of the 14 S&P settings shown, we indicated what configuration option was recommended by placing the advice-source directly adjacent to the choice in the survey (Figure 3). For example participants in the social influence conditions (cybersecurity experts and the public) would see “recommended by [“cybersecurity experts” or “a majority of people”]” while those in the non-social influence condition, Facebook Default, would see “Facebook default” right next to the corresponding option. (See all setting suggestions in Table 1.)

Individual Differences. Several individual differences were measured using previously validated scales. To evaluate participants' security attitudes, we used the SA-6, a six-item self-report scale [39]. We also adopted two existing scales to measure users' collective identity and individual identity [16]. Participants also completed Ellison *et al.*'s [35] Facebook Intensity scale, which includes questions on users' quantity of Facebook friends (response quantities modified to represent larger numbers for present-day users), average time spent on Facebook (validated by Facebook Research [37]), as well as 6 questions measuring users feelings towards Facebook and how integrated Facebook is in their daily lives. We also adapted self-reported metrics on participants' current level of interaction with security and privacy settings on Facebook [69]. An additional scale was created to measure participants' Facebook user habits, categorizing them as either passive users, active users, or neither. This was done in order to observe whether any Facebook use patterns correlated with users' S&P configuration behavior, and specifically

to determine whether Facebook use patterns could be used to predict the chance of users adherence to non-personal social influence suggestions.

To create the Passive Facebook User and Active Facebook User scales, we devised a set of 21 items which participants ranked on a 5 point Likert scale (strongly disagree to strongly agree). The items were piloted twice. A factor analysis from the first pilot (N=77) reduced the number of valid items to 8. These 8 items were supported by the second pilot (N=29). For the final study detailed in this paper, a factor analyses and internal reliability test created a Passive Facebook User scale of 4 items ($0.7 \leq \alpha \leq 0.9$) and an Active Facebook User scale of 3 items ($0.8 \leq \alpha \leq 1.0$). We've included the final items in our Passive and Active Facebook User scales in the Appendix (see table 5 in section A of the appendix).

4 DATA ANALYSIS

4.1 Quantitative Analysis

To evaluate our hypotheses, we encoded participants' selection of the social influence suggestions as a binary variable, *S&P suggestion*. The *S&P suggestion* variable was coded 1 for if participants selected the social influence setting configuration, 0 for if they did not. (Note that both social influence groups, experts and the public, agreed on configurations for all 14 settings.) We encoded a second binary variable noting whether participants selected the Facebook default suggestion: 1 for if they selected the Facebook default setting configuration, 0 for if they did not.

To assess H1, H2, and H5, we employed a random-intercepts logistic regression model, with the aforementioned *S&P best-practice*

variable as the dependent variable, the participant ID as a random intercept term to account for repeated observations, participants' social influence and reflective writing conditions for the setting in question as the independent variables, and the participant's psychometric measures as covariates. To assess H3 and H4, similar models were run except with the psychometric measures then serving as respective independent variables in place of the reflective writing condition. We fit all models using R's lme4 package [9].

4.2 Qualitative Coding

Across the 104 participants who each reflected on 7 of the 14 settings, there were 728 reflections total. 5 reflections were labeled as incoherent and unintelligible by researchers and removed for quality assurance. The remaining 723 were coded according to Table 4. Of the 723 written reflections, the average length was 189 characters. To analyze the results of the reflective writing intervention, initially two researchers read the reflection responses independently and made notes to represent an emerging code book. After initially open-coding the responses, researchers met several times to iteratively discuss, compare, and scrutinize codes. Inter-rater reliability was not calculated, as for open responses where there was a disagreement the coders consulted to come to an agreed upon code after critical and detailed discussions, and multiple rounds of open-coding [58]. This process resulted in a codebook of distinct experiences motivating participant's setting configuration choice. Responses were also scored for depth, length, and detail.

5 RESULTS

Below we present how users' decisions to configure Facebook S&P settings varied with non-personal social influence and reflective writing. This section is structured in order of our hypotheses. First we present our findings regarding non-personal social influence and how effective it was with regards to participant adherence to social advice compared to participants' adherence to the Facebook Default. Next we present our findings related to the effect of individual differences on the influence of non-personal social suggestions. We then describe both the overall affect of the reflective writing intervention on users' setting configuration in addition to the interaction effect of reflective writing and the Facebook default. Finally we discuss how the results from the qualitative analysis of reflective writing responses inform our understanding of what motivated users to look past the non-personal social influence suggestions. The collected demographic data is displayed in the Appendix (see table 6 in section B of the appendix). The sample was skewed male, containing 70.2% male participants, and was relatively young with 71.2% being younger than 40 years old. 42.3% of participants held a bachelor's degree and the majority of participants (64.4%) reported holding a full time job.

5.1 Intervention Effects on S&P Decision Making

Non-personal Social Influence. Participants in the non-personal social influence conditions (expert and public) were 62.8% more likely than those in the control to choose the suggested best practice configuration ($OR=1.93, p<0.01$). This result is in contrast to the Facebook default condition for which there were no significant

effects on the outcome of users' configuration decisions, supporting **H1**. A pairwise Tukey-HSD post-hoc test showed no difference in adherence to suggestions in the expert vs. public conditions ($Est=0.03, p=0.9$), meaning **H2** was not supported.

Factor	Est	SE	OR	Prob	p-value
Intercept	0.75	0.19	2.12	+0.680	9.11e-05***
SI (Expert & Public)	0.66	0.24	1.93	+0.628	0.007**
FB Default	-0.30	0.27	0.74	-0.425	0.259
Reflection	-0.66	0.22	0.51	-0.340	0.003**
SI*Reflection	-0.22	0.28	0.80	-0.444	0.425
FB Default*Reflection	0.99	0.32	2.69	+0.729	.002**

Table 2: Logistic regression results for interaction model to explain selection of S&P suggestion (i.e. whether participants chose the expert and public recommended setting configuration). Est=Estimated Coefficient, SE=Standard Error, OR=Odds Ratio, Prob=Probability, SI=Social Influence, FB=Facebook.

+/ Probability indicates increased/decreased likelihood of the selection of S&P best-practice

* $p<0.05$, ** $p<0.01$, *** $p<0.001$

Factor	Est	SE	OR	Prob	p-value
Intercept	-0.72	0.13	0.48	-0.328	6.12e-08***
SI (Expert & Public)	-0.37	0.15	0.69	-0.408	0.012*
FB Default	0.12	0.17	1.12	+0.527	0.51
Reflection	0.07	0.11	1.08	+0.518	0.53

Table 3: Logistic regression results for main effects model to explain selection of Facebook default suggestion (i.e. whether participants chose the Facebook default setting configuration). Note that an interaction model was also run, for which there were no significant interaction effects. In contrast to the expert and public suggestions (Table 2), the Facebook default suggestion had no significant effect on participants' setting selection. With non-personal social influence manipulations, participants were significantly less likely to select the Facebook default setting configuration. Est=Estimated Coefficient, SE=Standard Error, OR=Odds Ratio, Prob=Probability, SI=Social Influence, FB=Facebook.

+/ Probability indicates increased/decreased likelihood of the selection of Facebook default suggestion

* $p<0.05$, ** $p<0.01$, *** $p<0.001$

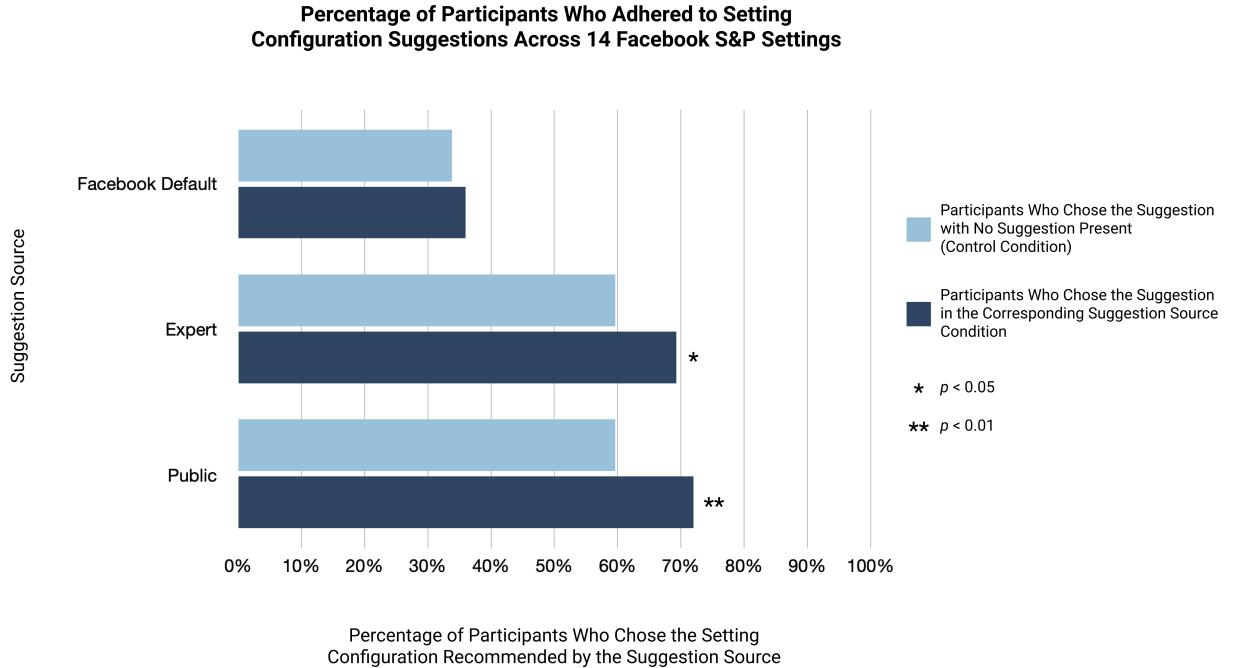


Figure 4: Horizontal bar graph depicting the percentage of participants who adhered to a given suggestion across the different suggestion source conditions. The percentage of participants who adhered to advice in the non-personal social influence conditions (expert and public) was roughly twice (experts=69.3%; public=72.0%) than that of the Facebook Default (36.0%). Compared to the control condition in which participants did not see any suggestions at all, participants in the non-personal social influence conditions chose the suggested configuration about 10% more (59.62% vs. 69.3%, 72.0%). The experts condition manipulation had a significance value of $p < 0.05$, and the public condition manipulation had a significance value of $p < 0.01$. In the Facebook default condition, the difference was about 2% (33.8% vs. 36.0%), with no statistically significant difference. See Table 1 for the list of all settings and suggestions.

Individual Differences that Affected User Adherence to Suggestions. Participants' security attitudes, as measured by SA-6 [39] had no significant relationship with participants' outcomes, failing to reject the null hypothesis for **H3**. However, participants with a collective identity score above the median were 60.6% more likely to follow non-personal social influence suggestions ($OR=1.54, p < 0.05$), confirming **H4**. Furthermore, we found that participants with a Passive Facebook Use score above the median were 34.2% less likely to follow the non-personal social influence suggestions ($OR=0.52, p < 0.05$).

Reflective Writing. We examined the influence of reflective writing by looking at the likelihood of selecting the social influence conditions. We found that with reflective writing, participants were 34.0% less likely to adhere to social suggestions ($OR=0.51, p < 0.01$), confirming **H5**. We also found reflection increased safety of setting choices – in the Facebook default condition participants with reflective writing were 72.9% more likely to choose the S&P best configuration than the suggested default ($OR=2.69, p < 0.01$). No differences in the effectiveness of reflective writing were found based on its order of appearance (e.g. when it occurred first vs. last).

Reflective writing was manipulated within subjects with participants asked to reflect on the first or second half of the 14 total settings they viewed (Fig. 2). We divided the settings into two

groups in order to balance the number of participants reflecting on each setting and to account for any potential order effects of reflection or that setting group. This meant an equal number of participants reflected on or did not reflect on each setting in the first and second half of the session.

5.2 How and Why Reflective Writing Impacts Decision-Making

After open-coding participants' reflective writing responses (see Table 4), we analyzed those codes against setting configuration outcomes. Participants' whose reflections centered around their cybersecurity experiences (see row 1 of table 4 for example) were 80.3% more likely to make a setting selection in line with S&P best-practices and adhere to the suggestions of the experts and the public ($OR=4.08, p < 1e-4$). We also found three significant relationships among participants who did *not* follow the suggestions of social influence advice-sources. Participants who expressed social experience related motivations for their configuration preference, such as prioritizing their experience networking with others online or concern for their reputation among their friends online (see row 2), were 34.9% less likely to follow the suggestion of the non-personal social influence advice-sources ($OR=0.53, p < 1e-4$). For participants who we coded as having user experience related motivations (see

row 4), we found they were 29.3% less likely to choose the social influence suggestion ($OR=0.42, p<0.05$). Finally for participants who expressed indifference to their setting configuration (see row 3), we found that they were 8.3% less likely to adhere to the social influence suggestion ($OR=0.09, p<1e-4$).

6 DISCUSSION

Our findings indicate that S&P suggestions from public groups may be an effective, scalable and broadly applicable method for improving end-user S&P decisions. Indeed, compared to our real world control, the Facebook default, participants were found to be much more likely to adhere to non-personal social influence suggestions. Our results specifically support the use of two sources of non-personal social influences, experts (social influence by authority) and the public (social influence by social proof), with both groups having a similar effect size. When observing how the individual differences of participants influenced their adherence to non-personal social influence suggestions, users with a high collective identity were more likely to adhere to social suggestions. Finally, we found evidence that reflective writing, possibly through a form of behavioral self-persuasion, was effective at encouraging users to configure safer S&P settings aligned with recommended practices. We also found that reflective writing hampers the positive persuasive effect of social influence. However, when paired with the Facebook default instead of social influence, reflection led participants to make more private and secure setting selections.

6.1 Practical Implications for Deploying Behavioral Persuasion

Deployment of Non-personal Social Influence. Prior literature shows that social influence serves as an effective means of guiding the S&P decisions of end-users, but has focused on the influence of peers and family. Our study looks beyond personal connections as social referents, offering evidence that non-personal social influence is an effective technique for behavioral persuasion in S&P. While the metadata needed to show end-users personalized peer recommendations is often unavailable or difficult to gather for many applications (e.g., social graphs displaying the number of friends with similar behaviors), non-personal advice is easy to collect for both websites and S&P experts to crowd source and display themselves. Future work could explore the deployment of non-personal social influence into existing or fabricated privacy setting interfaces. Non-personal social influence may also be explored in the context of a browser plugin, as researchers could crowd source suggestion data from the general public or experts for various websites, and centralize the settings in one place making these settings easy to view and configure. Drawing on prior research on the use of social navigation systems in usable S&P interfaces [34], researchers could investigate the use of non-personal social influence suggestions as default settings in the context of a website. Relatedly, while prior studies have only observed the efficacy of social influence cues presented prior to configuration, our results indicate that social influence is also effective at the time of setting configuration, suggesting that social influence may be fluidly integrated into setting selection designs.

Possible Explanations for the Effects of Reflective Writing. We also test a technique for behavioral self-persuasion that to our knowledge has yet to be studied in the context of S&P, reflective writing, and demonstrate its efficacy in influencing users either towards or away from recommended practices. Our findings show that reflective writing hampers the positive persuasive effect of social influence. Based on previous work on construal level theory (CLT) [7, 43, 87], we may infer that reflective writing narrowed the psychological distance between participants and the outcomes of their S&P decisions. Since psychological distance is egocentric [87] and social influence is inherently centered around others, reflective writing would naturally be more effective than social influence at narrowing this distance and strengthening participants' commitment to their original setting selection. This is one possible explanation for why reflective writing alone worked *against*, as opposed to *with*, the social influence recommendations. This idea is further supported by the power of self-persuasion [4].

In contrast to its interaction in the social influence conditions, reflective writing significantly strengthened participants' likelihood to choose the S&P best-practice setting configuration in the Facebook default condition. This result may possibly be explained by the backfire effect [62]. According to the backfire effect, individuals are more steadfast in their opinion (i.e., the setting configuration participants chose in their reflections) when their opinion is challenged (i.e., when the Facebook default suggestion differed from the participants' chosen configuration). Being exposed to the Facebook default suggestions may have pushed individuals further away from the default suggestion and toward S&P best-practices for this reason.

Deployment of Reflection. Given the recent push in the domain of S&P towards behavioral persuasion interventions, our work not only presents the efficacy of reflective writing in persuading users towards engaging in certain behaviors, but also shows how reflective writing can coax users away from behaviors, and potentially reduce one's likelihood of being manipulated. Future work may explore the efficacy of reflective writing in subverting the effect of dark patterns in S&P, encouraging users away from making uninformed, unintended choices about their personal data or privacy to their own detriment to the benefit of organizations that collect that data [2]. In practice, it may be difficult to smoothly integrate reflective writing into real-world S&P decision making contexts, however one environment where it might be useful is in the classroom. Young adults need to develop an awareness of S&P, however, existing academic programs struggle to educate young adults on S&P awareness. Because of its potential to protect individuals from being exploited and push users towards more secure decisions, reflective writing could serve as an effective exercise for existing S&P education programs aiming to train young adults. Specifically through ingraining in them the importance of slowing down and thinking through their decisions (be it in the context of an configuring settings on an SNS website, or elsewhere). Because our sample didn't include anyone under the age of 18, it would be beneficial for researchers to test whether reflective writing holds the same efficacy with children and teens – as it could potentially be extended to K-12 programs specifically as an exercise for S&P awareness training. It may also be worth exploring more succinct

Theme	%	Definition	Example Quotation
Cybersecurity Experience	57.1	Concern for security of content and information on Facebook.	<i>"Enable [Review posts you're tagged in before they appear on your timeline] in case the post contains information I don't want shared such as my location or something about my children." P7, S12</i>
Social Experience	21.6	Concern for interpersonal access of content and information, or how their social life would be affected.	<i>"Enable [Review posts you're tagged in before they appear on your timeline]. I wouldn't want to be seen in something vulgar or embarrassing without my knowledge." P12, S12</i>
Unaffected Experience	11.4	Frames decision around the fact that their overall experience would be unaffected regardless of what they choose.	<i>"I would disable [Choose 3-5 friends to contact if you get locked out] because I don't think I would care if I had to make a new account." P57, S2</i>
User Experience	11.7	Concern for how/whether their user experience on Facebook may be hindered by a decision.	<i>"I'd disable [Choose 3-5 friends to contact if you get locked out]. It sounds too frustrating and time consuming, even if it's a decent security feature." P50, S2</i>
Vacillate Experience	3.7	Participant was unable to come to a conclusion as they are unsure how the setting would affect their experience, or they otherwise cannot choose a configuration.	<i>"I'm not sure what [Do you want search engines outside of Facebook to link to your profile?] would mean. I assumed everyone on Facebook was accessible via most search engines." P55, S5</i>
Inconclusive	3.1	Participant did not provide enough information to determine the main experiential motivator behind their configuration decision.	<i>"Yes I would like Facebook to recognize you in photos and videos." P36, S6</i>

Table 4: Reflective writing code definitions and examples. The percentage (%) is out of 723 reflections.
 P=Participant ID, S=Setting number (see Table 1 for corresponding settings).

ways of narrowing the psychological distance between users and S&P events, such as asking them to engage in process-based mental simulations in place of a writing exercise.

6.2 Limitations & Future Work

Study Design. In order to properly engage in mental simulation, the reflective writing prompt in our study required users to choose an initial setting configuration. This introduces a limitation to our study as setting choice and written reflection may be unintentionally conflated. Future work may consider completely isolating participants' setting choices from written reflections (i.e. randomly assign participants specific setting choices to reflect on). Researchers may also consider varying the order of choice and reflection and then proceed with a two-sample hypothesis test, as writing a reflection prior to receiving a suggestion may have affected participants' reaction to the suggestions.

Another limitation of our study is that we do not compare non-personal social influence suggestions to personal social influence suggestions. While it was beyond our means to incorporate this

factor into our controlled experiment, future research may explore this comparison more feasibly by involving deception: Instead of garnering real recommendations from participants' close relationships, researchers may assert that participants' friends suggest a specific setting configuration, without acquiring the actual data. This ought to be executed thoughtfully as participants may question the validity of these assertions or may otherwise feel that their privacy has been violated.

This deception and assertion approach may also be used to help isolate the effects of the suggestion source from the content of the suggestions. For example, we found participants to be equally receptive to expert and public recommendations, with those recommendations being identical. The Facebook default condition appeared to be less influential, however this condition provided recommendations which were framed differently (e.g. "Facebook Default" as opposed to "Facebook recommends"). Future work may utilize minor deception to explore how participants react to identical recommendations from all three groups, and how varying these recommendations affect user behavior.

In a second round of coding, we counted 65 (9%) reflective writing responses which clearly demonstrated a misunderstanding of the setting function, however, there was no significant relationship between these misunderstandings and setting configuration decisions. Our study is limited in that we did not distinctly measure participants' understanding of each Facebook setting. Future work may further explore how the efficacy of non-personal social influence and reflective writing vary based on participant understanding. Our study also implemented social influence only at the time of a participants' setting configuration selection. Future work may compare the effects of behavioral persuasion before and during decision-making.

Recruiting & Sample Representation. Participants were not uniformly distributed across all of our conditions. Furthermore, our sample was conducted solely with participants residing in the United States. The gender demographics of our survey experiment experiment significantly skewed towards male. This gender distribution is not representative of the general population, nor is it representative of Facebook user gender statistics. Prior literature has well established gender difference in privacy perceptions and behaviors[45], and as such our results may not be widely generalizable. Overall our surveys were also skewed young and all participants were recruited solely from one online platform, MTurk. As such we introduced several potential biases into our work; our crowd-sourced participants are accustomed to volunteering their participation in research, and on the whole tend to be more tech-savvy than the general population [48, 68, 74]. Despite this, past work comparing an MTurk sample to a census-representative suggests that online S&P studies on MTurk can still act as a reasonable approximation of the S&P behaviors of the general public [74].

Our work may have also introduced additional bias in that our recruitment message called for Facebook users to participate in a survey on their settings. Therefore our sample is also likely biased towards those more inclined or interested in Facebook's approach towards S&P settings. Future work may explore whether these trends appear in larger populations, older populations.

7 CONCLUSION

Prior work has proven the efficacy of social influence from personal connections in guiding end-user S&P behavior, however this personal social influence is difficult to scale because the social metadata necessary to show individuals, e.g., their friends' configuration of specific S&P settings is often unavailable. The results in our study suggest that non-personal social influences, which are considerably easier to source, are significantly influential and as such a promising method for behavioral persuasion in S&P. Furthermore, our findings also suggest that expert and non-personal social referent groups are equally effective – we did not observe a significant difference when comparing the effect on users' setting configuration outcome between the two. Relatedly, our work also highlights that users with high collective identity are more likely to be influenced by non-personal social influence, than users without. This result shows that individual differences can affect users' receptiveness to behavioral persuasion interventions, in particular non-personal social influence.

Additionally, we provide support for a promising behavioral self-persuasion technique that has yet to be explored in S&P – reflective writing. We found that participating in reflective writing can push users towards, or away, from recommended behaviors, with the latter showing promise for it's potential to prevent users from being manipulated. We also describe three potential motivations users have for setting configuration preferences that lead them to look past non-personal social influence suggestions, based on the codes developed from the reflective writing responses. Future research may build on this finding by further exploring the potential for reflective writing to be used in the presence of dark patterns to sway users away from manipulation or in the context of S&P awareness training in young adults. In all, our work serves as an extension of the growing discussion around social cybersecurity and persuasive techniques for manipulating end user security and privacy decision-making behaviors and we believe that these findings can inform future designs that leverage behavioral persuasion in general and for S&P interfaces, in particular.

ACKNOWLEDGMENTS

We would like to thank Haojian Jin, and Cori Faklaris for their assistance in reviewing the work. We would also like to acknowledge the generous support of this project by the National Science Foundation under grant saTC-1704087.

REFERENCES

- [1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 1–41.
- [2] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514.
- [3] Alessandro Acquisti and Jens Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE security & privacy* 3, 1 (2005), 26–33.
- [4] Elliot Aronson. 1999. The power of self-persuasion. *American Psychologist* 54, 11 (1999), 875.
- [5] Sonal Arora, Rajesh Aggarwal, Pramudith Sirimanna, Aidan Moran, Teodor Grantcharov, Roger Kneebone, Nick Sevdalis, and Ara Darzi. 2011. Mental practice enhances surgical technical skills: a randomized controlled study. *Annals of surgery* 253, 2 (2011), 265–270.
- [6] Austin S Baldwin, Alexander J Rothman, Mark W Vander Weg, and Alan J Christensen. 2013. Examining causal components and a mediating process underlying self-generated health arguments for exercise and smoking cessation. *Health Psychology* 32, 12 (2013), 1209.
- [7] Heenkenda Mudiyanselage Ruwan Jayawickrama Bandara, Mario Fernando, and Md Shahriar Akter. 2018. Is the privacy paradox a matter of psychological distance? An exploratory study of the privacy paradox from a construal level theory perspective. (2018).
- [8] Brian Barrett. 2016. You Should Go Check Facebook's New Privacy Settings. *Wired* (2016). <https://www.wired.com/2016/06/go-check-facebooks-new-privacy-settings/>
- [9] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- [10] Adam Beaumont, M Angela Sasse, and Mike Wonham. 2008. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 New Security Paradigms Workshop*. 47–58.
- [11] Stefan F Bernittner, Iris van Ooijen, and Barbara CN Müller. 2017. Self-persuasion as marketing technique: the role of consumers' involvement. *European Journal of Marketing* (2017).
- [12] Michael S. Bernstein, Eytan Bakshy, Moira Burke, and Brian Karrer. 2013. Quantifying the Invisible Audience in Social Networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Paris, France) (CHI '13)*. Association for Computing Machinery, New York, NY, USA, 21–30. <https://doi.org/10.1145/2470654.2470658>
- [13] Nick Bilton. 2010. Price of Facebook Privacy? Start Clicking. *New York Times* CLIX, 55039 (2010), B8 N. <https://www.nytimes.com/2010/05/13/technology/>

personaltech/13basics.html?refer=

[14] Robert Bond, Christopher Fariss, Jason Jones, Adam Kramer, Cameron Marlow, Jaime Settle, and James Fowler. 2012. A 61-Million-Person Experiment in Social Influence and Political Mobilization. *Nature* 489 (09 2012), 295–8. <https://doi.org/10.1038/nature11421>

[15] Bryan L Bonner, Michael R Baumann, and Reeshad S Dalal. 2002. The effects of member expertise on group decision-making and performance. *Organizational Behavior and Human Decision Processes* 88, 2 (2002), 719–736.

[16] Marilyn B Brewer and Ya-Ru Chen. 2007. Where (who) are collectives in collectivism? Toward conceptual clarification of individualism and collectivism. *Psychological review* 114, 1 (2007), 133.

[17] Laura F Bright, Kristen Leah Sussman, and Gary B Wilcox. 2021. Facebook, trust and privacy in an election year: Balancing politics and advertising. *Journal of Digital & Social Media Marketing* 8, 4 (2021), 332–346.

[18] Laura F Bright, Gary B Wilcox, and Hayley Rodriguez. 2019. # DeleteFacebook and the consumer backlash of 2018: How social media fatigue, consumer (mis) trust and privacy concerns shape the new social media reality for consumers. *Journal of Digital & Social Media Marketing* 7, 2 (2019), 177–188.

[19] Moira Burke, Cameron Marlow, and Marilyn Lento. 2009. Feed Me: Motivating Newcomer Contribution in Social Network Sites. *CHI 2009*. <https://doi.org/10.1145/1518701.1518847>

[20] Karoline Busse, Julia Schäfer, and Matthew Smith. 2019. Replication: No one can hack my mind revisiting a study on expert and non-expert security practices and advice. In *Fifteenth Symposium on Usable Privacy and Security ({{SOUPS}} 2019)*. 117–136.

[21] Nichola Callow, Ross Roberts, Lew Hardy, Dan Jiang, and Martin G Edwards. 2013. Performance improvements from imagery: evidence that internal visual imagery is superior to external visual imagery for slalom performance. *Frontiers in human neuroscience* 7 (2013), 697.

[22] Raquel Castaño, Mita Sujan, Manish Kacker, and Harish Sujan. 2008. Managing consumer uncertainty in the adoption of new products: Temporal distance and mental simulation. *Journal of Marketing Research* 45, 3 (2008), 320–336.

[23] Baiyuan Chen and Justin Marcus. 2012. Students' self-presentation on Facebook: An examination of personality and self-construal factors. *Computers in Human Behavior* 28, 6 (2012), 2091–2099.

[24] Chee W Chow, Graeme L Harrison, Jill L McKinnon, and Anne Wu. 1999. Cultural influences on informal information sharing in Chinese and Anglo-American organizations: an exploratory study. *Accounting, Organizations and Society* 24, 7 (1999), 561–582.

[25] Robert B Cialdini and Noah J Goldstein. 2004. Social influence: Compliance and conformity. *Annu. Rev. Psychol.* 55 (2004), 591–621.

[26] Robert B Cialdini and Lloyd James. 2009. *Influence: Science and practice*. Vol. 4. Pearson education Boston, MA.

[27] Robert B Cialdini, Wilhelmina Wosinska, Daniel W Barrett, Jonathan Butner, and Małgorzata Gornik-Durose. 1999. Compliance with a request in two cultures: The differential influence of social proof and commitment/consistency on collectivists and individualists. *Personality and Social Psychology Bulletin* 25, 10 (1999), 1242–1253.

[28] Scott N Cole, Debbie M Smith, Kathryn Ragan, Robert Suurmond, and Christopher J Armitage. 2021. Synthesizing the effects of mental simulation on behavior change: Systematic review and multilevel meta-analysis. *Psychonomic Bulletin & Review* (2021), 1–24.

[29] Sauvik Das. 2017. *Social Cybersecurity: Reshaping Security Through An Empirical Understanding of Human Social Behavior*. Ph. D. Dissertation. Carnegie Mellon University.

[30] Sauvik Das, Tiffany Hyun-Jin Kim, Laura A Dabbish, and Jason I Hong. 2014. The effect of social influence on security sensitivity. In *10th Symposium On Usable Privacy and Security ({{SOUPS}} 2014)*. 143–157.

[31] Sauvik Das, Adam D.I. Kramer, Laura A Dabbish, and Jason I Hong. 2014. Increasing security sensitivity with social proof: A large-scale experimental confirmation. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. 739–749.

[32] Sauvik Das, Adam D.I. Kramer, Laura A. Dabbish, and Jason I. Hong. 2015. The Role of Social Influence in Security Feature Adoption. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work; Social Computing (Vancouver, BC, Canada) (CSCW '15)*. Association for Computing Machinery, New York, NY, USA, 1416–1426. <https://doi.org/10.1145/2675133.2675225>

[33] Morton Deutsch and Harold B Gerard. 1955. A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology* 51, 3 (1955), 629.

[34] Paul DiGioia and Paul Dourish. 2005. Social navigation as a model for usable security. In *Proceedings of the 2005 symposium on Usable privacy and security*. 101–108.

[35] Nicole B Ellison, Charles Steinfield, and Cliff Lampe. 2007. The benefits of Facebook “friends”: Social capital and college students' use of online social network sites. *Journal of computer-mediated communication* 12, 4 (2007), 1143–1168.

[36] Pardis Emami Naeini, Martin Degeling, Lujo Bauer, Richard Chow, Lorrie Faith Cranor, Mohammad Reza Haghagh, and Heather Patterson. 2018. The influence of friends and experts on privacy decision making in IoT scenarios. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–26.

[37] Sindhu Kiranmai Ernala, Moira Burke, Alex Leavitt, and Nicole B Ellison. 2020. How well do people report time spent on Facebook? An evaluation of established survey questions with recommendations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.

[38] Jennifer Edson Escalas. 2004. Imagine yourself in the product: Mental simulation, narrative transportation, and persuasion. *Journal of advertising* 33, 2 (2004), 37–48.

[39] Cori Faklaris, Laura A. Dabbish, and Jason I. Hong. 2019. A Self-Report Measure of End-User Security Attitudes (SA-6). In *Fifteenth Symposium on Usable Privacy and Security ({{SOUPS}} 2019)*. USENIX Association, Santa Clara, CA. <https://www.usenix.org/conference/soups2019/presentation/faklaris>

[40] Shane Frederick. 2005. Cognitive reflection and decision making. *Journal of Economic perspectives* 19, 4 (2005), 25–42.

[41] Simson Garfinkel and Heather Richter Lipford. 2014. Usable security: History, themes, and challenges. *Synthesis Lectures on Information Security, Privacy, and Trust* 5, 2 (2014), 1–124.

[42] Hana Habib and Lorrie Cranor. 2019. Website privacy options aren't much of a choice since they're hard to find and use. <https://theconversation.com/website-privacy-options-arent-much-of-a-choice-since-theyre-hard-to-find-and-use-124631>. Accessed: 2020-09-30.

[43] Cory Hallam and Gianluca Zanella. 2017. Online self-disclosure: The privacy paradox explained as a temporally discounted balance between concerns and rewards. *Computers in Human Behavior* 68 (2017), 217–227.

[44] Geert Hofstede and Gerard Hendrik Hofstede. 1984. *Culture's consequences: International differences in work-related values*. Vol. 5. sage.

[45] Mariea Grubbs Hoy and George Milne. 2010. Gender differences in privacy-related measures for young adult Facebook users. *Journal of interactive advertising* 10, 2 (2010), 28–45.

[46] Philip G Inglesant and M Angela Sasse. 2010. The true cost of unusable password policies: password use in the wild. In *Proceedings of the sigchi conference on human factors in computing systems*. 383–392.

[47] Ruogu Kang. 2015. *Incognito Online: Why and How People Hide Their Information*. *unpublished doctoral dissertation, Carnegie Mellon University* (2015).

[48] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy attitudes of mechanical turk workers and the us public. In *10th Symposium On Usable Privacy and Security ({{SOUPS}} 2014)*. 37–49.

[49] Johannes Klumpe, Oliver Francis Koch, and Alexander Benlian. 2020. How pull vs. push information delivery and social proof affect information disclosure in location based services. *Electronic Markets* 30, 3 (2020), 569–586.

[50] Adam D.I. Kramer. 2012. The Spread of Emotion via Facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Austin, Texas, USA) (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 767–770. <https://doi.org/10.1145/2207676.2207787>

[51] James M Leonhardt, Todd Pezzuti, and Jae-Eun Namkoong. 2020. We're not so different: Collectivism increases perceived homophily, trust, and seeking user-generated product information. *Journal of Business Research* 112 (2020), 160–169.

[52] Eden Litt. 2013. Understanding social network site users' privacy tool use. *Computers in Human Behavior* 29, 4 (2013), 1649–1656.

[53] Yabing Liu, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing Facebook Privacy Settings: User Expectations vs. Reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (Berlin, Germany) (IMC '11)*. Association for Computing Machinery, New York, NY, USA, 61–70. <https://doi.org/10.1145/2068816.2068823>

[54] Jeroen GB Loman, Barbara CN Müller, Arnoud Oude Groote Beverborg, Rick B van Baaren, and Moniek Buijzen. 2018. Self-persuasion on facebook increases alcohol risk perception. *Cyberpsychology, Behavior, and Social Networking* 21, 11 (2018), 672–678.

[55] Baozhou Lu, Weiguo Fan, and Mi Zhou. 2016. Social presence, trust, and social commerce purchase intention: An empirical research. *Computers in Human Behavior* 56 (2016), 225–237.

[56] Michelle Madejski, Maritza Johnson, and Steven M Bellovin. 2012. A study of privacy settings errors in an online social network. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*. IEEE, 340–345.

[57] Albert E Mannes, Jack B Soll, and Richard P Larrick. 2014. The wisdom of select crowds. *Journal of personality and social psychology* 107, 2 (2014), 276.

[58] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.

[59] Tamir Mendel and Eran Toch. 2017. Susceptibility to social influence of privacy behaviors: Peer versus authoritative sources. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. 581–593.

[60] Marina Micheli, Elissa M Redmiles, and Eszter Hargittai. 2020. Help wanted: Young adults' sources of support for questions about digital media. *Information, Communication & Society* 23, 11 (2020), 1655–1672.

[61] Mainack Mondal, Güneş Su Yılmaz, Noah Hirsch, Mohammad Taha Khan, Michael Tang, Christopher Tran, Chris Kanich, Blase Ur, and Elena Zheleva. 2019. Moving beyond set-it-and-forget-it privacy settings on social media. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 991–1008.

[62] Brendan Nyhan and Jason Reifler. 2010. When corrections fail: The persistence of political misperceptions. *Political Behavior* 32, 2 (2010), 303–330.

[63] Ted O'Donoghue and Matthew Rabin. 2000. The economics of immediate gratification. *Journal of behavioral decision making* 13, 2 (2000), 233–250.

[64] Hyun Jung Oh and Robert LaRose. 2015. Tell me a story about healthy snacking and I will follow: Comparing the effectiveness of self-generated versus message-aided implementation intentions on promoting healthy snacking habits among college students. *Health communication* 30, 10 (2015), 962–974.

[65] Kenneth Olmstead and Aaron Smith. 2017. Americans and cybersecurity. *Pew Research Center* 26 (2017), 311–327.

[66] Kenneth Olmstead and Aaron Smith. 2017. What the public knows about cybersecurity. *Pew Research Center* 22 (2017).

[67] Ben Parr. 2011. Facebook's Big Privacy Changes: An Overview. *Mashable* (2011). <https://mashable.com/archive/facebook-privacy-changes-guide#TMdc0NovxZqX>

[68] Eyal Peer, Laura Brandomarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70 (2017), 153–163.

[69] Andrew Perrin. 2018. Americans are changing their relationship with Facebook. *Pew Research Center* 5 (2018).

[70] Emilee Rader, Rick Wash, and Brandon Brooks. 2012. Stories as informal lessons about security. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*. 1–17.

[71] Tashfiq Rahman, Rohani Rohan, Debajyoti Pal, and Prasert Kanthamanon. 2021. Human Factors in Cybersecurity: A Scoping Review. In *The 12th International Conference on Advances in Information Technology*. 1–11.

[72] Lee Rainie. 2018. Americans' complicated feelings about social media in an era of privacy concerns. *Pew Research Center* (2018).

[73] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2016. How i learned to be secure: a census-representative survey of security advice sources and behavior. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 666–677.

[74] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2019. How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1326–1343.

[75] Elissa M Redmiles, Amelia R Malone, and Michelle L Mazurek. 2016. I think they're trying to tell me something: Advice sources and selection for digital security. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 272–288.

[76] Reuters. 2008. Facebook Plans New Privacy Controls. *New York Times* CLVII, 54254 (2008), C3 N. <https://www.nytimes.com/2008/03/19/technology/19facebook.html?fta=y&referter=>

[77] P Wesley Schultz, Jessica M Nolan, Robert B Cialdini, Noah J Goldstein, and Vladas Griskevicius. 2007. The constructive, destructive, and reconstructive power of social norms. *Psychological science* 18, 5 (2007), 429–434.

[78] Josphana Shibchurn and Xiangbin Yan. 2015. Information disclosure on social networking sites: An intrinsic–extrinsic motivation perspective. *Computers in Human Behavior* 44 (2015), 103–117.

[79] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. 2013. The post that wasn't: exploring self-censorship on facebook. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 793–802.

[80] Carmen Stavrositu and Jinhee Kim. 2018. Self-Persuasion Through Mobile Applications: Exploring Different Routes to Health Behavioral Change. *Cyberpsychology, Behavior, and Social Networking* 21, 8 (2018), 516–522.

[81] Katherine Strater and Heather Richter Lipford. 2008. Strategies and struggles with privacy in an online social networking community. *People and Computers XXII Culture, Creativity, Interaction* 22 (2008), 111–119.

[82] Frederic D Stutzman, Ralph Gross, and Alessandro Acquisti. 2013. Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of privacy and confidentiality* 4, 2 (2013), 2.

[83] Shelley E Taylor and Sherry K Schneider. 1989. Coping and the simulation of events. *Social cognition* 7, 2 (1989), 174–194.

[84] Anja Thieme, Rob Comber, Julia Miebach, Jack Weeden, Nicole Kraemer, Shaun Lawson, and Patrick Olivier. 2012. "We've bin watching you" designing for reflection and social persuasion to promote sustainable lifestyles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2337–2346.

[85] Melissa C Thomas-Hunt and Katherine W Phillips. 2004. When what you know is not enough: Expertise and gender dynamics in task groups. *Personality and Social Psychology Bulletin* 30, 12 (2004), 1585–1598.

[86] Debora Viana Thompson, Rebecca W. Hamilton, and Petia K. Petrova. 2009. When Mental Simulation Hinders Behavior: The Effects of Process-Oriented Thinking on Decision Difficulty and Performance. *Journal of Consumer Research* 36, 4 (04 2009), 562–574. <https://doi.org/10.1086/599325> arXiv:<https://academic.oup.com/jcr/article-pdf/36/4/562/19296650/36-4-562.pdf>

[87] Yaacov Trope and Nira Liberman. 2010. Construal-level theory of psychological distance. *Psychological review* 117, 2 (2010), 440.

[88] Tsai-Hsuan Tsai, Hsien-Tsung Chang, Ya-Ching Chang, and Yung-Sheng Chang. 2017. Personality disclosure on social network sites: An empirical examination of differences in Facebook usage behavior, profile contents and privacy settings. *Computers in Human Behavior* 76 (2017), 469–482.

[89] Aggeliki Tsohou, Maria Karyda, and Spyros Kokolakis. 2015. Analyzing the role of cognitive and cultural biases in the internalization of information security policies: Recommendations for information security awareness programs. *Computers & security* 52 (2015), 128–141.

[90] Emily A. Vogel and Monica Anderson. 2019. Americans and Digital Knowledge. *Pew Research Center* (2019).

[91] Stanford A Westjohn, Peter Magnusson, George R Franke, and Yi Peng. 2021. EXPRESS: Trust Propensity across Cultures: The Role of Collectivism. *Journal of International Marketing* (2021), 1069031X211036688.

[92] Hu Xie, Elizabeth A Minton, and Lynn R Kahle. 2016. Cake or fruit? Influencing healthy food choice through the interaction of automatic and instructed mental simulation. *Marketing Letters* 27, 4 (2016), 627–644.

[93] Rachid Zeffane. 2020. Gender, individualism–collectivism and individuals' propensity to trust: A comparative exploratory study. *Journal of Management & Organization* 26, 4 (2020), 445–459.

[94] Min Zhao, Steve Hoefler, and Gal Zauberan. 2007. Mental simulation and preference consistency over time: The role of process-versus outcome-focused thoughts. *Journal of Marketing Research* 44, 3 (2007), 379–388.

A APPENDICES

Statement (Passive)	α
lurk on other people	0.841
learn about people by looking through their profiles	0.694
look at the profiles of people I am Facebook friends with, but not very close to	0.758
check out the profiles of people I am not Facebook friends with	0.865
Statement (Active)	α
make posts that document my life	0.887
make posts that let others know what I'm doing	0.932
monitor the reactions on my posts or comments	0.739

Table 5: Passive and Active Facebook user type statements ranked on a 5-point Likert Scale. All statements were presented in a randomized order within a matrix to finish the sentence: "I like to use Facebook to..."

		Default	Expert	Public	Control	Total
Gender	Female	10	5	10	6	31
	Male	16	22	16	19	73
Age	18-29	8	10	5	4	27
	30-39	10	12	12	13	47
	40-49	5	2	5	5	17
	50-59	1	1	2	1	5
	>=60	2	2	2	2	8
Education	Postgraduate degree	0	4	3	2	9
	Some postgraduate	1	0	1	0	2
	Bachelor's degree	15	9	10	10	44
	Associate's degree	2	3	1	5	11
	Some college	4	2	7	2	15
	High school graduate	4	9	4	6	23
Employment	Full-time	16	18	16	17	67
	Part-time	6	4	4	5	19
	Unemployed & looking	1	2	1	1	5
	Unemployed & not looking	0	1	3	2	6
	Retired	0	1	1	0	2
	Disabled	3	1	1	0	5
Income	\$100k+	1	2	1	1	5
	\$75k-99k	2	1	0	1	4
	\$50k-74k	5	4	8	11	28
	\$25k-49k	7	12	10	6	35
	<\$25k	11	8	7	6	32
Ethnicity	Black/African	3	0	3	1	7
	Caucasian	21	27	19	22	88
	East Asian	1	0	1	0	2
	Latinx/Hispanic	0	0	3	0	3
	Middle Eastern	1	0	0	0	1
	Native American	1	1	0	1	3
	South Asian	0	0	0	1	1