

A Typology of Perceived Triggers for End-User Security and Privacy Behaviors

Sauvik Das
Georgia Institute of Technology
sauvik@gatech.edu

Laura A. Dabbish
Carnegie Mellon University
dabbish@cs.cmu.edu

Jason I. Hong
Carnegie Mellon University
jasonh@cs.cmu.edu

Abstract

What triggers end-user security and privacy (S&P) behaviors? How do those triggers vary across individuals? When and how do people share their S&P behavior changes? Prior work, in usable security and persuasive design, suggests that answering these questions is critical if we are to design systems that encourage pro-S&P behaviors. Accordingly, we asked 852 online survey respondents about their most recent S&P behaviors ($n = 1947$), what led up to those behaviors, and if they shared those behaviors. We found that social “triggers”, where people interacted with or observed others, were most common, followed by proactive triggers, where people acted absent of an external stimulus, and lastly by forced triggers, where people were forced to act. People from different age groups, nationalities, and levels of security behavioral intention (SBI) all varied in which triggers were dominant. Most importantly, people with low-to-medium SBI most commonly reported social triggers. Furthermore, participants were four times more likely to share their behavior changes with others when they, themselves, reported a social trigger.

1 Introduction

A longstanding goal in usable security and privacy is to bridge the gap between behaviors that experts recommend (pro-S&P behaviors) and those that end-users actually adopt [9, 18, 29]. However, these pro-S&P behaviors remain rare. For example, as of early 2018, fewer than 10% of Google account holders had enrolled in two-factor authentication and at least 17% of Google users re-used their account passwords [33].

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2019.
August 11–13, 2019, Santa Clara, CA, USA.

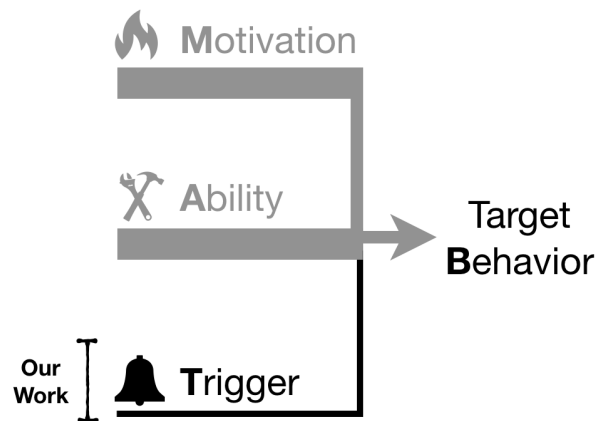


Figure 1: A popular model in behavioral psychology suggests that human behavior is a product of motivation, ability and trigger [20]. Prior work has extensively documented end-user motivation and ability to adopt S&P behaviors, but has less to say about behavioral *triggers* in-the-wild. We begin to systematically typify these behavioral triggers that lead to behavior change in S&P.

Recent Pew surveys found that only 12% of Internet users in the U.S. use password managers [36] and only 22% of smartphone users both use screen locks and regularly update their phone [4]. Additionally, Ion, Reeder and Consolvo showed that the pro-S&P behaviors that experts recommend only thinly overlap with the behaviors that non-experts find important and adopt [29].

The first step in bridging this disconnect between what experts *recommend* and what end-users *practice* is to understand what triggers pro-S&P behaviors when they *do* occur. Drawing from literature in behavioral psychology and persuasive design, a popular model suggests that behavior is a product of motivation, ability and trigger (also referred to as *prompt*) [20], as shown in Figure 1. That is, people perform

a behavior when they want to (motivation), believe they can (ability), and feel like they should *at that moment* (trigger). This framing helps identify gaps in our existing understanding of end-user S&P behavior. While there is extensive prior work that systematically typifies factors that inhibit people’s *motivation* and *ability* to adopt S&P behaviors (e.g., they are difficult [47], time consuming [24], not relevant to the task at hand [15]), there has been comparatively less work typifying the *triggers* that prompt S&P behaviors, in general.

Prior qualitative work has, however, identified both observed and hypothesized S&P behavioral triggers (e.g., [9, 10, 37, 40]). Surveying this prior work, we synthesized a set of three broad trigger types that often precede S&P behaviors—social, forced, and proactive. *Social* triggers are direct social interactions that prompt behavior change, e.g., a friend providing advice or observing others’ security behaviors. *Forced* triggers are non-social, and capture external stimuli or situations that necessitate behavior change outside of the end-users’ own volition: e.g., experiencing a personal data breach or an employer requiring one to update one’s passwords regularly. Finally, *proactive* triggers are also non-social, and capture internal processes that lead to volitional or goal-oriented behavior change: e.g., unprompted one-off decisions to enable a screen lock or routine password updates. Note that these trigger types primarily capture S&P behavioral triggers that are *perceptible in the moments leading up to behavior*. There are other “triggers” that could influence end-user S&P behaviors over longer periods of time that are harder to perceive in the moment: e.g., social norms or cultural attitudes towards openness, transparency and privacy. In this work, we primarily focus on the former as our goal is to provide actionable insights for researchers and designers.

While prior work helps categorize these in-the-moment behavioral triggers, what’s missing is a generalization of these qualitative findings to a broader sample. Indeed, what are the relative frequencies of these trigger types? How do they differently manifest for individuals from different demographic backgrounds and attitudes towards S&P? How often people actually share their S&P behaviors with others? Answering these questions is important if we are to design effective interventions that encourage pro-S&P behaviors — particularly for users who have low-to-medium security behavioral intention (SBI), or intention to behave in a manner consistent with expert-recommended security and privacy advice [18]. Moreover, while prior work has found that social triggers were particularly promising in motivating S&P behaviors for non-experts [9, 10], it remains unclear how often and for what behaviors different people encounter social triggers.

Our primary contribution is to address these gaps in the literature through an online survey ($n = 852$) in which we asked participants to report on recent S&P behaviors and what led to those behaviors. Through this process, we aim to address the following research questions:

- **RQ1:** How relatively frequent are the social, forced and proactive triggers that lead to S&P behaviors?
- **RQ2:** How does the relative frequency of social, forced and proactive triggers for S&P behaviors differ across people from different demographic backgrounds and levels of SBI?
- **RQ3:** How often and why do people share their S&P behaviors with others, and what factors correlate with this sharing?

Overall, we found that social triggers were most numerous—39% of all reported triggers were social, compared to 34% proactive and 26% forced. However, this trigger distribution varied significantly across different types of behaviors and individuals from different age groups, nationalities and levels of security behavioral intention. Perhaps most importantly, *people with low-to-medium SBI were far more likely to report changing their S&P behaviors in response to a social trigger*. Conversely, people with higher SBI were far more likely to report updating their behaviors proactively. We also found that participants were four times more likely to share their behaviors with others when their own behavior was preceded by a social trigger. In sum, our findings offer a unique new perspective in explaining end-user SP behaviors which opens up several promising new threads of research for designing tools that encourage pro-S&P behaviors.

2 Related Work

A popular model of human behavior, the Fogg Behavior Model (FBM), provides a helpful framing for understanding how to encourage pro-S&P behaviors. In brief, behavior occurs if and only if one *wants* to adopt the behavior (motivation), is easily *able* to adopt the behavior (ability) and something *prompts action* (trigger) [20]. We divide our survey of related work in usable privacy and security as they relate to these three categories of the FBM.

2.1 Ability

A broad survey of the usable privacy and security literature suggests that there are at least two barriers that reduce people’s *ability* to adopt S&P behaviors: awareness and knowledge. First, many users lack awareness of relevant security threats and what can be done to protect themselves from those threats. For example, prior studies have found that insufficient awareness of security issues resulted in people constructing their own, often incorrect, model of security threats [2, 16, 46]. Stanton et al. found that a lack of awareness of basic security principles influenced a number of security mistakes, such as using a social security number as a password [44]. In an analysis of expert and non-expert users, Ion, Reeder and

Consolvo [29] found that non-experts were unaware of the strategies experts employed to protect themselves.

Security tools are also often too complex for end-users to operate [2, 47]. Indeed, for many security and privacy systems, there is a wide gulf of execution [35] between what users *want* and *know how* to do. For example, many users cannot distinguish legitimate versus fraudulent URLs, nor forged versus legitimate email headers [13]. Another study revealed how security features in Windows XP, Internet Explorer, Outlook Express, and Word applications are difficult for users to understand and utilize [21]. Wash found that many people hold “folk models” of computer security that are often incorrect, which leads to ignoring security advice [46]. More recently, the Pew Internet Research center found that the majority of Internet users have strong misconceptions about basic cybersecurity concepts [36].

2.2 Motivation

In addition to being unable, people may also simply not want to follow recommended security advice [9]. This lack of motivation can be attributed to a number of key psychological principles. First, stringent security measures are often antagonistic towards the specific goal of the end user at any given moment [15, 41]. For example, strong e-mail account security (e.g., using two-factor authentication), can delay a user access to her email for an intolerable amount of time [17]. Thus, users may reject SP advice when they expect or experience it to be too time-consuming or require too much effort [2, 15, 28, 41].

Furthermore, many people may understand security threats in the abstract but may not believe they, themselves, are at risk [2, 46]. Herley argues that this perspective may be rational, as the expected monetized cost of a lifetime of following commonly recommended security advice (e.g., reading suspicious URLs) may be orders of magnitude higher than the expected monetized loss a compromised account [25]. Furthermore, while the benefits of security features are abstract and delayed (i.e., protection from a potential threat sometime in the future), the costs are immediate and concrete (i.e., additional time or effort now and forevermore) [1]. Indeed, security claims are often unfalsifiable — irrespective of present behavior, there is no guarantee of future security [26]. Furthermore, prior work has found that there may be a social stigma associated with use of expert recommended security tools and advice that further lowers people’s motivation to be secure [9, 11, 22].

2.3 Triggers

While prior work provides a rich foundation for understanding motivation and ability in the context of end-user S&P behaviors, we have less understanding of the triggers that prompt S&P behaviors and how they vary in frequency and effectiveness across individuals. This is not to say that there is

no work on behavioral triggers in S&P. For example, there is much work on improving adherence to and compliance with security warnings, as these warnings are commonly ignored and begrudged [3, 5, 10, 16, 19]. Similarly, there is a wealth of information about how to design privacy notices for different use-cases and scenarios (for a review, see [42]). Much of the research on security warnings and privacy notifications is centered around urging end-users to react to an external prompt — what we call “forced” triggers.

In a recent qualitative study, Das et al. introduced a typology of *social* triggers for S&P behaviors [9]. They found that nearly 50% of all reported S&P behavior changes were the result of an implicit or explicit social interaction with others [9]. Among these social triggers were “observing others”, “sharing access with others” and “receiving advice from others.” Others have also noted the broad efficacy of social triggers for S&P behavior change. Both Rader and Redmiles et al. found that informal word-of-mouth stories were effective social triggers for S&P behaviors [37, 40].

Prior work has also shown that people can sometimes be *proactive* in their S&P behaviors. For example, a series of studies found that people can be proactive in self-censoring content or otherwise adjusting their social media privacy settings to avoid later regrets [8, 43, 45].

There is also evidence that the efficacy and frequency of different S&P triggers might vary across individuals. Redmiles et al. [38] found that people from different socioeconomic backgrounds may respond to differently to advice received from different sources (e.g., notices from the workplace vs. informal stories from friends). There is also a growing body of evidence documenting how people from different nationalities and cultural contexts can have different S&P attitudes and behaviors [30, 32].

To date, most of our knowledge of S&P behavioral triggers, in general, is piecemeal — assembled together from an ensemble of studies that either implicitly note the presence of these triggers or that more thoroughly study a specific trigger. To our knowledge, we are the first to quantitatively explore and synthesize S&P behavioral triggers generally. In doing so, we provide insights on: (i) the relative frequency of in-the-moment, perceived triggers that inspire S&P behavior change; (ii) how those triggers might vary across different S&P behaviors; and, (iii) how different individuals respond to different triggers. Armed with this understanding, we make empirically grounded suggestions on how to effectively design behavioral triggers that encourage pro-S&P behaviors.

3 Method

We conducted an online survey on the Amazon Mechanical Turk (AMT) platform ¹. We selected AMT partially because of the ease of recruiting a large sample on the platform, and

¹<https://mturk.com>

partially because the biases of AMT samples are well studied [23,30,39]. To ensure high-quality responses, we included two attention-check questions, or questions for which participants are given specific instructions on how to answer to gauge if they are carefully reading questions [23]. We only discuss participants who passed these attention checks. The specific questions we asked in our survey is provided in Appendix A. For brevity, we provide a high-level overview of the questions.

Behavior change questions: We started by asking participants which, if any, of the following four behaviors they did in the past 6 months. The behaviors we selected were:

- *Mobile Auth:* enabling or changing one’s method of authenticating into a mobile device (e.g., smartphone, laptop, tablet or other portable electronic device);
- *App Uninstallation:* uninstalling a smartphone application, specifically for privacy or security reasons;
- *Password Updates:* changing or updating a password for an online account; and,
- *Facebook Privacy:* updating one’s Facebook account privacy settings.

We selected *these* behaviors because they were chosen in the closest related prior work [9], a qualitative exploration that found that nearly 50% of all reported behavioral triggers for the aforementioned behaviors were social. We wanted to compare our own results to that benchmark. Given that the selected behaviors still represent a diverse subset of S&P behaviors, we believe that our results should generalize as well as any other subset of S&P behaviors.

If participants had not recently done any of the aforementioned behaviors, they were allowed to manually specify a different S&P behavior they recalled doing in the past 6 months. They could answer remaining questions in reference to this “other” behavior.

Trigger questions: For each S&P behavior participants recalled having done in the past 6 months, we next asked participants which, if any, of a set of behavioral triggers preceded their decision to perform the behavior. The options we presented were synthesized from a survey of related work. Participants were also able to manually write-in a different trigger if the provided options were insufficient.

We categorized each of these triggers into three higher-level categories that we synthesized from a reading of prior work and a discussion among the authors — social, forced and proactive. Social triggers are those that involve a direct social interaction either with somebody the participant knew personally or with whom the participant could observe and/or interact (i.e., experiencing a security breach from someone who the participant knew, lending one’s device to someone else, observing others around them, or receiving advice). Forced triggers are non-social, and suggest the presence of an external catalyst that the participant did not specifically seek

or desire (i.e., a warning dialog, an organizational policy, experiencing a personal data breach from a stranger). Finally, proactive triggers, while also non-social, involve conative processes internal to the participant or voluntarily seeking out information that directly leads to behavior change (i.e., habit or routine, no specific stimulus, reading news articles, actively looking through device settings or options).

We note that this higher-level taxonomy may have blurred boundaries: some triggers, like changing one’s PIN due to lending one’s device to a friend, could be considered either “social” or “proactive”. However, we categorized each individual trigger into one higher-level category using the following well-defined process. First, if the trigger reflected any direct social influence or interaction, we categorized it as “social”, even if it might also be “forced” or “proactive.” We make this distinction based on findings from prior work which suggest that social triggers are uniquely motivating [9]. If the trigger did *not* reflect a direct social process, we categorized it as “forced” if the behavior change was either mandatory or forced by circumstance. Otherwise, we categorized it as “proactive,” as a non-social, non-forced trigger suggests that participants made the change voluntarily either because of routine, because of personal preference, or because they actively sought out information.

Table 1 shows a list of all trigger options we presented, their mapping to the higher-level categories of social, forced and proactive, as well as the overall percentage of participants who reported having experienced the trigger prior to enacting the behavior. Table 2 shows the distribution of the higher-level trigger types, both overall and across individual behaviors.

Social context questions: For each of the social triggers participants selected, we asked additional questions to uncover the social context of those triggers. For example, if participants selected the “Received Advice” trigger, they were asked to specify their relationship with the person from whom they received the advice: friend, family member, significant other, colleague or other. If participants selected other, they were allowed to manually write-in a description of their relationship with that person. Participants who did not select any social trigger would not see any of these questions.

Sharing questions: We also asked participants if they shared their behavior with others. If they did share, we asked them to specify with whom (e.g., friend, family member, etc.) and how they shared the change (e.g., through face to face conversation, phone call, SMS or email). We then asked participants why they shared the change, giving them a range of options largely derived from prior work [9]. Examples reasons include “I noticed they were being insecure” and “I felt obligated to protect them”. For brevity, we omit the complete list of responses here, but list them in Appendix A. We also discuss participants’ rationale to share (Table 5.3) and not share in more detail in our results. Participants were, again, allowed to manually write-in an answer if none of the provided options were sufficient.

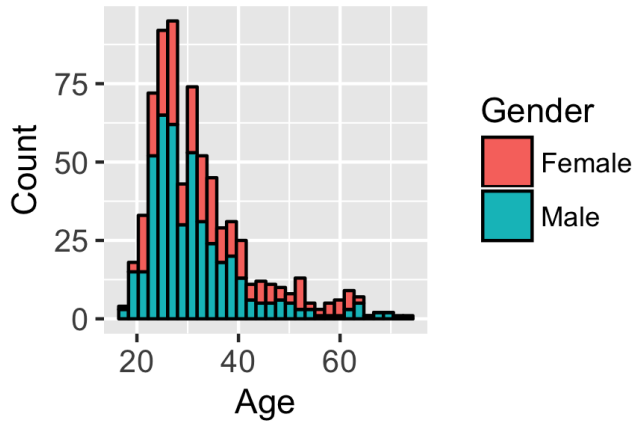


Figure 2: Distribution of participant ages and genders. Our participants were 33 years old, on average, and 63% self-reported as male.

Security behavioral intention questions: We next asked participants to answer Egelman and Peer’s SeBIS questionnaire [18] to measure their security behavioral intention (SBI). In addition, we asked participants a number of other questions about their general security knowledge and computer literacy derived from prior work.

Demographic questions: Finally, we asked participants to self-report a number of demographic dimensions, such as their age, gender and nationality. They were allowed to opt-out of providing any of this information, but nearly all participants answered all of the demographic questions.

3.1 Ethics and Compensation

Prior to data collection, we had our study approved by the Carnegie Mellon University Institutional Review Board (IRB). The survey took participants about 20 minutes to complete on average, and we paid participants \$3.50 (translating to an hourly wage of \$10.50). All collected data was anonymized — no identifiers were collected, and payment was facilitated through AMT.

4 Sample

Overall, we received responses from 1070 participants, 852 of whom both passed the attention-check quality tests and completed the entire questionnaire. Accordingly, we were left with $n = 852$ high-quality, complete responses.

4.1 Demographics and behavioral intention

Our participants had a mean age of $\mu = 33$ (sd: 10), ranging from 18 to 74. In addition, 537 participants self-reported

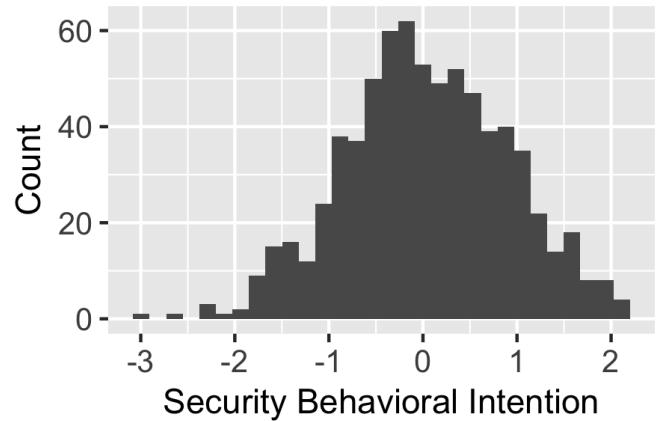


Figure 3: Distribution of security behavioral intention across all participants.

as male (63%), 312 as female (36%) and 3 preferred not to answer. Figure 2 shows participants’ age and gender distributions. Given the constraints of the AMT platform, our participants came mostly from the United States (449) and India (323). Approximately 47% of our participants reported a primary occupation that was “Computer Science related” or “Other engineering or technology related”. However, only 10 reported occupations directly related to cybersecurity. Finally, 96% of participants reported being native English speakers.

To facilitate later analysis, we used a factor analysis to reduce the dimensionality of the 16-item SeBIS questionnaire [18] into one higher level construct we refer to as “security behavioral intention” (SBI). This single factor captured 17% of the variance in the responses to the SeBIS questionnaire, with each item being correlated with the construct in the expected direction (i.e., positively coded questions positively correlated, negatively coded questions negatively correlated). Figure 3 shows a histogram of the SBI in our population. Note that this method of dimensionality reduction has been previously used to facilitate analysis with the SeBIS [12].

4.2 Raw response counts for behaviors and triggers

Behaviors: Out of our 852 participants, in the 6 months preceding the survey, 454 (53%) reported having changed their mobile / laptop authentication settings, 427 (50%) reported changing their Facebook privacy settings, 378 (44%) reported uninstalling a mobile application for security and privacy reasons, and 688 (81%) reported changing a password for an online account or device. Overall, 807 participants (95%) reported doing at least one of the aforementioned four S&P behaviors and reported on 1947 behaviors, in total.

Triggers: Table 1 shows the distribution of triggers se-

lected for different behaviors. Across the 1947 behavior changes in our dataset, participants reported 2954 triggers leading up to those changes. A large majority of the triggers that lead to a security or privacy behavior were covered by the options we provided, which were based off a survey of prior work. Indeed, only 2% of reported triggers, overall, warranted manual specification that was not covered by our initial typology. Upon deeper investigation of these manual entries, most correlated strongly with the available triggers choices. The most commonly specified trigger that was not well covered by the questionnaire responses was “habit / routine” — a number of participants reported periodically, habitually or routinely updating their passwords, browsing their Facebook privacy settings, etc. We considered habitual / routine updating of S&P behaviors to be a proactive trigger.

5 Results

Recall that we had three high-level research questions we wanted to answer: (RQ1) How relatively frequent are social, forced and proactive triggers for S&P behaviors? (RQ2) How does the relative frequency of social, forced and proactive triggers for S&P behaviors differ across people from different demographic backgrounds and levels of SBI? and (RQ3) How often and why do people share their S&P behaviors with others, and what factors correlate with sharing? We present empirical answers to each of these questions.

5.1 RQ1: Relative trigger frequency

Table 2 shows the relative trigger frequency across all four behaviors. Overall, 1153 (39%) of reported triggers were social, 1005 (34%) were proactive and 773 (26%) were forced. To test if these rate differences were significant, we ran a logistic regression correlating trigger presence with trigger type (social, forced, proactive) and included random-intercept terms for distinct users and distinct behaviors to account for repeated observations (i.e., multiple behaviors per user, multiple triggers per behavior). We found that each of the pairwise rate differences were statistically significant ($p < 0.001$).

While different S&P behaviors vary in how often they are prompted by social, forced and proactive triggers, social triggers were the most frequent overall. This result highlights the importance of understanding and leveraging social influence to encourage better S&P behaviors. Indeed, out of all three of the higher-level trigger types, social triggers may be the most actionable. While many usable security interventions have attempted to make people more proactive about their security and privacy to little avail, the design space for encouraging greater social interaction in security and privacy is sizable and but is only just beginning to be explored [6, 10, 34].

We also found that our participants reported being surprisingly proactive in engaging with their security and privacy. Indeed, proactive triggers were the second most frequently

reported triggers leading to S&P behaviors. Also surprisingly, forced triggers were least frequent. This duality of results is promising, in theory — we want people to be more proactive about S&P and to avoid forcing compliant behaviors. However, since few people use two-factor authentication [33] or password managers [36] or regularly update their software [29], there is clearly much room for improvement.

One limitation in interpreting these results is that because people could select multiple triggers leading up to a behavior change, it’s difficult to say which trigger played the most important role in convincing someone to change their behavior. Accordingly, the best we know is that these triggers could have played *some* role. A more general limitation is that because participants may have been several months removed from the event, their memory of the relative order of these triggers and their behavior may be muddled.

5.2 RQ2: Individual differences

We next empirically modeled how S&P behavioral triggers varied across individuals and behaviors using a series of random-intercepts logistic regressions. Specifically, we modeled how likely a participant was to report a social trigger, a forced trigger and a proactive trigger given their age, gender, nationality, security behavioral intention and the type of behavior they reported having changed. Due to location restrictions of the AMT platform, we filtered out 83 participants who did not identify as being from the U.S. or India as we did not have enough data for other nationalities.

We used a random-intercepts term for each participant to account for repeated observations. We calculated the six pairwise comparisons between the four different behaviors using a contrast matrix with R’s multcomp package [27]. Significance levels were corrected using the Bonferroni method. Table 3 shows the results.

Coefficients for the numeric covariates (i.e., age, SBI) indicate a change in log odds that a participant reported a particular trigger leading up a behavior change. A positive coefficient implies that the log odds of a participant reporting a particular trigger increases as the predictor variable increases by one standard deviation, while a negative coefficient implies the opposite. For example, the social trigger regression in Table 3 shows that age has a negative coefficient ($b_{age}^{social} = -0.10$). Thus, for every one-standard deviation increase in age, the model estimates that a participant’s log odds to have reported a social trigger should decrease by 0.10 (i.e., younger people are more likely to report social triggers).

For categorical covariates (i.e., behavior types, gender, nationality), coefficients represent the difference in log odds to have experienced a particular trigger between participants at different levels of the covariate. For example, Table 3 shows that the coefficient for a participant from the U.S. to report a proactive trigger versus a participant from India is $b_{US}^{proactive} = 0.81$. As the coefficient is positive and large, we

Behavior trigger	Abbrev	Type	Mobile Auth	App Uninstall	Password Update	Facebook Privacy
I directly experienced a security breach from someone I know	Breach by Known	Social	5%	6%	4%	5%
I allowed someone to use my device or account previously	Shared Access	Social	15%	N/A	8%	21%
I observed people around me doing this	Observed Others	Social	8%	14%	7%	11%
Someone I know advised me to do this	Received Advice	Social	14%	16%	9%	12%
Other Social	Other Social	Social	2%	< 1%	< 1%	< 1%
I directly experienced a security breach from a stranger	Breach by Stranger	Forced	5%	5%	9%	6%
My device or account prompted me to do this	Device Prompt	Forced	9%	9%	22%	9%
My organization required me to do this	Org Prompt	Forced	6%	3%	6%	3%
Other Forced	Other Forced	Forced	<1%	< 1%	3%	2%
I looked through settings / options to do this	Browsed Settings	Proactive	13%	N/A	5%	15%
Nothing really happened	No Trigger	Proactive	7%	26%	13%	5%
I read a news article about the security vulnerability or recommending a best practice	Read News	Proactive	15%	11%	13%	12%
Other Proactive	Other Proactive	Proactive	<1%	< 1%	< 1%	3%

Table 1: Behavioral triggers, classified into three higher-level types: social, forced, and proactive. Trigger rates for each behavior are provided in the last four columns. The dominant trigger(s) for each behavior is highlighted in green.

	Mobile Auth	App Uninst.	Changed Pwd	FB Priv.	Over-all
Social	375 (43%)	131 (43%)	273 (29%)	374 (48%)	1153 (39%)
Re-active	179 (21%)	66 (19%)	375 (39%)	153 (20%)	773 (26%)
Pro-active	315 (36%)	122 (36%)	312 (33%)	256 (33%)	1005 (34%)

Table 2: Trigger frequency across all four S&P behaviors individually and overall. The dominant trigger type for each behavior is highlighted in green.

can conclude that participants from the U.S. are much more likely than participants from India to report a proactive trigger leading up to a S&P behavior.

Generally, Table 3 shows that there are many significant correlations between behaviors, demographics and how likely one is to report a social, forced or proactive trigger leading up to a S&P behavior. We discuss each key finding, in turn.

Security Behavioral Intention: There was a strong correlation between SBI and the triggers participants’ reported leading up to their behaviors. Unsurprisingly, people with higher SBI were more likely to report proactive triggers ($b_{sbi}^{proactive} = 0.36, p < 0.001$), while people with lower SBI were more likely to report forced triggers ($b_{sbi}^{forced} = -0.25, p < 0.001$). Perhaps most importantly, we also found that people with lower SBI were more likely to report a social trigger ($b_{sbi}^{social} = 0.12, p < 0.05$).

Figure 4 shows the relationship between SBI and a participants’ likelihood of reporting a social, forced or proactive trigger. The likelihood is calculated from our estimated random-intercepts logistic regression model, which also takes into account participants age, gender, nationality and the behavior type. The trend lines are fit using a Gaussian Additive Model, which allows us to model non-linearities in the relationship. We can see a clear trend — controlling for all of the other covariates, people with low-to-medium security behavioral intention (-1.5 to 0.5) are much more likely to report a social trigger leading up to a S&P behavior. From Figure 3, we know that most people (> 65%) fall into this low-to-medium range.

Taken together, we found strong empirical evidence sug-

	Social	Forced	Proactive
Intercept	0.31	-1.38**	-1.59**
<i>Individual comparisons</i>			
SBI	-0.12*	-0.25**	0.36**
Age	-0.10*	-0.06	0.15*
Male (vs. Female)	-0.17	0.09	0.10
US (vs. India)	-0.87**	0.09	0.81**
<i>Behavior comparisons</i>			
Pwd (vs. App Uninst.)	-0.44*	1.05**	-0.24
MAuth (vs. App Uninst.)	0.13	0.12	0.12
FB (vs. App Uninst.)	0.37*	0.07	-0.14
MAuth (vs. Pwd)	0.57**	-0.93**	0.35*
FB (vs. Pwd)	0.81**	-0.99**	0.10
FB (vs. MAuth)	0.24	-0.06	-0.26

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

Table 3: Logistic regression coefficients comparing how often social, forced and proactive triggers were reported as behavioral triggers for different participants and for different behaviors. Bonferonni correction was used to account for multiple testing. Baseline comparison groups are indicated in parentheses for categorical variables. We used R’s multcomp package to compute the six pairwise differences for the four behaviors.

gesting that social triggers are *especially* effective S&P behavioral triggers for the majority of people who have low-to-medium security behavioral intention. Yet, to date, most end-user facing security and privacy systems do not take into account social factors or encourage social interaction. A strong implication for design, then, is to create systems that encourage greater social interaction so that it is easier to reach people with low-to-medium SBI.

Age and Gender: There were strong correlations between age and the triggers that reportedly lead up to S&P behaviors. Younger people were more likely to report social triggers ($b_{age}^{social} = -0.10, p < 0.05$) and older people were more likely to report proactive triggers ($b_{age}^{proactive} = 0.15, p < 0.01$). We found no significant correlations between gender and behavioral triggers. Additional research may be needed to determine causality, but our results suggest that some level of age-based personalization may be needed to trigger pro-S&P behaviors.

Nationality: We found a strong correlation between self-reported nationality and reported S&P behavioral triggers. People from the U.S. were much more likely to report being individually proactive about their security ($b_{U.S.}^{proactive} = 0.81, p < 0.001$), whereas people from India were much more likely to report social triggers ($b_{U.S.}^{social} = -0.87, p < 0.001$).

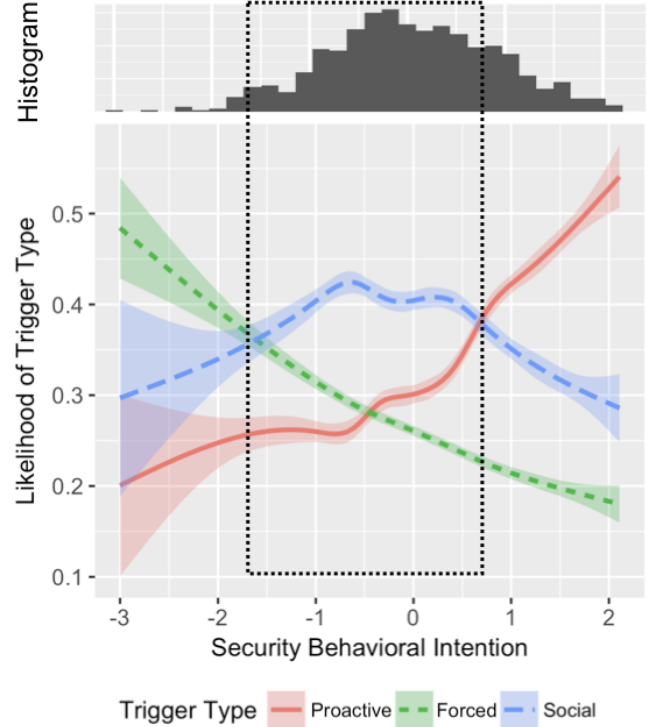


Figure 4: Estimated likelihood of reporting a social, forced or proactive trigger for participants with different SBI. The trend lines, and the 95% confidence intervals, were fit to a Gaussian Additive Model. People with high SBI were more likely to report a proactive trigger, while the majority of people with low-medium SBI were more likely to report social triggers. The dashed boxed outlines the SBI range in which social triggers were most prevalent to facilitate cross-referencing with the SBI histogram above.

Figure 5 visualizes the distribution of the estimated likelihood of reporting a social, forced or proactive trigger for participants from the U.S. vis-a-vis those from India. The estimated likelihoods are calculated from the logistic regression in Table 3, and take into account the other covariates in the regression. We can see a clear separation in the social and proactive distributions, with the former favoring people not from the U.S. and the latter favoring people from the U.S.

These findings echo those of prior work modeling differences in the privacy attitudes of Mechanical Turk workers in India vs. the US [30]. While it’s tempting to attribute these effects to cultural differences, our findings do not imply causality. Additional research will be necessary to tease apart the effect of culture from other confounding factors such as, for example, the work contexts of AMT workers in the U.S. versus those in India.

Behaviors: Different behaviors had significantly differing trigger distributions. While the raw numbers are pre-

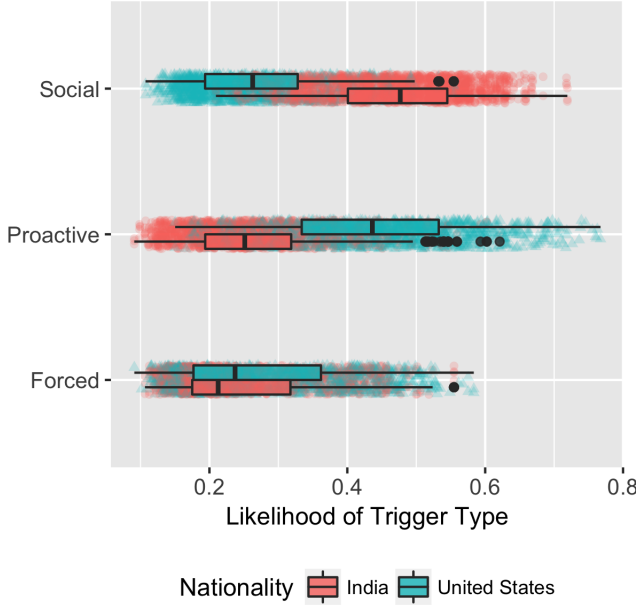


Figure 5: We plot the estimated likelihood that people from the U.S. and from India were to report social, proactive or forced triggers. These likelihoods are estimated from the random-intercepts logistic regressions shown in Table 3, additionally accounting for age, sbi and behavior type. People from the U.S. were more likely to report proactive triggers, while people from India were more like report social triggers.

sented in Table 2, the regression analysis uncovered statistically significant differences across behaviors controlling for age, gender, nationality and SBI. Mobile authentication changes were significantly more likely to have a reported social ($b_{MobvPwd}^{social} = 0.57, p < 0.001$) and proactive ($b_{MobvPwd}^{proactive} = 0.35, p < 0.001$) trigger than changing passwords. Changing passwords was significantly more likely to have a reported forced trigger than mobile authentication changes ($b_{MobvPwd}^{forced} = -0.93, p < 0.001$), changing Facebook privacy settings ($b_{FBvPwd}^{forced} = -0.99, p < 0.001$) and uninstalling applications ($b_{PwdrvApp}^{forced} = 1.05, p < 0.001$). Uninstalling applications was more likely to have a reported social trigger than changing passwords ($b_{PwdrvApp}^{social} = -0.44, p < 0.01$). Finally, changing Facebook privacy settings was more likely to have a reported social trigger than uninstalling applications ($b_{FBvApp}^{social} = 0.37, p < 0.05$) and changing passwords ($b_{FBvPwd}^{social} = 0.81, p < 0.001$).

5.3 RQ3: Sharing patterns

Conversations and interactions about security are rare and avoided by both experts and non-experts alike [9, 11]. Yet, social triggers cannot be produced without some form of active

	Mobile Auth	App Uninst.	Changed Pwd	FB Priv.	Overall
Overall Shared	137 (30%)	173 (46%)	81 (12%)	222 (53%)	613 (32%)
Family	66	40	43	64	213
Friend	91	44	82	89	306
Colleague	42	16	12	12	82
S.O.	38	17	32	54	141
Other	3	1	4	3	11

Table 4: Number of people who shared their behavior changes across different behaviors and overall. Participants could select multiple audiences. The first row indicates the total number of those behaviors that were shared.

or passive social interaction. Accordingly, we next wanted to understand when and why people share their security behaviors with others to see if there may be untapped opportunities to encourage greater sharing.

Table 5.3 shows how many participants decided to share their reported behavior changes with others, both overall and with specific other relations (e.g., friends, family, colleagues and significant others). Overall, 32% of reported behavior changes were shared with others — primarily with friends and to a lesser degree with family and significant others. This overall sharing rate is in line with prior work on people’s willingness to share news articles about security and privacy, which found that 29% of MTurkers reported sharing such articles with friends and family [12]. We suspect the actual rate of sharing S&P behaviors may be lower in practice, but that the behavior changes participants were reporting on were especially salient and thus more likely to be shared.

We found a large difference in the sharing rate of different S&P behaviors. The most shared behavior was updating one’s Facebook privacy settings (53% share rate). This result was unsurprising, given the inherent social nature of Facebook and its salient privacy settings. Conversely, changing passwords was least likely to be shared (12% share rate). This contrast suggests that there remains a significant opportunity to develop systems that encourage more explicit social interactions between individuals, especially for behaviors made outside of a social platform such as Facebook. Indeed, as people with low-to-medium SBI appear to respond especially well to social triggers and are rarely proactive, a high-level goal should be to encourage more social interactions and greater observability of S&P behaviors more generally, albeit with the ability to maintain individual privacy as desired.

We next wanted to explore why people *did* and *did not* share their behaviors with others. If we have a better understanding of the reasons people share their S&P behaviors, we

	Mobile Auth	App Uninst.	Changed Pwd.	FB Priv.
I noticed they were being insecure	15%	14%	12%	33%
They learned about a new security tool	14%	9%	9%	N/A
I felt obligated to protect them	13%	17%	18%	N/A
They experienced a breach	12%	11%	11%	N/A
They had to set up a new device, account or tool	7%	4%	6%	N/A
They read a news article about security	11%	9%	8%	23%
I just wanted to talk about my recent change	15%	22%	21%	43%
They noticed that I made a change	12%	13%	12%	N/A
No reason	1%	0%	0%	0%
Other	1%	2%	2%	1%

Table 5: Reasons people decided to share that they had made a security and privacy behavior change with others. Many people mentioned being vigilant of others’ S&P and feeling obligated to protect them. These rows are highlighted in green.

may be able to design targeted systems and interventions that encourage more explicit social interactions. Table 5.3 lists why people elected to share their behaviors with others.

The most commonly reported reason to share was non-descriptive: “I just wanted to talk about my recent change.” We included this option for participants who could not select a more specific reason for why they shared their behavior. However, the second and third most commonly reported reasons across all behaviors was that people felt an obligation to protect others and because participants were vigilant of other’ being insecure. These findings suggest that people often share their S&P behaviors with others because they feel a sense of accountability or obligation for the security of their friends and loved one, as has been alluded to in past work [9]. However, as has been previously reported, there are very few systems in place that allow people to act on this sense of accountability for their friends and loved ones [11]. Furthermore, the low observability of S&P behaviors places a strong burden on early adopters to explicitly share their behaviors with others if those behaviors are to spread.

	Coefficient	p-value	
Intercept	-0.16	0.62	
Social Trigger?	2.31	<0.001	**
<i>Individual differences</i>			
SBI	0.07	0.41	
Age	0.002	0.79	
Male (vs. Female)	-0.10	0.52	
US (vs. India)	-1.10	<0.001	**
<i>Behavior differences</i>			
Pwd (vs. App)	-2.83	<0.001	**
MAuth (vs. App)	-1.94	<0.001	**
FB (vs. App)	-0.65	0.01	*
MAuth (vs. Pwd)	0.89	<0.001	**
FB (vs. Pwd)	2.19	<0.001	**
FB (vs. Mob)	1.23	<0.001	**

Table 6: Regression coefficients comparing how the decision to share one’s new security behavior correlates with one’s SBI, demographics, whether or not the behavior was socially triggered, and the type of behavior being shared. Bonferonni correction was applied. Baseline comparison groups are indicated in parentheses for categorical variables. We used R’s multcomp package to compute the six pairwise differences for the four behaviors.

The primary reasons *not to share*, unsurprisingly, centered around a general lack of desire to share (38%) and an assumption that other people did not need to know anything about one’s S&P behaviors (34%). If we are to increase the prevalence social triggers, these results suggest that we should make S&P systems that encourage social sharing and that are more easily observable so that early adopters do not need to explicitly share their behaviors.

To better understand what factors lead to sharing S&P behaviors, we ran a mixed-effects logistic regression correlating if a participant shared their reported S&P behavior with their age, gender, SBI, nationality, whether or not their behavior was socially triggered, and the type of behavior. The results are shown in Table 6. Coefficients can be interpreted in the same way as in the models reported in Table 3. We found strong, significant correlations as outlined below.

Nationality: People from the U.S. were far less likely to share than people from India. ($b_{U.S.}^{share} = -1.10, p < 0.001$). This lack of sharing could also explain the stark difference in social triggers as a catalyst for behavior in the U.S. versus India, but further research is necessary for this to be conclusive.

Behavior type: All pairwise differences between the sharing rates of distinct behavior types were significant. Com-

bined with the raw counts of sharing by behavior presented in Table 5.3, it looks like changing Facebook privacy settings is shared most frequently, followed by app uninstallations, mobile authentication changes and, finally, password updates.

Behavior prompted by a social trigger: Finally, if participants reported changing their behavior as a result of a social trigger, they were much more likely to share information about that behavior with others ($b_{social}^{share} = 2.31, p < 0.001$). Concretely, 56% of behaviors that had a reported social trigger were shared with others, compared to just 14% of behaviors that were not — a four-fold increase.

6 Discussion

6.1 Summary of Results and Contributions

Most generally, we found that social triggers (39%), in which people were influenced by others, were the most frequent reported catalysts for S&P behaviors. Proactive triggers (34%), where people individually decided to make a change independent of an external prompt or breach, were second most frequently reported. Finally, forced triggers (26%), where people made a change in response to a specific breach or news event, were least frequently reported.

While our aggregate results paint a simple picture, once we drilled down into differences between people from different backgrounds and across different behaviors, we uncovered a more nuanced story. Specifically, we found that individual and behavioral differences correlate strongly with which triggers participants reported. Indeed, people with high security behavioral intention were most likely to report proactive triggers, but people with *low-to-medium* SBI, who make up the vast majority, were much more likely to report changing their behavior in response to a social trigger. Demographics also correlated with reported behavioral triggers — younger people and people from India were much more likely to report changing their behavior in response to a social trigger, while older people and people inside the U.S. were much more likely to report changing their behavior proactively.

In analyzing when and why people shared their own security and privacy behaviors with others, we found that people who themselves reported social triggers were far more likely to share their behaviors with others. We also found that people in India were much more likely to share their behaviors with others than people in the U.S., and that different behaviors are shared at different rates — specifically, uninstalling applications for security and privacy reasons was shared most often, followed by updates to Facebook privacy settings, changes to mobile device security and, lastly, password updates.

Finally, we also found that while most people do not share their S&P behaviors with others because they just do not want to, when people *do* share their behaviors they do so because they feel a sense of accountability for or obligation to protect their friends and loved ones.

6.2 Design Implications

Our work contributes the first large quantitative analysis comparing the relative frequency of self-reported S&P behavioral triggers and how those triggers vary across individuals from different backgrounds and behavior types. We now reflect on some actionable design implications.

Designing security and privacy systems that encourage social interaction: The highest-order bit of our results is a hypothesis — to encourage more widespread use of pro-S&P behaviors by non-experts, these behaviors should be designed to be more passively observable or to encourage greater active social interaction. In other words, we hypothesize that to encourage pro-S&P behaviors, we should design systems and interventions that facilitate social triggers.

The basis for this hypothesis is two-fold: first, social triggers were the most frequently cited prompts for S&P behaviors, in aggregate, and were *especially so* for people with lower security behavioral intention; and, second, people who reported changing their behavior as a result of a social trigger were four times more likely to share their own behaviors with others, in turn. Accordingly, by making more social systems we may be able to bootstrap a feedback loop in which social triggers lead to behavior change, which, in turn, should lead to even more social triggers.

We note that our call to make security more social is not new — prior work has also made similar suggestions [9, 11, 14, 31]. Still, our work adds to an emerging chorus of research illustrating the importance of considering social factors in the design of end-user facing S&P systems.

How can we design systems and interventions that encourage more social sharing? Prior work suggests that by making security systems that are more observable (i.e., a system that is easily seen by others when it is used), cooperative (i.e., a system that allows people to work together towards mutually beneficial ends) and stewarded (i.e., a system that allows one person to act for the benefit of others), people are more likely to both actively engage in social interactions about S&P as well as passively observe others' S&P behaviors [7]. Of course, such systems should also respect the individual privacy preferences of those who would prefer not to be identifiable in social cues to others. For end-user communities who would prefer their individual S&P behaviors to be private, aggregate social cues where no individual is identified may be one effective path forward [10, 14].

Exploring a broader design space for S&P triggers: Fogg defines three types of behavioral triggers for persuasive design [20]: *sparks*, which motivate people with high ability but low motivation; *facilitators*, which simplify action for people with high motivation but low ability; and, *signals*, which serve as reminders for people who already have high motivation and ability. Many existing S&P warnings and notifications are *signals*. Sparks and facilitators also pose interesting opportunities for S&P, as few end-users have both

high motivation and high ability to engage in pro-S&P behaviors. An example of a *spark* that encourages S&P behaviors is Das et al.’s social proof notifications, which informed Facebook users of the number of their friends who used optional security tools on Facebook [10]. An example of an effective *facilitator* that simplifies S&P behaviors comes from Akhawe and Felt’s redesign of the Chrome SSL warning to simplify exiting out of suspicious webpages [3].

This prior work has only begun to explore a rich design space for sparks and facilitators. We foresee opportunities to co-opt an end-users’ social and environmental contexts to create better sparks and facilitators (e.g., trending S&P behaviors in the wake a publicized incident, aggregated social proof cues of others’ S&P behaviors in the same room).

Personalized behavioral triggers: Our results also illustrate that behavioral triggers may need to be personalized to people from different cultural contexts, demographic backgrounds, and levels of SBI. The growing body of literature on modeling individual differences in S&P paints a nuanced picture of the varying desires, attitudes and assumptions of different groups of people with respect to S&P. Our results inform the need to individually tailor triggers that prompt people to act in a manner consistent with expert S&P advice.

6.3 Limitations

As with any study, ours has a number of limitations that are important to keep in mind when interpreting the results. First, our dataset has biases. Specifically, our sample over-represents males and people in technology related fields and occupations. We suspect this is the result of a self-selection bias in who decided to fill out our survey, as other AMT studies tend to have more gender balance (e.g., [12, 30]). The upshot is that our population probably over-represents those with high SBI. In turn, we expect that due to this sample skew, proactive triggers should account for a smaller proportion of reported behavioral triggers in a more representative sample. While these biases are important to consider in interpreting our results, our choice of sample still provides generalizable insights. Indeed, while some prior work suggests that MTurk workers tend to be more concerned about privacy than a U.S.-census representative sample [30], more recent work found that an MTurk sample was more representative of the U.S. population in terms of privacy and security experience, knowledge and advice sources than a census-representative web panel [39].

A second limitation is that our survey is primarily based on self-report and recollection. We asked participants about recent behaviors that occurred but it’s likely that their recollection of these behaviors and their triggers is imperfect. This limitation may also contribute to higher-than-expected reporting of proactive triggers — people who cannot recall what factors lead up to a behavior change may simply attribute the change to their own independent judgment. In future work, it would be useful to catch behavior changes closer to the

moment those behaviors occur, perhaps through a diary study.

Our data captures a limited subset of S&P behaviors, though we expect it to generalize as well as any other subset. There are many other S&P behaviors we did not ask about — e.g., two-factor authentication enrollment, software updates, and usage of password managers.

The typology of S&P behaviors we explored in this study only capture triggers that are perceived *in-the-moment*. There may be other catalysts for behavior change that play a longer-term role in influencing end-user SP behaviors: e.g., social norms and cultural attitudes.

Finally, our categorization of individual behavioral triggers into “social”, “forced” and “proactive” is one of a number of other possible groupings. While our categorization was based on a synthesis of prior work and a thorough discussion amongst the authors of this paper, other groupings of the triggers may also be valid and could, through analysis, offer other insights into S&P behavioral triggers.

7 Conclusion

We conducted a large online survey to answer questions about what triggers good S&P behaviors, how that varies across individuals, and how people share their S&P behaviors with others. Social triggers were the most frequently reported behavioral triggers for pro-S&P behaviors, especially among those with low-to-medium security behavioral intention. We also found that participants were four times more likely to share their own S&P behaviors with others when their behaviors were also reported to be socially triggered. This result suggests the possibility of a feedback loop: if we can design behaviors that encourage social interaction, we may be able to trigger additional behavior change which, in turn, should encourage even more social interaction. People from different age groups and nationalities differed in which triggers they reported as prompting their S&P behaviors. Older people and people in the U.S. were more likely to respond to proactive triggers, while younger people and people in India were more likely to respond to social triggers. In summary, we contribute a general typology of in-the-moment, perceived S&P behavioral triggers and identify how those triggers vary across different individuals and behaviors. In turn, this contribution opens up fruitful new opportunities for the design of behavioral triggers meant to encourage pro-S&P behaviors.

Acknowledgment

This research was generously supported, in part, by the National Science Foundation through grants SaTC-1755625 and CNS-1704087. Tiffany Hyun-Jin Kim made important contributions to this work, particularly in helping design the survey. This work was also significantly improved through a dialogue with our anonymous reviewers, whose effort we appreciate.

References

- [1] Alessandro Acquisti and Jens Grossklags. Losses, Gains, and Hyperbolic Discounting: Privacy Attitudes and Privacy Behavior. In J. Camp and R. Lewis, editors, *The Economics of Information Security*, pages 179–186. 2004.
- [2] Anne Adams and Martina Angela Sasse. Users are not the enemy. *Communications of the ACM (CACM)*, 42(12):40–46, dec 1999.
- [3] Devdatta Akhawe and Adrienne Porter Felt. Alice in warningland: a large-scale field study of browser security warning effectiveness. In *Proc. USENIX Sec'13*, pages 257–272, 2013.
- [4] Monica Anderson and Kenneth Olmstead. Many smartphone users don't take steps to secure their devices. Technical report, Pew Research Center, 2017.
- [5] Cristian Bravo-Lillo, Lorrie F. Cranor, Julie Downs, Saranga Komanduri, Robert W. Reeder, Stuart Schechter, and Manya Sleeper. Your Attention Please: Designing security-decision UIs to make genuine risks harder to ignore. In *Proc. SOUPS'13*, 2013.
- [6] Lynne M Coventry, Debora Jeske, John M Blythe, James Turland, and Pam Briggs. Personality and social framing in privacy decision-making: A study on cookie acceptance. *Frontiers in psychology*, 7:1341, 2016.
- [7] Sauvik Das. Social cybersecurity: Understanding and leveraging social influence to increase security sensitivity. *it - Information Technology*, 58(5):237–245, jan 2016.
- [8] Sauvik Das, Eiji Hayashi, and Jason Hong. Exploring Capturable Everyday Memory for Autobiographical Authentication. In *Proc. UbiComp'13*, 2013.
- [9] Sauvik Das, Hyun Jin Kim, Laura A. Dabbish, and Jason I. Hong. The Effect of Social Influence on Security Sensitivity. In *Proceedings of the 10th Symposium on Usable Privacy and Security (SOUPS'14)*, 2014.
- [10] Sauvik Das, Adam D.I. Kramer, Laura A. Dabbish, and Jason I. Hong. Increasing Security Sensitivity With Social Proof: A Large-Scale Experimental Confirmation. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS '14)*, pages 739–749, New York, New York, USA, 2014. ACM Press.
- [11] Sauvik Das, Adam D.I. Kramer, Laura A. Dabbish, and Jason I. Hong. The Role of Social Influence in Security Feature Adoption. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*, pages 1416–1426, New York, New York, USA, 2015. ACM Press.
- [12] Sauvik Das, Joanne Lo, Laura Dabbish, and Jason I Hong. Breaking! A Typology of Security and Privacy News and How It's Shared. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–12, New York, New York, USA, 2018. ACM Press.
- [13] Rachna Dhamija, J D Tygar, and Marti Hearst. Why phishing works. In *Proc. CHI '06*, number April, pages 581–590, New York, New York, USA, 2006. ACM Press.
- [14] Paul DiGioia and Paul Dourish. Social navigation as a model for usable security. In *Proc. SOUPS '05*, pages 101–108, New York, New York, USA, 2005. ACM Press.
- [15] Paul Dourish, Rebecca E. Grinter, Jessica Delgado de la Flor, and Melissa Joseph. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing*, 8(6):391–401, sep 2004.
- [16] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proc. CHI '08*, page 1065, New York, New York, USA, 2008. ACM Press.
- [17] Serge Egelman, David Molnar, Nicolas Christin, Alessandro Acquisti, Cormac Herley, and Shriram Krishnamurthi. Please Continue to Hold: An empirical study on user tolerance of security delays. In *Proc. WEIS'10*, 2010.
- [18] Serge Egelman and Eyal Peer. Scaling the Security Wall: Developing a Security Behavior Intentions Scale (SeBIS). In *Proc. CHI'15*, pages 2873–2882, New York, New York, USA, 2015. ACM Press.
- [19] Adrienne Porter Felt, Alex Ainslie, Robert W Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettis, Helen Harris, and Jeff Grimes. Improving SSL Warnings. In *Proc. CHI'15*, pages 2893–2902, 2015.
- [20] BJ Fogg. A behavior model for persuasive design. In *Proceedings of the 4th International Conference on Persuasive Technology - Persuasive '09*, page 1.
- [21] SM Furnell, A Jusoh, and D Katsabas. The challenges of understanding and using security: A survey of end-users. *Computers & Security*, 25(1):27–35, 2006.
- [22] Shirley Gaw, Edward W Felten, and Patricia Fernandez-Kelly. Secrecy, Flagging, and Paranoia: Adoption Criteria in Encrypted E-Mail. In *Proceedings of the SIGCHI*

- conference on Human Factors in computing systems (CHI '06), pages 591–600, New York, New York, USA, 2006. ACM Press.
- [23] Joseph K. Goodman, Cynthia E. Cryder, and Amar Cheema. Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples. *Journal of Behavioral Decision Making*, 26(3):213–224, 2013.
- [24] Marian Harbach, Emanuel von Zezschwitz, Andreas Fichtner, Alexander De Luca, and Matthew Smith. It's a Hard Lock Life: A Field Study of Smartphone (Un)Locking Behavior and Risk Perception. In *SOUPS '14: Proceedings of the Tenth Symposium On Usable Privacy and Security*, pages 213–230, 2016.
- [25] Cormac Herley. So long, and no thanks for the externalities. In *Proc. NSPW '09*, pages 133–144, New York, New York, USA, 2009. ACM Press.
- [26] Cormac Herley. Unfalsifiability of security claims. *Proceedings of the National Academy of Sciences*, 113(23):6415–6420, 2016.
- [27] Torsten Hothorn, Frank Bretz, Peter Westfall, Richard M. Heiberger, Andre Schuetzenmeister, and Susan Scheibe. Simultaneous Inference in General Parametric Models, 2017.
- [28] Philip G. Inglesant and M Angela Sasse. The true cost of unusable password policies. In *Proc. CHI'10*, pages 383–392, New York, New York, USA, 2010. ACM Press.
- [29] Iulia Ion, Rob Reeder, and Sunny Consolvo. "...no one can hack my mind": Comparing Expert and Non-Expert Security Practices. In *Symposium on Usable Privacy and Security (SOUPS)*, pages 327–346, jan 2015.
- [30] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. Privacy Attitudes of Mechanical Turk Workers and the U.S. Public. In *Proc. SOUPS'14*, pages 37–49, 2014.
- [31] Heather Richter Lipford and Mary Ellen Zurko. Someone to watch over me. In *Proceedings of the 2012 workshop on New security paradigms - NSPW '12*, page 67, 2012.
- [32] Alexander De Luca, Sauvik Das, Martin Ortlieb, Iulia Ion, and Ben Laurie. Expert and Non-Expert Attitudes towards (Secure) Instant Messaging. In *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, 2016.
- [33] Grzegorz Milka. The Anatomy of Account Take-Over. In *USENIX ENIGMA*, 2018.
- [34] James Nicholson, Lynne Coventry, and Pam Briggs. Can we fight social engineering attacks by social means? assessing social salience as a means to improve phishing detection. In *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*, pages 285–298, 2017.
- [35] Donald A Norman and Stephen W Draper. *User centered system design: New perspectives on human-computer interaction*. CRC Press, 1986.
- [36] Kenneth Olmstead and Aaron Smith. Americans and Cybersecurity. Technical report, Pew Research Center, 2017.
- [37] Emilee Rader, Rick Wash, and Brandon Brooks. Stories as informal lessons about security. In *Proc. SOUPS '12*, New York, New York, USA, 2012. ACM Press.
- [38] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 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 - CCS'16*, pages 666–677, New York, New York, USA, 2016. ACM Press.
- [39] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples. In *IEEE Symposium on Security Privacy (S&P'19)*, page 0. IEEE, 2019.
- [40] Elissa M. Redmiles, Amelia R. Malone, and Michelle L. Mazurek. 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)*, pages 272–288. IEEE, may 2016.
- [41] M.A. Sasse. Computer security: Anatomy of a Usability Disaster, and a Plan for Recovery. In *Proc. CHI Workshop on HCI and Security Systems*. Citeseer, 2003.
- [42] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security ({SOUPS} 2015)*, pages 1–17, 2015.
- [43] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. The Post that Wasn't: Exploring Self-Censorship on Facebook. In *Proceedings of the 2013 conference on Computer supported cooperative work - CSCW '13*, pages 793–802, New York, New York, USA, 2013. ACM Press.
- [44] JM Stanton, P Mastrangelo, KR Stam, and Jeffrey Jolton. Behavioral Information Security: Two End User Survey Studies of Motivation and Security Practices. *AMCIS*, (August):2–8, 2004.

- [45] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorie Faith Cranor. "I regretted the minute I pressed share": A Qualitative Study of Regrets on Facebook. In *Proc. SOUPS 2011*, page 1, New York, New York, USA, 2011. ACM Press.
- [46] Rick Wash. Folk models of home computer security. In *Proc. SOUPS '10*, page 1, New York, New York, USA, 2010. ACM Press.
- [47] Alma Whitten and J.D. Tygar. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In *Proc. SSYM'99*, pages 14–28, 1999.

A Survey Questionnaire

Page 1: Shown to all participants.

1. What type of a cell phone do you have?
 - iPhone, Android, or other Smartphone
 - Non-smartphone cell phone
 - I don't know what kind of a cell phone I have
 - I don't own a cell phone

Page 2: Shown only to participants who selected 'iPhone, Android or other smartphone' in Page 1, Question 1 (P1Q1).

1. How do you mainly use your phone? Select all that apply.
 - Make phone calls
 - Check emails
 - Access social networking sites, such as Facebook, Twitter, Instagram, etc.
 - Access the Internet
 - Shopping, such as Amazon, Netflix, iTunes, etc.
 - Banking
 - Play games
 - Other: *Manual write-in*
2. Have you done any of the following in the past 6 months? Check all that apply.
 - Enabled or changed authentication on any of your mobile devices (e.g., 4-digit PIN, Android 9-dot, password, fingerprint, face recognition on your phone, laptop, tablet or other portable electronic device)
 - Updated your Facebook privacy settings
 - Uninstalled a smartphone app for privacy or security reasons

- Changed a password on an online account

Page 3: Shown to all participants

This study requires you to share your opinion. It is important that you take the time to read all instructions and questions carefully before you answer them. Previous research has found that some people do not take time to read everything that is displayed in the questionnaire. The questions below serve to test whether you actually take time to do so. If you read this, please answer 'two' on the question 4, add two to that number and use the result as the answer on question 5. Thank you for participating and taking time to read all instructions.

1. How many email addresses do you maintain?
2. How many social media accounts do you maintain?

Page 4: Shown to participants who fail P3 attention checks.

1. Is there someone to whom you truly want to talk about the recent change? Previous research has found that some people do not take time to read questions and answer options carefully. This question serves to test whether you actually take the time to do so. If you read this, please select 'colleague'. Thank you for taking time to read all instructions.
 - Friend
 - Family member
 - Significant other
 - Colleague
 - Other
 - None of the above

If this attention check is also failed, participants are disqualified and sent to the final page.

Page 5: Shown to participants who did not recall engaging in any of the behaviors listed in P2Q2.

1. Do you recall the most recent security or privacy behavior that you have changed on your mobile device or on the Internet? Please describe it briefly.
 - *Open response*

Page 6: Shown for each behavior participants selected in P2Q2 or P5Q1.

1. (Brief description of behavior being asked about). Did any of the following happen before you made the change? Please select all that apply.

- I directly experienced a security breach from a stranger
 - I directly experienced a security breach from someone I know
 - I allowed someone to use my device / account previously
 - (*Facebook privacy update only*) I noticed that my Facebook activities were visible to unintended people
 - (*App uninstallation only*) I noticed that the app required unusual permissions
 - I observed / heard about other people doing this
 - Someone I know advised me to do this
 - The device prompted me to do this prior to use
 - My organization required me to do this
 - I read a news article about the security vulnerability or recommending a best practice
 - I looked through settings/options for my mobile device
 - Other (required): *Manual write-in*
 - Nothing in particular happened before this change
2. (if selected option: 'I directly experienced a security breach from someone I know' in P6Q1) Who breached security on you? Please select all that apply.
- Friend
 - Family member
 - Significant other (spouse / boyfriend / girlfriend)
 - Colleague
 - Other (required): *Manual write-in*
 - I don't remember
3. (if selected option: 'I allowed someone to use my device / account previously' in P6Q1) Who used your device / account previously? Please select all that apply.
- Friend
 - Family member
 - Significant other (spouse / boyfriend / girlfriend)
 - Colleague
 - Other (required): *Manual write-in*
 - I don't remember
4. (if selected option: 'I allowed someone to use my device / account previously' in P6Q1) Who used your device / account previously? Please select all that apply.
- Friend

- Family member
 - Significant other (spouse / boyfriend / girlfriend)
 - Colleague
 - Other (required): *Manual write-in*
 - I don't remember
5. (if selected option: 'I observed people around me doing this' in P6Q1). You observed people around you doing this. Who did you observe? Please select all that apply.
- Friend
 - Family member
 - Significant other (spouse / boyfriend / girlfriend)
 - Colleague
 - Other (required): *Manual write-in*
 - I don't remember
6. (if selected option: 'Someone I know advised me to do this' in P6Q1). Who advised you to make this change? Please select all that apply.
- Friend
 - Family member
 - Significant other (spouse / boyfriend / girlfriend)
 - Colleague
 - Other (required): *Manual write-in*
 - I don't remember

Page 7: Shown for each behavior participants selected in P2Q2 or P5Q1.

1. After you made the change, did you talk about it to anyone else? Who did you talk with most recently?
- Friend
 - Family member
 - Significant other (spouse / boyfriend / girlfriend)
 - Colleague
 - Other (required): *Manual write-in*
 - I didn't talk about this with anyone.
2. (if selected option 'I didn't talk about this with anyone' in P7Q1) Why did you decide not to talk about this to anyone? Please select all that apply.
- I didn't feel comfortable to talk about security
 - I assumed that people already knew about this
 - I assumed that people didn't need to know about this
 - I just didn't want to talk about this to anyone

- I didn't have a chance to talk about this to anyone yet
 - Other (required): *Manual write-in*
3. (if selection any option *except* 'I didn't talk about this with anyone' in P7Q1) What channel did you use to talk about the change most recently?
- Face to face conversation
 - Phone call
 - Text message or email
 - Facebook
 - Twitter
 - Other (required): *Manual write-in*
4. (if selection any option *except* 'I didn't talk about this with anyone' in P7Q1) What prompted you to talk about the change with them? Please select all that apply.
- I noticed they were being insecure
 - They noticed my change
 - They learned about a new security tool
 - I felt obligated to protect them
 - They experienced a security or privacy breach
 - They had to set up a new device, account, or security tool
 - They read a news article about security
 - I just wanted to talk about my recent change
 - Other: *Manual write-in*
 - None of the above
5. (if selection any option *except* 'I didn't talk about this with anyone' in P7Q1) What did you talk about in your conversation? Please select all that apply.
- I shared a notification or warning of a potential security or privacy threat
 - I demonstrated insecure behavior
 - I shared instructions on how to change insecure behavior
 - I shared specific advice
 - I shared a story about an experience I had
 - I shared my emotional venting
 - I just talked about a security event
 - Other: *Manual write-in*
 - I just talked about the change I made

1. How would you evaluate your computer literacy level?
- Very low: I don't know much about computers (1)
 - Low (2)
 - Neither high nor low (3)
 - High (4)
 - Very high: I know a lot about computers (5)
2. How would you evaluate your Internet literacy level?
- Very low: I don't know much about how the Internet works (1)
 - Low (2)
 - Neither high nor low (3)
 - High (4)
 - Very high: I know a lot about how the Internet works (5)
3. How many hours per week are you on the Internet for reasons other than work (both using the smartphone, tablets, or computers)?
- 0 to 10 hours
 - 10 to 20 hours
 - 20 to 30 hours
 - 30 to 40 hours
 - More than 40 hours
4. How many different online communities (e.g., reddit), social networks (e.g., Facebook), or online groups (e.g., email list) do you read or post in regularly?
- None
 - 1
 - 2 to 4
 - 5 or more
5. How many hours per day do you spend on sharing and reading content on social networking sites (e.g., Facebook, Twitter, Google+, Instagram, etc.)?
- 0 to 1 hour
 - 1 to 3 hours
 - 3 to 6 hours
 - 6 to 9 hours
 - More than 9 hours
6. Please rate your familiarity with the following concepts or tools on the following scale:
- I never heard about this
 - I heard about this but I don't know what it is

- I know what this is but I don't know how it works
- I know generally how it works
- I know very well how this works

- IP address
- Cookie
- Secure Socket Layer (SSL) / Transport Layer Security (TLS)
- Virtual Private Network (VPN)
- Encryption
- Proxy server
- Tor
- Privacy settings for your web browser
- Private browsing mode in browsers

7. Please indicate whether you think each statement is true or false. Please select "I'm not sure" if you don't know the answer.

- Private browsing mode in browsers prevents websites from collecting information about you.
- Login cookies can store username/id and a random string in your web browser to keep the user signed in.
- No one, except for the sender and intended receiver, can reveal the content of an encrypted message.
- Tor can be used to hide the source of a network request from the destination.
- A Virtual Private Network (VPN) is the same as a Proxy server.
- IP addresses can always uniquely identify your computer.
- HTTPS is standard HTTP with SSL / TLS to preserve the confidentiality of network traffic.
- A proxy server cannot be tracked to the original source.

Page 9: Shown to all participants who passed the attention checks. SEBIS questions from Egelman and Peer [18].

1. Please indicate how often you have done the following on the following scale:

- Never
- Rarely
- Sometimes
- Often
- Always

- I set my computer screen to automatically lock if I don't use it for a prolonged period of time.
- I use a password/passcode to unlock my laptop or tablet.
- I manually lock my computer screen when I step away from it.
- I use a PIN or passcode to unlock my mobile phone.
- I do not change my passwords, unless I have to.
- I use different passwords for different accounts that I have.
- When I create a new online account, I try to use a password that goes beyond the site's minimum requirements.
- I do not include special characters in my password if it's not required.
- When someone sends me a link, I open it without first verifying where it goes.
- I know what website I'm visiting based on its look and feel, rather than by looking at the URL bar.
- I submit information to websites without first verifying that it will be sent securely (e.g., SSL, "https://", a lock icon).
- When browsing websites, I mouseover links to see where they go, before clicking them.
- If I discover a security problem, I continue what I was doing because I assume someone else will fix it.
- When I'm prompted about a software update, I install it right away.
- I try to make sure that the programs I use are up-to-date.
- I verify that my anti-virus software has been regularly updating itself.

Page 10: Shown to all participants who passed the attention checks.

1. While using the Internet, have you ever done any of the following? Please check all that apply.

- I have used a temporary username or email address.
- I have used a fake name or username.
- I have given inaccurate or misleading information about myself.
- I have set my browser to disable or turn off cookies.
- I have cleared cookies and browser history.

- I have used a service that helped me browse the web anonymously, such as a proxy server, Tor, or a virtual personal network (VPN).
 - I have sent encrypted e-mails.
 - I have decided not to use a website because they asked for my real name.
 - I have deleted something I posted in the past.
 - I have asked someone to remove something that was posted about me online.
 - I have used a public computer to browse anonymously.
2. If we ask you to perform the following actions now, can you do it without getting help from others? Please answer on the following scale.
- Yes I can do this without getting help from others
 - Probably but I may need help from time to time
 - No I need help from others to do this
- (a) Change authentication on mobile devices
- (b) Change Facebook privacy settings
- (c) Change passwords of your online account
- (d) Check permission requests when downloading an app on mobile devices
3. Have you ever done any of the following? Please select all that apply.
- I have turned off the automatic connections to free Wi-Fi on my mobile device(s)
 - I have looked for "https" when browsing or shopping on my mobile devices
 - I have turned on login approvals on my Facebook account
 - I have enabled secure browsing on my Facebook account
 - I have kept the same password for an online account after logging in using a public computer
 - I have clicked a URL link on an email and entered my username and password

Page 11: Shown to all participants who passed the attention checks.

1. What is your gender?*

 - Male
 - Female
 - Non-conforming
 - Prefer not to answer

2. What is your age?
3. What is your current relationship status?
4. Are you a parent or guardian of any children under 18 years of age?
5. How many adults (age 18 or older) currently live in your household, including yourself? *Optional manual write-in*
6. What is the highest level of school you have completed or the highest degree you have received?
7. Out of the following, which best describes your major (if you are a student) or occupation (if you are a professional)?*

 - Cybersecurity related
 - Computer Science related
 - Other engineering or technology related
 - Other: *Manual write-in*

8. What nationality do you most identify with?
9. Do you have native or fluent proficiency in English?
10. Are you Hispanic or Latino?
11. What is your race? Please select all that apply.

 - American Indian or Alaska Native
 - Asian
 - Black or African American
 - Native Hawaiian or other Pacific Islander
 - White
 - Prefer not to answer