The Effect of Social Influence on Security Sensitivity

Sauvik Das, Tiffany Hyun-Jin Kim, Laura A. Dabbish, Jason I. Hong
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA, 15213
sauvik@cmu.edu, hyunjin@cmu.edu, dabbish@cmu.edu, jasonh@cs.cmu.edu

ABSTRACT

Despite an impressive effort at raising the general populace's *security sensitivity*—the awareness of, motivation to use, and knowledge of how to use security and privacy tools—much security advice is ignored and many security tools remain underutilized. Part of the problem may be that we do not yet understand the social processes underlying people's decisions to (1) disseminate information about security and privacy and (2) actually modify their security behaviors (e.g., adopt a new security tool or practice). To that end, we report on a retrospective interview study examining the role of social influence—or, our ability to affect the behaviors and perceptions of others with our own words and actions—in people's decisions to change their security behaviors, as well as the nature of and reasons for their discussions about security. We found that social processes played a major role in a large number of privacy and security-related behavior changes reported by our sample, probably because these processes were effective at raising security sensitivity. We also found that conversations about security were most often driven by the desire to warn or protect others from immediate novel threats observed or experienced, or to gather information about solving an experienced problem. Furthermore, the *observability* of security feature usage was a key enabler of socially triggered behavior change—both in encouraging the spread of positive behaviors and in discouraging negative behaviors.

1. INTRODUCTION

There are many reasons why security advice is often ignored and many security tools are left unutilized [17]. Some prior work suggests that many believe they are in no danger of experiencing a security breach [1] and are *unaware* of both threats and the security tools available to protect against those threats. Other work suggests that many choose not to use security tools and follow security advice because doing so is often antagonistic towards the immediate goal of end users—a complex password that usually requires three attempts to get right *prevents* a user from doing what she actually wants to do: e.g., authenticating into social media. Herley further argues it may even be economically *rational* for users to ignore security advice, as the expected cost, in time, of a lifetime of following security advice might actually be higher than the expected loss a user would suffer if his account actually was compromised [17]. Thus, many people lack the *motivation* to behave securely. Still others suggest that security tools are simply too difficult to use [26,34], so many people do not have the *knowledge* required to operate them. Taken together, it appears that the lack of what we call *security sensitivity*—the awareness of, motivation to use, and knowledge of how to use security and privacy tools—is a large barrier to increasing the uptake of security tools and the following of security advice.

Symposium on Usable Privacy and Security (SOUPS) 2014, July 9-11, 2014, Menlo Park, CA.

Prior work has looked at improving all parts of the security sensitivity stack—for example, through games for security education [28], browser extensions to make users more aware of phish [35], more effective user interfaces for security tools [19], and faster or simpler ways to authenticate users [31]. Security sensitivity, nevertheless, remains low.

We argue that part of the problem is that we do not yet understand the *social* processes underlying people's decisions to communicate about security and adopt security tools. In other words, security behaviors—as any human behavior—should be viewed within the context of a social system. Indeed, social psychology and sociology literature illustrates that social influence, or our ability to affect other people's perceptions and behaviors with our words and actions [6], plays a central role in how people behave—even specifically in changing their behavior or adopting a new technology or idea [6,25]. Rogers' highly influential *diffusion of innovations* work, for example, has shown that social influence drives technology adoption [25]. Social processes, thus, should undoubtedly affect a user's decision to follow security advice or adopt a security tool.

Nevertheless, the effect of social influence on decisions and communications about security and privacy remains understudied. Indeed, we do not yet know *how* social influence affects behavior change with regards to security and privacy. Likewise, we know little about the nature of conversations about security and privacy, through which this influence should occur. Understanding how social influence affects security related behavior change and communication could improve our understanding of why security sensitivity remains low, and could help inform the design of social interventions that can raise security sensitivity. To that end, we report on a retrospective interview study aimed at investigating the following research questions:

Q1: What role does social influence play in an individual's decisions to use, discontinue use, and explore security tools and privacy settings?

Q2: Under what circumstances do people communicate about security and privacy?

In our interviews, we probed participants about their experiences with regards to mobile phone authentication, mobile application installation and uninstallation, and social media privacy settings. We also asked participants to recall specific conversations they had about cybersecurity and online privacy.

Our findings suggest social processes played a major role in a large number of privacy and security related behavior changes reported by our interviewees, probably because these processes were effective at raising all points of the security sensitivity

stack—awareness, motivation and knowledge. However, different triggers for socially driven behavior change varied in the extent to which they raised awareness, motivation and knowledge about security tools and behaviors. In addition, conversations our participants had about security and privacy were most often instigated by the desire to (1) warn or protect others from immediate or novel threats observed or experienced, and (2) to gather information about solving an immediate problem. One particularly salient theme that arose from our interviews is that the *observability* of security feature usage was a key enabler of socially triggered behavior change and conversation—in encouraging the spread of positive behaviors, discouraging negative behaviors, and getting participants to talk about security. Taken together, our results suggest that: (1) there is a substantial and often overlooked social process that strongly affects security-related behavior change, and (2) in order to maximally raise security sensitivity, security and privacy tool usage should be more observable and amenable to conversation.

2. RELATED WORK

2.1 Security Sensitivity

Prior work in usable privacy and security alludes to three reasons underlying why much security advice is ignored and many security features remain unused: lack of awareness, motivation, and knowledge. First, many users lack the *awareness* of security threats and the tools available to protect themselves against those threats. For example, a study by Adams and Sasse found that insufficient awareness of security issues caused users to construct their own model of security threats that are often incorrect, resulting in security breaches [2]. Stanton et al. found that a lack of awareness of basic security principles even influenced experts to make naïve security mistakes, such as using a social security number as a password [30]. Users who are unaware of a threat cannot take measures to avoid the threat, and users who are not cognizant of the tools available to protect themselves from these threats cannot use those tools to actively defend themselves.

Second, users—even those who are aware of security and privacy threats and the preventive tools that combat those threats—often lack the *motivation* to utilize security features to protect themselves [2,12]. The lack of motivation to use security features is not entirely surprising, as stringent security measures are often antagonistic towards the specific goal of the end user at any given moment [10,26]. For example, while a user might want to access her Facebook, a complex password that usually requires three attempts to get right *prevents* her from accessing Facebook for an intolerable amount of time [11]. Thus, users often reject the use of security and privacy tools when they expect or experience them to be weighty [2,14,18,26]—and security features are often weighty.

Furthermore, many security threats remain only an abstract threat to most individuals [2,16,24]: Bob may know, conceptually, that there are security risks to using the same simple password across accounts, but does not believe that he is, himself, in danger of experiencing a security breach. Additionally, this perspective may be economically rational, as the expected cost, in time, of a lifetime of following security advice might actually be higher than the expected loss a user would suffer if his account actually was compromised [17]. Finally, the benefits of security features are often invisible, as users are often not cognizant of the *absence* of a breach that otherwise would have occurred without the use of a security or privacy tool. In all, it is unsurprising that many users lack the motivation to explicitly use security tools: to do so would mean to incur a frustrating complication to everyday interactions in order to prevent an unlikely threat with little way to know whether the security tool was actually effective.

Third, security tools are often too complex to operate for even aware and motivated end-users, suggesting that users lack the *knowledge* to actually utilize security tools [34]. Indeed, there is a wide gulf of execution for most security features for most users. For example, many users cannot distinguish legitimate vs. fraudulent URLs, nor forged vs. legitimate email headers [8]. Also, a study revealed how security features in Windows XP, Internet Explorer, Outlook Express, and Word applications are difficult for users [13]. And, Wash found that many people hold "folk" models of computer security that are often misguided, and use these incorrect models to justify ignoring security advice [32].

In sum, prior work in usable privacy and security suggests that there are at least three large obstacles inhibiting the widespread use of security and privacy tools: the *awareness* of security threats and tools, the *motivation* to use security tools, and the *knowledge* of how to use security tools. We refer to this layered stack as *security sensitivity* for ease of discussion, as it encapsulates how likely a user is to seek information about and use security tools. Note, however, that the concept is not necessarily novel, as prior work has alluded to such a stack in security specifically [12], and in the adoption of technology more generally [7,25].

2.2 Social Influence and Security Sensitivity

In his seminal work on the diffusion of innovations, Rogers claimed that new technology gets widely adopted through a process by which it is communicated through members of a social network [25]. Rogers argues that primarily *subjective perceptions*, not scientific or empirical fact, get communicated through social channels, and that these perceptions are key to the success of an innovation. He further outlines that preventative innovations—or innovations, like security and privacy tools, that prevent undesirable outcomes from happening in the future—typically have low adoption rates, probably because of their lack of *observability*, or the invisibility of their use and benefits.

Other work in cognitive psychology has looked at the psychological mechanisms underlying social influence. For example, lots of prior work has demonstrated the potency of the concept of "social proof"—or our tendency to look to others for examples of how to act in uncertain circumstances [5,6]. For example, Milgram, Bickman, and Berkowitz [22] demonstrated the social proof principle when they showed that simply getting a small crowd of people—the more, the better—to look up at the sky on a busy sidewalk caused others to do the same.

Still other work has shown how social interventions can be powerfully effective at driving human behavior: for example, at reducing household energy consumption by showing people their neighbors' reduced energy consumption [27], reducing hotel guests' wasteful use of towels by showing them previous patrons chose to be less wasteful [15], and even in eliminating young children's phobia of dogs by showing them film clips of other children playing with dogs [3].

Taken together, the background literature suggests that social influence strongly affects people's behaviors and decisions; likely, also their security-related behaviors and decisions. And, indeed, prior work *has* alluded to the importance of social processes in raising security sensitivity. For example, DiGioia and Dourish [9] suggested that "social navigation"—or people's inclination to

	Age	Gender	Race	Occupation
P1	28	Male	Black	Customer Service
P2	22	Female	Asian	Unemployed
P3	22	Female	Black	Student
P4	22	Male	Black	Student
P5	27	Female	Asian	Unemployed
P6	29	Male	White	Programmer
P7	54	Female	White	Admin. Assistant
P8	31	Male	Indian	Unemployed
P9	30	Male	White	Software Developer
P10	37	Male	White	Graphic Designer
P11	54	Male	Black	Chef
P12	20	Female	Black	Student
P13	24	Female	Indian	Graduate Student
P14	25	Male	Indian	Graduate Student
P15	21	Male	Indian	Graduate Student
P16	22	Male	Indian	Graduate Student
P17	34	Female	Asian	Unemployed
P18	20	Male	Black	Student
P19	20	Male	White	Student

Table 1. Participant demographics.

look for cues on how to act—can be used to raise users' security sensitivity by showing them other users' actions in context. Rader et al.'s study on stories as informal lessons about security suggests that storytelling increases awareness of and motivation to guard against security threats [23]. In addition, Singh et al. outlined the common practice of sharing passwords and PINs [29]. On the other hand, Gaw et al. [14] found that many people believed that use of security features was an indication of paranoia, unless the user had an obvious reason for doing so. If there *is* a stigma of paranoia attached to using security features, then it is possible that, under some circumstances, social influence can work *against* security sensitivity (e.g., "only paranoid people encrypt their e-mail, and I'm not paranoid").

Nevertheless, the background literature on the social dimensions of security and privacy remains surprisingly thin. To our knowledge, little work has specifically looked at how social influence affects security sensitivity, and, in turn, enacts behavior changes related to privacy and security, or how people generally communicate about security and privacy (outside of Rader et al.'s study on security storytelling [23]). Yet, understanding how social influence affects security related behavior change and communication could improve our understanding of why security sensitivity remains as low as it is, and could even help inform the design of social interventions that raise security sensitivity. To that end, we look specifically at the social dimensions of security related behavior changes and communications in this work.

3. Methodology

3.1 Semi-Structured Interview Methodology

We constructed an IRB approved semi-structured interview protocol to probe participants about recent security related behavior changes and conversations. We elected a semi-structured approach so that we could concretize the discussion by directing participants' memories towards *changes* in behavior or *specific* instances of communication, while still allowing participants the flexibility to expand on the their undoubtedly unique experiences. Our interview protocol probed participants about recent changes in (1) *mobile authentication*, or whether and why participants enabled, disabled, or changed authentication on their smartphones (e.g., from PIN to Password); (2) *application installation and uninstallation*, or whether and why participants decided to uninstall or halt installing applications because of privacy and security concerns; and, (3) *online privacy settings in social media*, or whether and why participants changed their privacy settings on the social media platform they most commonly used. We chose to explore three categories to uncover general trends across different types of security tools, and we chose these three categories specifically because they represented a broad range of behaviors representative of common security and privacy decisions made by just about all people fairly regularly.

If participants reported a *specific* security-related behavior change, we asked them to explain further how the change was catalyzed—specifically, to discern between *social* and *non-social* catalysts for behavior change. Either way, we asked participants to explain, in detail, the context surrounding their decision to enact the change: Was the change brought about by a personal negative experience, or because of an article they read online? If they heard about a security incident through a friend, how did the friend broach the conversation? And, if a social process drove the change, we asked participants to clarify how the social process manifested—for example, did they seek out advice, or did a friend offer them unsolicited advice? We also asked participants whether and why they did or did not share their concerns, advice, or behavior change with anyone else.

We also asked participants if they could recall *specific* conversations they had about security and privacy. Did they ever share information about security or privacy? If so, what did they share, with whom, and why? By focusing on specific conversations about security and privacy (e.g., "I told my mother to update her privacy settings"), rather than general conversations (e.g., "People usually tell me to update my password"), we were often able to uncover the specific context of a conversation (e.g., a catalyst and goal for the conversation).

To capture security-related conversations that did not fit into the pre-constructed themes of mobile authentication, app installation, and social media privacy settings, at the end of the interview, we also asked participants more open-ended questions about conversations related to security and privacy.

We iteratively refined our protocol by piloting it with 5 people. All interviewers participated in the pilots in order to mitigate variation in delivery across interviewers and interview sessions. Questions that participants could not easily answer (e.g., hypotheticals) were culled through these iterations. Ultimately, our interview lasted approximately 45 minutes, and interviewees were compensated $10 to participate.

3.2 Recruitment

We recruited participants from an online recruitment tool that pairs research participants from the local area with research projects of interest. Participants were required to own a smartphone running Android or iOS, be an active user of any

social media service, and be at least 18 years old. We went through three rounds of recruitment to recruit a variety of occupations and ages across our sample. For example, in our first round of recruitment, we predominantly interviewed students in their mid-twenties. Thus, in subsequent recruitment rounds, we specifically recruited older non-students. We stopped recruiting additional participants once we believed we had sufficient diversity in occupation, age, and security proficiency to capture a large cross-section of experiences with security-related behavior change and communication. In our case, we appeared to reach this point after interviewing 19 participants—indeed, after the first 15, every additional participant echoed experiences very similar to those previously reported by others. Our recruitment solicitation is attached in Appendix B.

Our participants ranged in age from 20 to 54 years old (m=28.5, sd=10). Seven out of the 19 participants were female. Furthermore, as we tried to recruit participants from diverse backgrounds, 10 of our participants were non-students from many different professional backgrounds. All participants used an Android (n=12) or iOS (n=7) smartphone and were frequent Facebook users. Fifteen of the 19 participants reported using Facebook daily, while the remaining 4 reported that they checked Facebook at least a few times every week. Table 1 summarizes participant demographics. A more detailed description is in Table A1 of Appendix A.

3.3 Data Coding and Analysis

We recorded and transcribed, with consent, each interview, and used a qualitative data analysis program called Dedoose [37] to analyze the anonymized transcripts. We first partitioned each transcript into two sets of "excerpts". The first set of excerpts was a collection of all instances of an *action taken*, a *decision made*, or, more generally, a *behavior changed* related to security or privacy. As such, we will refer to this set of excerpts as the *behavior changes*. A representative example of behavior changes is P18's decision to rub-off the smudges on his Android device after a friend demonstrated that the smudges on his screen makes it easy for others to "crack" his Android 9-dot pattern:

"What I've been doing, I believe, after that scare with the nine dot, pretty much every time I turn off my phone, I put it in the pocket, I just kind of rub, just rub the smears off so you can't really see what direction I was going." (P18)

The second set of excerpts was a collection of all *specific* instances of communication about security and privacy, which we will refer to as the *communications*. An example excerpt comes from P14. After he received spam mail from a friend's e-mail account, he mentioned:

"I told my friend that this is something weird that came from your account. This is not what you would be probably into." (P14)

In total, from our 19 transcripts, we extracted n=114 behavior change excerpts, and m=118 communication excerpts. Excerpts were usually just answers to pointed questions, but to ensure robustness, two of the research group mutually agreed on all partition points for each excerpt.

We used these excerpts as our units of analysis—though, occasionally, we aggregated data across participants where it made sense (e.g., in determining how many participants actually changed their behavior as a result of a social process). We used an iterative, open coding process [21] to code the data, constructing codes where patterns naturally emerged and refining the codes iteratively until we reached consensus. Ultimately, we had two goals in mind through the coding process. The first was to understand the effect of social influence in driving *behavior changes*—which, in turn, means understanding the effect of social influence in modulating *security sensitivity*; and, the second was to better understand the triggers and reasons underlying *communications* about security and privacy.

Concretely, two researchers independently and openly coded a random subset of 20% of the excerpts from each of the *behavior changes* and *communications* excerpts. These openly generated codes were collaboratively synthesized into a set of high-level codes that three of the research team then used to code the remaining excerpts. Upon completion, the coding team discussed potential extensions to the coding scheme that arose from coding the new examples. If a change to the scheme was made, the coding team re-coded the full set of excerpts with the new scheme. We required two coding iterations to come to consensus.

From the 20% overlap of excerpts, overall inter-coder agreement was 85% for *behavior changes*, and 79% for *communications* (calculated as the number of overlapping excerpts where codes matched divided by the total number of overlapping excerpts). In cases of discrepancies, the coders discussed the discrepancies until agreement was reached, following standard practice. Inter-coder agreement for *each* applied code can be found in Table A3 in Appendix A, and all exceeded the 0.7 threshold commonly held to be acceptable in qualitative research [21].

4. RESULTS

4.1 Behavior Changes

First, we wanted to know if social processes often drove security related behavior changes, so we coded each *behavior change* excerpt as being driven by a *social* or *non-social* process. Excerpts were coded as being driven by a social process when the reason for the behavior change was social, and, importantly, if the social process was *clearly* reported by the participant in the transcript. For example, when asked about why he first enabled a PIN on his iPhone, P6 stated:

"When I first had a smartphone I didn't have a code, but then I started using one because everyone around me I guess had a code so I kind of felt a group pressure to also use a code." (P6)

As the underlying reason for the behavior change was a social process (observing one's friends) and was stated as such, we coded that behavior change as social. An example of a non-social behavior change comes, again, from P6. When asked why he changed his Twitter password, P6 responded:

"Diversification of passwords. I had the same password for every service so I wanted to pick a stronger password for… the service, yeah." (P6)

While P6 *could* have learned about the need for password diversification from friends, as he did not explicitly confirm this speculation, we coded the excerpt as non-social.

In all, out of the 114 behavior change excerpts, we coded a substantial 48 as being explicitly driven by some form of social influence. Furthermore, most participants (17 out of 19) reported at least one action taken, decision made, or behavior changed that was driven by social influence. Of note, however, is that the 48 examples of socially driven behavior change did not come uniformly from all of our participants. Notably P2 and P10 reported the largest number of socially driven changes at eight,

Trigger	N	Description	Example
Observed friends	14	Observing people around them engaging in a particular security behavior and emulated those people.	*"So when I was an undergrad I've been using it since then. And this four digit everybody started using it and it was a hype. And we had it." (P14)*
Social sensemaking	9	Discussing concerns with friends/loved ones to determine the right behavior.	*"I mean, like, one of my friends told me that you could alter the privacy settings so that, like, not everyone can look up your profile and not everyone can, like, try sending messages to you." (P15)*
Prank/ Demonstration	8	Friends/loved ones hacked into his/her account, demonstrating they were insecure.	*"Yeah, like my laptop was in my room. I walked out of my room and someone walked by and saw my Facebook and thought it would be funny to put something up." (P19)*
Security breach	6	Someone hacked into his/her account or information was shared too widely.	*"I did change that within the past week. The girlfriend was reading all of my mail, which is also a privacy concern" (P10)*
Sharing access	3	Sharing access to a device or account with another person leading to need for better security.	*"There are sometimes when you have to tell your friends what is my PIN number because they are a very good friend of yours and they have to make a call and I can't go every time and just unlock this for them." (P14)*

Table 2. Social triggers for behavior change derived from our iterative open coding process.

each. It is important to keep this bias in mind in any quantitative interpretation of our findings.

In all, these results suggest that social influence already plays a strong role in driving security and privacy related behavior change—even without any explicit social interventions. Next, we wanted to understand when and how social influence is effective at driving these behavior changes.

4.1.1 Social Triggers in Driving Behavior Change

To explore *when* social influence drove behavior change, we open coded the *triggers* for behavior change excerpts coded as "social". We found five primary social triggers for behavior change: *observing friends*, *social sensemaking*, *pranks and demonstrations*, *experiencing security breaches*, and *sharing access*. Table 2 lists all triggers, their frequency and their description. Next, to answer *how* social processes enacted behavior change, we also coded whether or not the socially driven behavior change examples in our dataset affected any part of the security sensitivity stack. Specifically, we asked the following:

Raised Awareness: Did the social process raise the participant's awareness of a new threat and/or security tool?

Raised Motivation: Did the social process raise the participant's motivation to protect him or herself against a security threat?

Raised Knowledge: Did the social process raise the participant's knowledge of how to use a security tool or method?

Importantly, we only answered "yes" to those questions if the *social* process mentioned in the excerpt was the reason for the heightened security sensitivity. For example, P16 mentioned that his Facebook account getting "hacked" resulted in him changing many of his passwords every 6 months at the advice of his friends, who he sought out for advice after the incident. In this example, the social process of P16 speaking with his friends raised his *knowledge* but not his *awareness* or *motivation*. It was the non-social process of experiencing a security breach that raised both his awareness and motivation.

For most (44 of 48) reported examples of socially driven behavior change, we found that the social process triggering the behavior change did, in fact, raise *some* form of security sensitivity. In fact, many examples raised *all* points of the security sensitivity stack. For example, P18 recalled advice he received on password composition after asking his friend to share a password:

"When I was working this summer, one of my co-workers told me about the whole algorithm thing. One, it just helps you I guess have different passwords. It helps you recall them easier based on I guess the type of profile. I guess you can cater, you can change your algorithm, depending on I guess what you want to be in it. But ever since I started using it." (P18)

In this example, the social process of P18 asking his friend about how to compose a password increased his *awareness* of a new method of password composition, his *motivation* to update his own method of password composition, and his *knowledge* of how to improve his method of password composition.

In the text to follow, we describe each social trigger we found in our data for security related behavior change. Furthermore, as a descriptive aid, we plotted how frequently different social triggers raised the different components of security sensitivity in Figure 1.

Observing friends (14/48 examples)

Most frequently, our participants reported changing their behavior after observing the actions of friends or others around them. In other words, participants changed their behavior after finding *social proof*—or, cues on how to act based on the actions of others [6]. For example, one participant in our sample adopted the 9-dot authentication method on his Android phone because his friends also used it. Additionally, as previously illustrated, P6 adopted a PIN because he felt "group pressure" to do so after observing everyone around him use authentication. This finding appears to be well supported by the background literature on technology adoption, which lists *observability* as a key criteria for an innovation to spread rapidly through social channels [25].

In certain cases, other forms of social influence apart from social proof appeared to be at play—specifically the social influence concepts of *liking*, or our tendency to follow the advice of those we like and those like us, and *authority*, or our tendency to follow the advice of those we consider to be authority figures [6]. For example, one participant indicated that she adopted a PIN code for her iPhone wholly because her mother, who she considered technically savvy, also had a PIN:

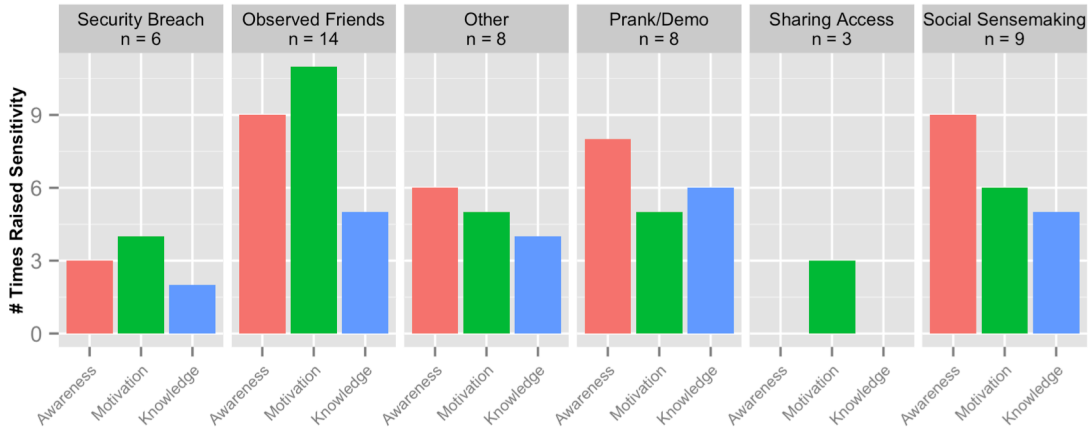

Figure 1. The number of times each social trigger for behavior change reported by our sample raised any of the three parts of the security sensitivity stack: awareness, motivation, or knowledge.

"My mother has-- she had an iPhone before I did, so she always had the block on hers, so I just kind of the... I think just because I saw her doing it, so it kind of just felt like it was something I had to do too." (P3)

Observation influenced behavior change for mobile authentication more often than the other specific topics we asked about in our interviews, probably because it is relatively easy to observe others authenticating onto their phones compared to observing others update their social media privacy settings or uninstall an app.

Looking at Figure 1, participants who observed others use security tools often were themselves *motivated* to start using those tools (11/14 examples). Furthermore, participants often became more *aware* of security tools after observing others' using those tools (9/14), but only occasionally gained *knowledge* of how to use the observed tools and methods by observing others (5/14).

Social Sensemaking (9/48 examples)

The second most frequent social trigger reported by our sample was *social sensemaking*—or, the process of making sense of a security system, tool, or threat by discussing concerns with others. We termed these triggers social sensemaking because they were similar in form and purpose to discussions, observed by Weick et al., among members of an organization who attempted to resolve uncertainty about recent novel events in their environment [33].

Participants often reported having discussions to resolve ambiguity in news and hearsay about security. The aim of these discussions was usually to find the correct or appropriate way to act to achieve the desired level of privacy or security within a system or with a security tool. In many cases, these discussions were prompted by a sudden infusion of uncertainty—for example, news articles about a novel security threat or gossip about anomalous security breaches others had experienced. Participants discussed these novel threats with others to share information about the threat, assess its veracity, and determine whether and how to change their behavior in response. For example, one participant in our dataset reported becoming more restrictive with posting to Facebook in response to a sudden, alarming, but unclear threat of all timeline posts becoming public:

"So yeah. I recently, like, a day or two, day before yesterday, I went through an ordeal. I don't know if it's fake or it's real, but somebody mentioned that all his private messages, they became public. Like, his messages with a friend. And it was like he had never thought of putting it on wall. And it suddenly opened his Facebook and everything was on his…I don't know if it's a real thing. And somebody mentioned in a comment that it happened with him as well, few days back." (P16)

P16's example is another illustration of *social proof* based social influence affecting an individual's security behavior: facing an ambiguous threat, P16 observed his friends for cues on how to act.

Social sensemaking also occurred when a participant wanted to understand a particular function within a system—for example, Facebook privacy settings. This need for specific information resulted in discussion and information sharing that exposed novel functionality or methods for protecting oneself against threats—often increasing participants' *knowledge* about the system (5/9 examples) and eventually leading to behavior change as a result. For example, one participant updated his privacy settings after a discussion that revealed novel system functionality:

"I mean, like, one of my friends told me that you could alter the privacy settings so that, like, not everyone can look up your profile and not everyone can, like, try sending messages to you. As in you can go to the privacy settings tab. And then, you could actually change it. Because I didn't know that you could do it, before. I mean, I just thought that it was default that everyone could look at your profile." (P15)

Social sensemaking also made participants more *aware* of available security tools (9/9), and the discussions would frequently *motivate* participants to act on their newly acquired knowledge (6/9).

Prank/Demonstration (8/48 examples)

The third most prevalent social trigger for the behavior changes reported by our participants was pranks and demonstrations—i.e., friends or loved ones cracking participant's accounts and devices as a prank, or to demonstrate that they were being insecure. Often, these pranks were explicit demonstrations to prove to the victim that their current security strategy or behavior was insecure. For example, one participant in our sample described a co-worker breaking into his phone to show the vulnerabilities of 9-dot authentication:

"One of my, when I was interning, engineering company, one of my friends and a fellow intern came to my desk, just unlocked my phone. I was surprised. I was like, "Hey, how'd you do it?" He put it against the sunlight and he saw I guess the smudges my finger left. He just followed the direction. Yeah, he had access to my phone." (P18)

Other prank examples reported were simply driven by opportunity—for example, a friend gaining unauthorized access to the participant's account because they left their Facebook account open on an unprotected device. Indeed, several of our participants were motivated to change their security behavior after their friends accessed their social media accounts and posted embarrassing information on their behalf. For example, one participant experienced this type of prank after leaving his laptop open and unprotected in his dorm room:

"Besides just my friends getting into my phone or on my Facebook and that's more from just me leaving my Facebook open or something if I walk out of the room and they just put up a funny status or something like or even just look through my messages or something like that. But nothing too threatening, more like practical joking side of it. But once that happens, I usually change my password immediately as would all of my other friends would too." (P19)

Pranks appeared to be quite effective at raising participants' security sensitivity. In all cases (8/8 examples), participants were made aware of a security threat and, in most cases, participants were instantly motivated (6/8) to update their behaviors to prevent a reoccurrence of the prank. Pranks aimed at demonstrating insecure behavior were also effective at raising participants' knowledge (5/8), as they were often followed up with direct or indirect lessons to prevent the breach from reoccurring—for example, the screen smudge "hack" reported by P18 taught him to wipe out the smudges from his phone screen periodically.

Experienced a security breach (6/48 examples)

Another prominent social trigger reported by our sample was experiencing a security breach—when participants or someone they knew had an account or device accessed by a stranger, or otherwise had information shared with unintended parties. In these examples, the victims of a security breach solicited advice from friends and loved ones, simultaneously spreading *awareness* (3/6 examples) of a new security threat, and *motivating* (4/6) behavior changes by grounding it in a real example of harm.

One participant initiated a new practice of updating his password on a monthly basis following his Facebook account getting breached, because his friend recommended that course of action:

"Because once I got my account hacked. And I was [doing my] bachelor's in a city, so yeah. After that I was more precautious regarding the same. And I'll keep changing my password, so on a monthly basis [because] My friends, actually they recommended me to do so. Like there's one of my friends used to do it. He said it's better to be safe than sorry, so…" (P16)

Sharing access (3/48 examples)

Another general social trigger reported by our participants was behavior change triggered by sharing a device or account with a friend or loved one—for example, modifying a password after allowing a friend to check their phone. These changes were a reflexive response to the fact that what participants desired to generally be private was now more widely available because of a transient need to share access. For example, one participant let her son use her phone and updated the passcode afterwards:

"One of my boys wanted to use my phone for something so I gave them my passcode. And not that I have anything that I don't care for them to see or anything, but after they did that then I changed it again because I just didn't want anybody to just-- I don't care if it's them or not. I don't want them to just be able to pick up my phone and do what they want with it." (P7)

While these triggers rarely raised awareness (0/3 examples) or knowledge (0/3), they seemed to be motivate participants to make a change (3/3).

Other triggers (8/48 examples)

Eight other instances of behavior change reported by our sample were triggered by other experiences, usually conversations or recommendations—for example, an authority figure recommending the use of authentication, as mentioned by P8 when asked why he first enabled mobile authentication:

"I think my boss at the time had it and he recommended it, because he leaves his phone at his desk." (P8)

Likewise, P10 mentioned adopting anti-virus software after receiving a recommendation from a friend who he considered a security expert, and P13 mentioned that she stopped using Google Chrome for financial transactions because two of her security expert friends informed her that the version of Chrome she used insecurely stored information. These recommendations often raised participants' awareness of, motivation to use *and* knowledge of how to use a new security tool or method.

Importantly, however, recommendations from authority figures didn't *always* result in behavior change. P13, for example, mentions that she ignored her boss's advice to have different passwords for different accounts because it would be hard to remember all those passwords. Nevertheless, the advice did raise her *awareness* of proper security practices.

P7 reported re-activating the PIN for her iPhone because a family member asked her why she deactivated it in the first place, urging her to reconsider. The conversation didn't raise her awareness or knowledge, but re-upped her motivation to use a security tool with a bit of social proof.

Interestingly, another participant mentioned installing anti-virus software on her laptop simply because she felt guilty, after conversing with others who attended her university's cybersecurity awareness fair, for not using software that her school provided:

"*I also felt guilty that I have all this free stuff I could install to protect my computer, and all this stuff I could do that's smart and I wasn't taking it." (P12)*

The guilt inspired behavior change reported by P12 is emblematic of the *reciprocity* principle of social influence, which suggests that people are more likely to follow the suggestions of those who did them a favor—even an unsolicited one [6].

Importantly, one participant reported how a social process urged her *against* behavior change (but was still responsible for a decision she made about security). P17 mentioned that she did not follow her security-expert husband's advice to delete unused and obscure online accounts because she noticed that her friends, who did not follow the advice, never experienced a security breach:

Catalyst	N	Description	Example
Observed Insecure Behavior	15	Noticed that someone was being insecure.	*"Right now I have ignored this storing passwords on my cell phone. He was like, 'Don't do this. It's dangerous.'" (P7)*
Observed Novel Behavior	11	Noticed a new security tool / method.	*"[I] see a lot of fancy password protection programs on [my co-workers] laptops. Like special files being encrypted. I'm like, "What's going on?" (P11)*
Sense of Obligation	15	Shared information out of obligation to protect others.	*"When I was younger, I remember my parents always telling me, like I'm sure everyone's parents tell them, to be very careful about who they give their Social Security number to. So, that's always like in my head, like if someone asks me for that, I'm just like, uh, no." (P14)*
Negative Experience	33	Experienced a security or privacy breach	*"Yes, my data got stolen. My photo got stolen on Facebook. I spoke to a couple of my friends. The only thing I could do was report abuse." (P6)*
Configuration	14	Had to set up security for a new device, account or security tool.	*"He was asking about Facebook, and he's a businessperson, so social media is somewhat of a new thing to him, and I think Facebook was-- he was just curious about it and how he could use it to kind of help his business and stuff like that. So…" (P20)*
News Article	15	Read a news article.	*"Well, before, I did not even know like I need to pay attention to this. Like I was aware of this, but I just did not know it was such a big deal. Then later, like I saw a topic, like online articles talking about that _____, talking about that, and that's when I went to the setting of like Facebook to change some." (P5)*

Table 3. Conversation catalysts derived from our iterative open coding process.

"I don't think it will be dangerous. Maybe I didn't see this kind of news or my friend didn't get some trouble when they didn't set password. Like, my friends sometimes they usually have a lot of different accounts, the same as me. But they didn't get any trouble. So I think maybe it will not be dangerous." (P17)

In this way, P17's friends' *lack* of a security breach offered her *social proof* that it's okay to ignore her husband's security advice.

4.1.2 Summary & Discussion

In summary, we analyzed 114 examples of behaviors changed, actions taken, or decisions made related to security and privacy, and found that social processes drove many (48) of those changes. We identified five common triggers for these socially driven changes, and found that these triggers were effective because they often raised participants' security sensitivity—usually awareness and motivation, but occasionally knowledge as well. These findings lend some support to the notion that social influence, especially in the form of *social proof*, *authority*, *liking*, and *reciprocity*, can be potent in raising security sensitivity—a result that bolsters the implications of prior work [14,23,29].

But, it remains unclear: is socially driven behavior change related to security and privacy as common as it could be? Socially driven change is the result of an interaction between two or more individuals—but those interactions are rare in the domain of security and privacy. Indeed, when asked why he didn't share his concerns about the U.S. government's pervasive surveillance (NSA PRISM) program, P11 stated: *"That's one thing I will never talk about."* Similarly, when asked about whether he has warned friends about a malicious smartphone application he uninstalled, P9 stated: *"Especially online. In person, it depends on the context. It does become a boring subject."*

The realization that conversations about security remain rare—and, thus, so too does the potential for socially driven behavior change related to security—begged the question: Under what circumstances do conversations about cybersecurity occur? To answer this question, we explored the 118 instances of *communication* about cybersecurity reported by our participants.

4.2 Communicating About Security

To understand the conditions under which conversations about security and privacy occur, we open coded excerpts about communication to surface triggering events for the interaction (*catalysts*) and the goal of the conversation (*conversation goal*).

4.2.1 Catalysts for Security Related Communication

We observed six primary catalysts for security related conversations in our dataset summarized in Table 3.

Insecure behavior

Some participants started a conversation about security in response to observing what they believed was insecure behavior, such as a friend or family member oversharing on social media:

"One of the reasons we talked about it is because I saw so many people post things on Facebook. A lot of times it's unnecessary things, you know, like just what they did today, "Oh, I had an amazing day," or, "I had a great dinner," and I was just talking to my husband, like why they-- I don't understand like they do that, like why they like to post things on Facebook to so-called to share." (P5)

Observing novel behavior

Relatedly, participants reported broaching conversations after observing novel security behavior or technology—for example, a new, visually appealing authentication technique. For example, one participant was stopped in a coffee shop and asked about the 9-dot authentication on his Android phone:

"We were just sitting in a coffee shop and I wanted to show somebody something and [they said], "My phone does not have that," and I was like, "I believe it probably does." (P10)

Sense of obligation

Obligations or responsibilities associated with a social role also prompted conversations about security. For example, parents lectured their children about security and privacy best practices (see example in Table 3 above), and managers informed their employees about how to manage company data because it was a

Goal	N	Description
Notify / warn	32	Notify or warn others of a potential security or privacy threat.
Prank/ Demonstrate	5	Demonstrate insecure behavior by hacking into a friend's account or device.
Share solutions	14	Share solutions, tools, and best practices (e.g., sharing how one composes his/her own password).
Vent	8	Seek social support / commiserate the experience.
Offer advice	19	Offer specific advice to others (e.g., update privacy settings, change password).
Seek advice	18	Ask for specific advice about security / privacy.
Storytelling	12	Topic was interesting/shocking/otherwise made for a good story.

Table 4. Conversation goals derived from our iterative open coding process.

part of their responsibilities. One participant described this type of interaction with his boss:

"When I was at work, I was given some sensitive documents, and I was told I couldn't send them over e-mail. I had to use a flash drive to move them over, encrypt them, then send them in e-mail." (P18)

This obligation included, in addition, a university's desire to protect its students. For example, one student talked about her university providing security solutions and advice in an annual security fair that she attended:

"They give us LoJack and all these different things you can get at the computer center. So we did talk about that. Like, locking up our computers and changing our passwords and stuff and being careful with the Wi-Fi." (P12)

Negative experiences

Negative experiences were the most common catalyst for security conversations reported by our sample. Indeed, many participants reported having conversations with friends and loved ones after experiencing a security breach. For example, one participant sought advice from friends after she received a friend request, on Facebook, from a fake profile using her own picture (see example in Table 3). Her friends recommended she report abuse in response to the attacks.

Configuration

Another frequent catalyst for discussion about security and privacy was configuring security and privacy settings on a new device, application or account. For example, one participant reported asking a friend for advice when a Facebook application asks for access to protected information:

"So there are many applications and Facebook would say that if you want to access them, there's a pop-up saying, "Allow," like, it will access all your information and stuff. So I asked him if I should go for it or not, and he tells me if it's worth going. Like, "Is it reliable or not?" (P16)

In general, participants frequently started conversations when setting or re-setting Facebook privacy settings (P13, P14, P16). In addition, many participants reported parents or older friends initiating conversations when they were setting up new computers or social media profiles for the first time (P4, P10, P15).

News articles

News articles or other press about security and privacy breaches also frequently triggered conversations. For example, one participant read and subsequently shared an article on social media about how over sharing could lead to identity theft and, more darkly, black market organ trading:

"I know there's like news talking about girls they are just so crazy about telling people on the social media where they are every minute, what they are doing every minute. So some criminals they actually use the information and just like kind of how do you say they found the girl according to her shared information online every minute. […] So I shared this article just to let my friends see just don't do it very often because I saw some of my friends on Facebook she did this really often like telling everybody what she was doing and what she had and where she was and like that." (P2)

4.2.2 Conversation Goals

We next analyzed our communication excepts for conversation goals to better understand what the conversation initiator wanted to achieve from the interaction—was it to warn others about potential threats, edify others about security tools or seek advice on how to configure security settings? During our open coding process, we identified seven distinct types of conversation goals, summarized in Table 4.

4.2.3 The Interaction of Catalysts and Goal

The interaction of conversation catalysts and goals provided enough context to answer the question: *under what circumstances do conversations about cybersecurity occur?*

To identify the most frequent conversations, we ran a cross tabulation of catalysts and conversation goal. For brevity, we focus here on the six most prevalent and interesting combinations, summarized in Table 6. These six combinations grouped into two broad categories of conversations, distinct in terms of their catalyst, focus and goal—*warnings* and *teachings*.

4.2.3.1 Warnings

Warnings were meant to raise awareness of a specific, immediate threat that had come to the attention of the conversation initiator. These warnings took three forms, varying in their catalysts, but resulted in a notification about a novel threat: cautionary tales, targeted warnings, and spreading the news.

Cautionary tales (10/118 examples)

The most common catalyst-goal combination reported by our participants was what we called cautionary tales—a conversation triggered by a negative experience on the part of the conversation initiator (or someone close to the initiator), with the goal of warning friends and loved ones about the threat. These conversations often involved sharing information about a recent security breach so that others could judge if their accounts or

Name	N	Catalyst	Content
Warning			
Cautionary tales	10	Negative experience	Notify / warn
Targeted warning	7	Insecure behavior	Notify / warn
Spreading the news	8	News article	Notify / warn
Teaching			
Lecturing	8	Sense of obligation	Offer advice
Configuration help	8	Configuration	Seek advice
Social learning	5	Novel behavior	Share solution

Table 5. The most frequent conversations about security and privacy, based on the catalyst and content.

information were in any danger. In several cases the conversation was a response to an out-of-character behavior on the part of a friend or family member. For example, after receiving odd requests for money from a friend via e-mail, one participant notified this friend that his email account was likely breached:

"Because, when I opened the e-mail, it said that they were, I think, they were in England and they didn't have enough money to come back to the States so can you send us some money, wire us some money, over, yeah. And if I'm not mistaken, I was probably the first to contact them that they were hacked. I'm like, 'This isn't right. Something strange'" (P11)

In another example, after his girlfriend illicitly accessed his e-mail account, one participant spoke to his friends to let them know that she may have read their conversations:

"It was just like, 'Hey, [my girlfriend's] been reading through our mail, like our conversations and stuff,' […] She probably read some of our conversations, not like she's going to get into your accounts." (P10)

Targeted warnings (7/118 examples)

Another common catalyst-goal combination we found was one where the conversation initiator issued a warning about potential security or privacy threats after observing others engaged in what they believed was risky behavior—what we call *targeted warnings*. For example, one participant described a friend warning her about the danger of not having a passcode:

"I was having a conversation with somebody and they were saying, 'Don't you have your passcode on there anymore?' And I said, 'No, it's a pain in the butt.' And they said, 'Well, it'd probably be a good idea if y- especially if you like leave it lay around on your desk or something like that. Or even if you're out in the evening and you have it on your purse, which most people now when they're out they have this thing right on the table where they are that somebody doesn't come by and grab it or whatever. That way they can do whatever they want with it.'" (P7)

Spreading the news (8/118 examples)

News articles about security breaches often resulted in conversations we refer to as *spreading the news*—conversations where the initiator attempted to warn friends and loved ones about a security threat outlined in a news article. These conversations sometimes included advice on how to change behavior to protect oneself from the new threat, but were usually just meant to raise awareness that a threat existed. For example, one participant talked about his contacts on Twitter discussing stories about Facebook privacy concerns without giving advice:

"Oh. Yes. People have said constantly on Twitter about how Facebook, it's not private anymore. Which is ironic, because neither is Twitter. So I've seen that, but no one has showed a article about being secure like with NSA and stuff." (P4)

As with other warnings, these conversations were often motivated by a desire to protect. For example, one participant described sharing a link to an article, through social media, about a credit card breach in order to warn her loved ones to be careful. Indeed, when asked why she shared one such news article, P2 said:

"To ask my beloved to actually pay attention to these things, to make sure they're okay. Their bank accounts are okay, if they actually do some shopping that day." (P2)

Conversations prompted by news articles also sometimes led to sharing best practices or details of privacy and security behaviors.

"We were just generally sitting around and somebody was like, 'Oh, this is an article about Facebook privacy stuff again. Let's look at it' 'Do you use this,' or 'I use that,' and 'Oh.' So really just comparing notes is the best way I can put it. Like we weren't overly scrutinizing each other's things. But like 'I found this to be effective.'" (P10)

4.2.3.2 Teachings

The other broad category of conversations we found was *teachings*. Teachings involved sharing security best practices or edifying others on how to protect themselves from security and privacy threats. In contrast to warnings, these conversations focused on sharing specific information about behaviors to enact in order to *solve* an immediate problem or *avoid* a future threat. Three conversations fell into this category: lecturing, configuration help, and social learning.

Lecturing (8/118 examples)

Conversations we referred to as *lecturing* involved advising others about security best practices, usually because the initiator felt a sense of obligation. Several of these conversations were between parents and children. Initially, parents offered children advice—for example, to not over share on Facebook. When children were older, however, they tended to be the ones lecturing their parents about privacy and security best practices. One participant described the litany of advice he gave to his parents about what to do and what not to do:

"I mean, I've spoken to my mom and dad about it. Like, I've told them, like, because I've told them to also use the same features that I do. Like having screen locks for phones and being more careful about passwords. And not logging into public computers and just leaving them without signing out." (P8)

Another type of lecturing was managers lecturing employees about security best practices to protect company data. For example, one participant described her boss asking her to regularly update her password:

"Actually, this was given to me by my manager, with whom I used to work. So he's the one who told me about this. He was like you

should change your password because it contains confidential information." (P13)

Another participant described his boss asking him to encrypt confidential files and transmit them physically on a USB flash drive rather than through email (P18).

Configuration help (8/118 examples)

Conversations about *configuration help* consisted of a conversation initiator soliciting advice on how to configure security and privacy settings for a new device or account. For example, one participant described helping his mother set up her new laptop with the appropriate security settings to keep her information safe (P19). Another participant described encouraging his mother to enable 9-dot authentication on her new Android phone to make sure no one else could access it:

"I mean, just the same reason that people shouldn't just look into her phone. Because, like, if it does not have a button, anyone can just, like, unlock and look at her messages and stuff." (P15)

Most frequently, configuration help conversations were about setting up the Facebook privacy settings (P1, P3, P4, P8, P19).

"If anything maybe my mom. I'm not sure directly security issues but she doesn't really know how to do Facebook that much so she'll ask me questions about it, in general, like how to post or, I guess, how to remove herself from something or certain things like that. So, I guess, I have given her advice in a way, just given her a few basic steps of set this as this just so you don't have--you're not completely open and public." (P19)

Social learning (5/118 examples)

In *social learning* conversations, conversation initiators observed novel security or privacy behaviors or tools—for example, a new way to compose passwords (P9, P10, P18) or a new type of authentication (P8)—that led to questions that allowed others to share information about the behavior. These conversations were opportunities for experts or early adopters to boast about their solutions for solving common security problems. For example, P18 asked a friend about sharing his Amazon account password, prompting the friend to share his password composition method:

"When I was working this summer, one of my co-workers told me about the whole algorithm thing. One, it just helps you I guess have different passwords. It helps you recall them easier based on I guess the type of profile. I guess you can cater, you can change your algorithm, depending on I guess what you want to be in it. But ever since I started using it." (P18)

4.2.4 Summary & Discussion

In analyzing the 118 conversations about security and privacy reported by our participants, we uncovered six common conversation catalysts (Table 3) and seven common conversation goals (Table 4). From these catalysts and goals, we identified six common catalyst-goal contexts (Table 5) that captured a large number of the security conversations reported by our sample, enabling us to answer the question: *under what circumstances do people generally talk about privacy and security?*

Broadly, the answer appears to be: to *warn* or to *teach*. Indeed, most commonly, our participants reported conversations about privacy and security to be educational experiences—either in sharing and receiving information about a novel security threat, or in sharing and receiving advice about how to solve a specific security problem or security best practices. This finding appears to confirm the notion that social processes *can* contribute to the heightening of security sensitivity, as these educational conversations often raised any or all of awareness, motivation or knowledge about security.

Observability, again, appeared to be a key driver of conversations—whether experts witnessing *insecure* behavior or non-experts witnessing *novel* behavior. In general, however, social learning may not have been as prevalent as would be ideal. Social learning conversations may represent the ideal context under which social influence *can* affect security sensitivity—novices interested in learning about security voluntarily ask for information from experts, thereby raising their own knowledge. In turn, experts are willing to share their information and don't feel that their efforts are wasted, as was implied by several of the security savvy participants we interviewed when asked why they don't share information about threats more often (P4, P9).

Unfortunately, many of our participants alluded to an illusory correlation [4] between security feature usage and paranoia, referring to their expert friends as "hyper-secure" (P5) and their actions as "above and beyond" (P18) or "nutty" (P1). Perhaps as a result of this negative perception towards those with high security sensitivity, many of the security savvy participants we interviewed mentioned that they avoided sharing information with their friends because the topic seemed socially inappropriate or unwelcome—as too preachy, for example. There is, thus, a substantial missed opportunity for experts to share knowledge with novices that only appears to be overcome when novices observe and query about interesting, novel behavior by the expert.

5. GENERAL DISCUSSION

Our results introduce a typology of social interaction around cybersecurity behavior and communication. First, we confirmed that social processes are an important influence on cybersecurity behavior change—a large number of behavior changes reported by our sample were driven at least partially through social processes. Specifically, we identified five common social triggers for security related behavior change—observing and learning from friends, social sensemaking (discussing ambiguous security threats with friends to determine the relevance of the threat and a clear course of action), pranks and demonstrations, experiencing a security breach and sharing access to a device with others. Furthermore, all social triggers for behavior change reported by our sample appeared to heighten security sensitivity in some way—either by increasing participants awareness of a new threat or security tool, motivating participants to protect themselves, or increasing participants knowledge of how to protect themselves.

We also found that conversations about security are primarily educational in nature, instigated mostly with a goal to *learn* or to *teach*. Many of our participants, for examples, reported having conversations about security to warn their friends and loved ones to be careful after experiencing a security breach, reading about a security threat on the news, or observing a friend's insecure behavior. Others reported specifically querying for security knowledge and advice after observing novel security behavior (e.g., the use of a new type of authentication), or if they had a specific and immediate security problem they wanted to solve (e.g., configuring the security settings of a new laptop).

Our results also emphasize the influential nature of a specific negative experience in raising the security sensitivity and, in turn, changing the cybersecurity behavior of victims and those around them. Interestingly, friends and loved ones appeared to at least indirectly take advantage of this fact, often breaking into others'

accounts to prove to that person that s/he was not fully protected. This notion of pranking by friends and family can also be considered as an effective way to create a *teachable moment*, something that past work on PhishGuru has found to be effective in teaching people about phishing attacks [20]. In other cases, pranks were not necessarily meant to directly educate victims, but were used as a form of "hazing". Either way, the breach elicited a similar reaction—both the victims of these negative experiences and the people around them who they shared the experience with became more aware of and motivated to address their own security vulnerabilities. These breaches also motivated participants to communicate with others to solve their problems.

The observability of security features and methods also proved to be important in driving behavior changes through social processes. Indeed, observing friends was the most frequent social trigger for behavior change. Nevertheless, most security features and methods are inherently unobservable and were rarely surfaced in our interviews—password composition methods, for example. When P18 learned of a new way to compose passwords from his expert friend, he immediately started utilizing this new composition policy. However, only two of our participants mentioned talking about password composition policies, suggesting there is much room for improvement in leveraging social processes to raise security sensitivity.

Observing novel or insecure behavior was also a key trigger for conversations about security and privacy, prompting novices to ask experts about novel behaviors and experts to warn novices about insecure behaviors. These conversations, again, were contingent upon the observability of the security feature or method. Experts could see the lack of mobile authentication on their friends' smartphones, but they could not see their friend's social media privacy settings, for example, and so conversations about social media privacy settings were rarely proactive—they were usually reactive, after someone encountered a breach.

However, simply increasing the observability of all security features may not be the best solution. First, security settings have historically been private—and for good reason. Indeed, past work by Gaw et al. [14] found that people who encrypted e-mail were often considered paranoid unless they were in a role where they handled sensitive company data, suggesting an illusory correlation [4] between security feature usage and paranoia. Our own interviews allude to a similar phenomenon, which appeared to be inhibit security experts from sharing their knowledge with others unless specifically asked. Indeed, as early adopters of security features are likely those who are especially concerned about their security—and, thus, are the most likely to be considered as paranoid by lay users—it is possible that making security decisions and behaviors perfectly observable might work *against* security sensitivity. After all, potential adopters may look at the present adopter list and find tenuous *social proof* that only "paranoid" people use a security feature. Second, we also saw evidence that social processes can work *against* a user following advice if it seems like none of their friends are affected by a threat. Likewise, it is possible that when a useful security feature has low current adoption, potential adopters might see the *absence* of adoption as *social proof* against using the feature.

To best leverage the positive effects of observability, therefore, it would seem that we want to facilitate more *social learning* conversations and *observing friends* behavior change. To that end, if we make security tools more visual and amenable to conversation while considering simple design for enhanced usability [36], non-experts can passively raise their *awareness* and *motivation* by observing their expert friends, and then raise their *knowledge* by voluntarily asking about security.

5.1 Limitations and Future Work

Our sample, although representative in many respects, is primarily from the US and young. Furthermore, as we solicited participants from only one online recruitment source, we could have introduced a systematic bias into our results—our participants were the type that generally volunteers for research projects. This means our results may not necessarily widely generalize, as is the case with most qualitative research. Thus, future work should examine whether the patterns and relationships identified in our data persist in a larger, representative sample of technology users. Our results are also limited to the communication and interaction instances participants could recall during our interview session—the so-called *recall* problem that afflicts retrospective interview studies [21]. Furthermore, as we only analyzed instances of behaviors changed, actions taken, and decisions made driven by social processes, we do not talk about the substantial number of non-social triggers for the same.

Our findings inform a breadth of potential future work, specifically in designing systems and interventions that *leverage* social influence processes to raise security sensitivity. For example, a key finding from our interviews was that the *observability* of security tool greatly facilitates its spread through social channels. Nevertheless, most security features are not observable, leaving little room for social spread and learning. Future work could introduce simple manipulations to increase the observability of security features and measure their effect on conversation frequency and behavior change, for example.

6. CONCLUSION

In summary, we qualitatively examined how social processes drive security-related behavior change and communications about security. Our findings suggest social processes played a major role in a large number of privacy and security related behavior changes reported by our interviewees, probably because these processes were effective at raising security sensitivity—the awareness of, motivation to use and knowledge of how to use security tools. In addition, conversations our participants had about security and privacy were most often instigated by the desire to (1) warn or protect others from immediate or novel threats observed or experienced and to (2) gather information about solving a privacy problem. One theme that arose from our interviews, especially, is that the *observability* of security feature usage was a key enabler of socially triggered behavior change and conversation—in encouraging the spread of positive behaviors, discouraging negative behaviors, and getting participants to talk about security. Altogether, our results suggest that there is a substantial and often overlooked social process that helps drive security related behavior change, and that in order to maximally raise security sensitivity, we should make security tool usage more observable and amenable to conversation. In addition, we believe our work provides a strong foundation for much needed further exploration into the social dimensions of cybersecurity behavior.

7. ACKNOWLEDGMENTS

This work was generously supported by NSF Award #1347186, the NDSEG Fellowship, and CMU's CyLab. We would also like to thank Samantha Finkelstein, Hsu-Chun Hsiao, and Ruogu Kang for helping with refining the interview protocol.

8. REFERENCES

[1] Acquisti, A. and Grossklags, J. Losses, Gains, and Hyperbolic Discounting: Privacy Attitudes and Privacy Behavior. In J. Camp and R. Lewis, eds., *The Economics of Information Security*. 2004, 179–186.

[2] Adams, A. and Sasse, M.A. Users are not the enemy. *CACM 42*, 12 (1999), 40–46.

[3] Bandura, A., Grusec, J.E., and Menlove, F.L. Vicarious Extinction of Avoidance Behavior. *Journal of Personality and Social Psychology 5*, 1 (1967), 16–23.

[4] Chapman, L.J. Illusory correlation in observational report. *Journal of Verbal Learning and Verbal Behavior 6*, 1 (1967), 151–155.

[5] Cialdini, R.B. and Goldstein, N.J. Social influence: compliance and conformity. *Annual Rev. of Psych. 55*, 1974 (2004), 591–621.

[6] Cialdini, R.B. *Influence*. Harper Collins, 2009.

[7] Davis, F.D. Perceived Usefulness , Perceived Ease Of Use , And User Accep. *MIS Quarterly 13*, 3 (1989), 319–340.

[8] Dhamija, R., Tygar, J.D., and Hearst, M. Why phishing works. *Proc. CHI '06*, ACM Press (2006), 581–590.

[9] DiGioia, P. and Dourish, P. Social navigation as a model for usable security. *Proc. SOUPS '05*, ACM Press (2005), 101–108.

[10] Dourish, P., Grinter, R.E., Delgado de la Flor, J., and Joseph, M. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing 8*, 6 (2004), 391–401.

[11] Egelman, S., Acquisti, A., Molnar, D., and Herley, C. Please Continue to Hold An empirical study on user tolerance of security delays. *Methodology*, (2010).

[12] Egelman, S., Cranor, L.F., and Hong, J. You've been warned. *Proc. CHI '08*, ACM Press (2008), 1065–1074.

[13] Furnell, S., Jusoh, A., and Katsabas, D. The challenges of undersatnding and using security: A survey of end-users. *Computers & Security 25*, 1 (2006), 27–35.

[14] Gaw, S., Felten, E.W., and Fernandez-Kelly, P. Secrecy, flagging, and paranoia. *Proc. CHI '06*, ACM Press (2006), 591–600.

[15] Goldstein, N.J., Cialdini, R.B., and Griskevicius, V. A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research 35*, 3 (2008), 472–482.

[16] Herley, C. and Oorschot, P. van. Passwords: If We're So Smart, Why Are We Still Using Them? *Financial Cryptography and Data Security*, (2009).

[17] Herley, C. So long, and no thanks for the externalities. *Proc. NSPW '09*, ACM Press (2009), 133–144.

[18] Inglesant, P.G. and Sasse, M.A. The true cost of unusable password policies. *Proc. CHI'10*, ACM Press (2010), 383–392.

[19] Kim, T.H.-J., Gupta, P., Han, J., Owusu, E., Hong, J., Perrig, A., and Gao D. OTO: Online Trust Oracle for User-Centric Trust Establishment. *Proc. CCS '12*, ACM Press (2012), 391–403.

[20] Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L.F., and Hong, J. Teaching Johnny not to fall for phish. *ACM Transactions on Internet Technology 10*, 2 (2010), 1–31.

[21] Miles, M.B. and Huberman, M. *Qualitative Data Analysis: An Expanded Sourcebook*. Sage Publications, Inc., 1994.

[22] Milgram, S., Bickman, L., and Berkowitz, L. Note on the drawing power of crowds of different size. *JPSP 13*, 2 (1969), 79–82.

[23] Rader, E., Wash, R., and Brooks, B. Stories as informal lessons about security. *Proc. SOUPS '12*, ACM Press (2012).

[24] Renaud, K. Evaluating Authentication Mechanisms. In L.F. Cranor and S. Garfinkel, eds., *Security and Usability*. O'Reilly Media, 2005, 103–128.

[25] Rogers, E.M. *Diffusion of innovations*. New York, New York, USA, 2003.

[26] Sasse, M.A. Computer security: Anatomy of a Usability Disaster, and a Plan for Recovery. *Proc. CHI '03 Wkshp on HCI and Security Systems*, Citeseer (2003).

[27] Schultz, P.W., Nolan, J.M., Cialdini, R.B., Goldstein, N.J., and Griskevicius, V. The constructive, destructive, and reconstructive power of social norms. *Psychological science 18*, 5 (2007), 429–34.

[28] Sheng, S., Magnien, B., Kumaraguru, P., et al. Anti-Phishing Phil. *Proc. SOUPS '07*, ACM Press (2007), 88–99.

[29] Singh, S., Cabraal, A., Demosthenous, C., Astbrink, G., and Furlong, M. Password sharing. *Proc. CHI '07*, ACM Press (2007), 895–904.

[30] Stanton, J., Mastrangelo, P., Stam, K., and Jolton, J. Behavioral Information Security: Two End User Survey Studies of Motivation and Security Practices. *AMCIS*, August (2004), 2–8.

[31] Suo, X., Zhu, Y., and Owen, G.S. Graphical passwords: A survey. *Proc. ACSAC'05*, IEEE (2005).

[32] Wash, R. Folk models of home computer security. *Proc. SOUPS '10*, ACM Press (2010), 1.

[33] Weick, K.E., Sutcliffe, K.M., and Obstfeld, D. Organizing and the Process of Sensemaking. *Organization Science 16*, 4 (2005), 409–421.

[34] Whitten, A. and Tygar, J.D. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. *Proc. SSYM'99*, (1999), 14–28.

[35] Zhang, Y., Egelman, S., Cranor, L., and Hong, J. Phinding Phish : Evaluating Anti-Phishing Tools. *Proc. NDSS'07*, (2007).

[36] Zurko, M. E. IBM Lotus Notes/Domino: Embedding Security in Collaborative Applications. In L.F. Cranor and S. Garfinkel, eds., *Security and Usability*. O'Reilly Media, 2005, 607–622.

[37] Dedoose. http://www.dedoose.com.

Appendix A: Additional Figures and Tables

Expanded Demographics

	Age	Gender	Race	Occupation	Phone OS	Mobile Auth	Social Media Usage
P1	28	Male	African American	Customer Service	Android	None	Daily
P2	22	Female	Asian	Unemployed	iOS	None	Daily
P3	22	Female	African American	Student	iOS	PIN	Daily
P4	22	Male	African American	Student	Android	None	Daily
P5	27	Female	Asian	Unemployed	iOS	None	Daily
P6	29	Male	White	Software Developer	iOS	None	Daily
P7	54	Female	White	Administrative Assistant	iOS	PIN	Weekly
P8	31	Male	Indian	Unemployed	Android	None	Weekly
P9	30	Male	White	Software Developer	Android	None	Weekly
P10	37	Male	White	Graphic Designer	Android	9-dot	Daily
P11	54	Male	African American	Chef	Android	None	Weekly
P12	20	Female	African American	Student	iOS	None	Daily
P13	24	Female	Indian	Graduate Student	Android	None	Daily
P14	25	Male	Indian	Graduate Student	Android	PIN	Daily
P15	21	Male	Indian	Graduate Student	Android	9-dot	Daily
P16	22	Male	Indian	Graduate Student	Android	9-dot	Daily
P17	34	Female	Asian	Unemployed	iOS	None	Daily
P18	20	Male	African American	Student	Android	9-dot	Daily
P19	20	Male	White	Student	Android	9-dot	Daily

Table A1. Expanded participant demographics.

Co-Frequency of Catalysts and Reasons for Conversations

	Offer Advice	Share Solution	Vent	Seek advice	Notify or Warn	Storytelling	Prank or Demonstrate	Other	**Total**
Sense of Obligation	**8**	2	0	0	2	3	0	0	**15**
Insecure Behavior	4	0	1	0	**7**	0	2	1	**15**
Negative Experience	3	3	5	7	**10**	2	2	1	**33**
Configuration	2	2	1	**8**	0	0	0	1	**14**
News Article	1	0	0	0	**8**	3	0	3	**15**
Observed Novel Behavior	0	**5**	0	3	0	2	0	1	**11**
Other	1	2	1	0	5	2	1	3	**15**
Total	**19**	**14**	**8**	**18**	**32**	**12**	**5**	**10**	**118**

Table A2. Co-frequency of catalysts for conversations about security and privacy (rows) and reasons for starting the conversation (columns).

Inter-Coder Reliability for Each Applied Code

Code	Inter-Coder Agreement
Behavior Change: Social or Non-Social	0.93
Behavior Change: Trigger Event	0.87
Behavior Change: Raised Awareness	0.87
Behavior Change: Raised Motivation	0.80
Behavior Change: Raised Knowledge	0.80
Communication: Catalyst	0.71
Communication: Reason	0.86

Table A3. Inter-coder agreement of codes applied in our analysis, calculated from a 20% overlap of coded excerpts by two coders.

Appendix B: Recruitment Materials

We solicited study participants through CBDR, an online research study participation pool maintained by Carnegie Mellon's Department of Social and Decision Sciences. Below we show the posting for our study.

Study Name: ($) Talk to us about cybersecurity

Description:

Participate in an interview about how you learn about and manage online privacy and cybersecurity—for example, about mobile phones, passwords and social media privacy settings. We are researchers in the Human-Computer Interaction Institute at Carnegie Mellon University. We are studying how people learn about and manage cybersecurity. Please bring your smartphone and laptop for the study. We may ask you to show us your smartphone's home screen, and we may ask you to log into your Facebook account using your own laptop.

Eligibility: You must be (1) 18 or over, (2) a regular Android or iOS smartphone user, and (3) a Facebook user

Duration: 45 minutes

Pay: 10 Dollars