Increasing Security Sensitivity With Social Proof: A Large-Scale Experimental Confirmation

Sauvik Das
Carnegie Mellon University
sauvik@cmu.edu

Adam D.I. Kramer
Facebook, Inc.
akramer@fb.com

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

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

ABSTRACT

One of the largest outstanding problems in computer security is the need for higher awareness and use of available security tools. One promising but largely unexplored approach is to use *social proof*: by showing people that their friends use security features, they may be more inclined to explore those features, too. To explore the efficacy of this approach, we showed 50,000 people who use Facebook one of 8 security announcements—7 variations of social proof and 1 non-social control—to increase the exploration and adoption of three security features: Login Notifications, Login Approvals, and Trusted Contacts. Our results indicated that simply showing people the number of their friends that used security features was most effective, and drove 37% more viewers to explore the promoted security features compared to the non-social announcement (thus, raising awareness). In turn, as social announcements drove more people to *explore* security features, more people who saw social announcements *adopted* those features, too. However, among those who explored the promoted features, there was no difference in the *adoption rate* of those who viewed a social versus a non-social announcement. In a follow up survey, we confirmed that the social announcements raised viewer's awareness of available security features.

Categories and Subject Descriptors

H.1.2 **[MODELS AND PRINCIPLES]:** User/Machine Systems—*Human factors*

General Terms

Experimentation, Security, Human Factors

Keywords

Social Cybersecurity, Facebook, Social Influence, Persuasion, Security Feature Adoption, Security

1. INTRODUCTION

In 2013, the Associated Press's Twitter account was compromised through a phishing scheme. The intruders misleadingly tweeted that President Obama was injured in a bombing [28], plummeting stock prices [20] and adversely affecting thousands. Moreover, this break-in could have been easily prevented with two-factor authentication—a security feature, available at that time, that requires entry of a pseudo-random code generated on a person's smartphone in addition to a password when authenticating [17].

This incident is just one example of how the underutilization of available security features can often have dire consequences, and illustrates how the need for higher *security sensitivity* [9]—the awareness of, motivation to use, and knowledge of how to use security tools—remains one of the largest outstanding problems in computer security today. Indeed, while two-factor authentication may not be necessary for every person for every service, widespread awareness and utilization of available security features is critically important.

Recent work suggests that one promising approach to widespread heightening of security sensitivity is through social proof—or, our tendency to look to others for cues on what to use and how to behave [6]. Much work in social psychology has shown that social proof is powerfully effective at driving human behavior: for example, at reducing household energy consumption by showing people their neighbors' reduced energy consumption [24], reducing hotel guests' wasteful use of towels by showing them that 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 [2]. In a small interview study, Das and colleagues [9] found that this result might extend into the security domain, as *observing others* use security tools and behaving securely was a key enabler for security related behavior change among their participants.

Historically, however, security feature usage has been kept confidential to preserve an individual feature-user's privacy, and this hiding of security feature use has both stifled the social diffusion of security features *and* made it difficult to test the effect of social interventions on increasing people's security sensitivity. Consequently, the security community has overlooked a potentially fruitful avenue for increasing security sensitivity, as there is a dearth of empirical data conclusively linking social-proof based interventions to heightened security sensitivity.

Here, we share among the first results experimentally confirming whether and how social proof can be used to raise security sensitivity. We designed a set of 7 social-proof based security announcements that can preserve the privacy of individuals who use security features and provide their friends with social proof that others they know use security features. All social announcements informed viewers that their friends used security features, but the seven variations differed in their *specificity* (i.e., showing viewers exactly how many of their friends used security features versus just saying that "some" of their friends used security features) and *framing* (i.e., using keywords such as "only" or "over" to prime viewers' interpretation of the text).

Then, to test the efficacy of social proof on increasing people's awareness of and use of security features, we showed n=50,000 people who use Facebook one of 8 security announcements—our 7 variations of social proof and 1 non-social control—intended to increase the awareness and adoption of three Facebook security

CCS'14, November 03 - 07 2014, Scottsdale, AZ, USA

ACM 978-1-4503-2957-6/14/11$15.00.
http://dx.doi.org/10.1145/2660267.2660271

features: Login Notifications, Login Approvals, and Trusted Contacts (described below).

We found that while all of our social-proof based interventions were effective, simply showing people the specific number of their friends that used security features without any subjective framing was most effective—driving 37% more viewers to explore the promoted security features compared to the non-social announcement (thus, raising awareness). Furthermore, the effect of social proof strengthened when a viewer had more friends who already used security features. In turn, as social announcements drove more people to explore security features, more people who saw social announcements adopted those features, too. However, comparing just those who clicked on any of the announcements, there was no difference in the *adoption rate* of those who viewed any of the social announcements relative to the non-social announcement. Finally, in a follow up survey, we confirmed that social announcements can at least indirectly raise people's awareness of the availability of additional security features.

2. BACKGROUND

Prior work in usable security alludes to three main reasons underlying why many security features remain unused: the need for greater awareness, motivation, or knowledge. Das and colleagues [9] coin this three-layered stack *security sensitivity*.

First, many people 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 people to construct their own model of security threats that is often incorrect, potentially leaving them vulnerable to security breaches [1]. Second, many people—even those who are aware of security and privacy threats and the preventive tools to combat those threats—often lack the *motivation* to utilize security features to protect themselves [1,13]. 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 [12,23]. Finally, security tools are often too complex to operate for even those who are aware and motivated, suggesting that many people lack the *knowledge* to actually utilize security tools [27]. Indeed, there is a wide gulf of execution for most security features for most people. For example, many cannot distinguish between legitimate and fraudulent URLs or email headers [10].

Efforts have been made at improving all parts of the security sensitivity stack—for example, through games for security education [25], browser extensions to make people more aware of phish [29], more effective user interfaces for security tools [11], and simpler ways to authenticate [8]. Security sensitivity, nevertheless, could be much higher. We take the stance that because people look to others around them for cues on how to act in uncertain circumstances [6], we can offer them *social proof* that their friends use security tools to heighten at least their *awareness of* and *motivation to use* security tools.

Prior work in cognitive psychology has demonstrated the potency of social proof. For example, Milgram, Bickman, and Berkowitz [19] 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. More recent studies on online platforms such as Facebook have similarly alluded to the potency of social proof. Kramer [18] showed that users were more likely to share emotional content matching the emotional valence of content shared by friends in the past few days, and Burke and colleagues [4] showed that social learning plays a substantial role in influencing how newcomers to Facebook use the platform. Notably, Bond and colleagues [3] found that simply showing people that their Facebook friends voted was sufficient to increase voter turnout in the 2010 U.S. Congressional elections.

Others have looked at the effect of social processes in the adoption of technology, specifically. Indeed, in his seminal work on the diffusion of innovations, Rogers argued that new technology gets widely adopted through a process by which it is communicated through members of a social network [22]. He further outlines that preventative innovations—or innovations, like security and privacy tools, that prevent undesirable outcomes from happening in the future—typically have lower adoption rates, probably because of their lack of *observability* (i.e., the invisibility of their benefits and use).

Still other work has shown that there is, indeed, a social component to peoples' perceptions about and use of security tools. Rader and colleagues showed that people often learn about security from informal stories told by one another [21]. Singh and colleagues outlined the common practice of sharing passwords and PINs, emphasizing social practices [26]. And, Das and colleagues found that many behavior changes related to security and privacy are driven by social processes, and found that the observability of security feature usage among strangers and friends was a key component in increasing security sensitivity [9].

Nevertheless, while all this background work *alludes* to the potential efficacy of social proof in heightening security sensitivity, there is a substantial lack of work that has employed social cues to elicit security related behavior change. Part of the problem is that security feature usage has historically been kept secret to preserve the privacy of individual feature-users. Still, as social channels are the primary way through which innovations spread [22], the hiding of social meta-data surrounding security feature usage has undoubtedly inhibited both the widespread adoption of security features *and* research in studying social cues as a way to heighten security sensitivity.

The little empirical data we *do* have about the effects of social influence on security related behavior change comes from work that only treated the social dimension in passing. Egelman and colleagues [14] included a simple social condition in their study on the effects of various types of password meters on convincing people to create stronger passwords. They found that a "peer pressure" password meter that showed participants how strong their passwords were relative to other "users" performed no better in increasing the strength of participants' composed passwords, as compared to a standard password meter that told participants whether their passwords were "weak", "medium" or "strong". However, Egelman and colleagues' "peer pressure" password meter measured participants' passwords relative to strangers' passwords for a completely different service, and provided little feedback as to whether a given meter reading was important enough to act upon (is it good or bad that my password is better than 50% of "others"?). In addition, their social intervention could only have an affect on participants' motivation—the part of security sensitivity that will likely prove most difficult to increase.

Taken together, all this prior work strongly suggests that increasing the observability of friends' security feature use can heighten people's security sensitivity, though Egelman and colleagues' [14] null result with their peer pressure password meter suggests that the *specificity* and *framing* of social information may moderate its effect. To test these conjectures, in this work, we sought to answer the following questions: (1) Does increasing the observability of security feature usage drive the

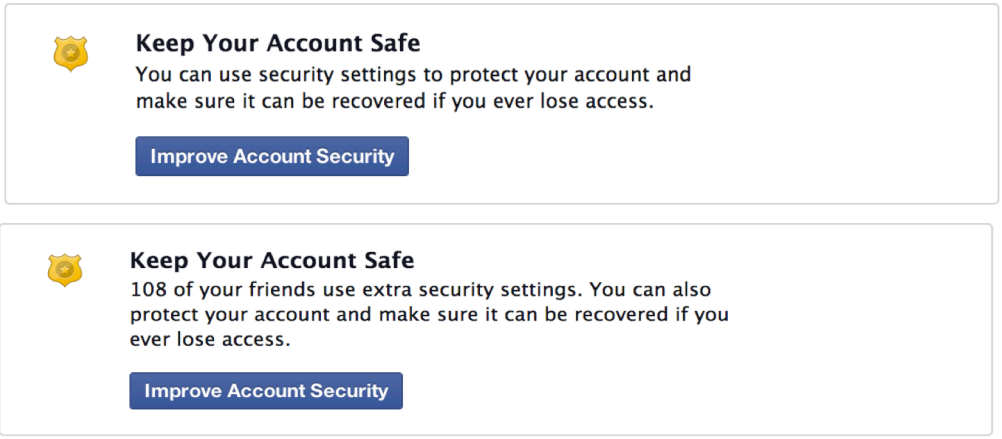

Figure 1. Image of the control (top) and Raw # (bottom) social prompts rendered onto users' news feeds.

exploration and adoption of security features? (2) Does the framing of social information affect the exploration and adoption of security features? For example, is it more effective to frame the social information in a way that suggests that not enough of one's friends use security features, and thus the viewer should lead the way? Or is it more effective to frame the social information in a manner that suggests that many of the viewer's friends already use extra security settings, so the viewer should join these savvy friends? and, (3) Does specificity in the social cue matter? Is it enough to simply inform users that "some" of their friends use extra security features rather than directly inform users about the exact numbers?

While it has historically been impossible, or at least very difficult, to answer these questions because of the confidentiality of security feature use, today, with the rich and nigh-complete social meta-data on platforms such as Facebook, we can design simple social cues that show non-adopters social proof that their friends use security features while preserving the individual privacy of those same security-feature users. To that end, in the first large-scale study on raising security sensitivity with social proof, we measure the effect of showing people simple social cues on security feature exploration and adoption on Facebook.

3. SOCIAL PROMPT EXPERIMENT

In our initial experiment, we showed 50,000 people who use Facebook one of eight announcements, pinned at the top of their Facebook newsfeed, informing them about the availability of extra security features on Facebook. Seven of these announcements included a social cue informing viewers that their friends also used security features, but varied in their *specificity* (i.e., showing the exact number of friends versus just saying "some" friends) and *framing* (i.e., priming the interpretation of the social cue with keywords such as "only" and "over"). None of the announcements revealed any information about individual feature users, however, thus providing aggregated social proof *without* surfacing *who* was using *which* features. We measured whether the nature of the text in the announcement (social vs. non-social, the framing and specificity of the social proof text) led to greater exploration of available security features and greater adoption of security features—or, increased *awareness of* and *motivation to use* security features, respectively.

3.1 Methodology

People in our sample who logged on to Facebook between November 4th, 2013 and November 8th, 2013 were shown one of eight announcements informing them that they can use extra security features to protect their Facebook accounts. The announcements were rendered at the top of their newsfeeds—the portion of Facebook's user interface where people are directed when they first log in, where they see an assortment of content shared by their friends. All announcements contained a call-to-action button (labeled "Improve Account Security") that directly linked people who clicked on the button to an interstitial that explained the benefits of the three security features we promoted (described below) and allowed viewers to enable the features.

Announcements were shown at most three times to the same person over the course of the four days, in order to mitigate the effect of greater exposure to those who were more active.

3.1.1 Experiment Groups

We designed and implemented four social framings to test not only whether and how social-proof cues can increase people's security sensitivity, but also if the specificity and framing of those cues matter. We refer to these framings as "Over", "Only", "Raw", and "Some". The "Over" framing informed viewers that more than a certain number or percent of their friends use extra security features, priming viewers to interpret the social cue as there being abundant social proof that others they know use security features: i.e., "many people do this, so I should too." The "Only" framing takes a contrasting approach, framing the social cue in a manner that suggests that only a few of a viewer's friends use security features so they should be among the first of their friends to secure their account. The "Raw" framing eliminates the subjective framing altogether and simply presents the viewer with the quantity of her friends who use security features. Finally, the "Some" framing is intentionally ambiguous: informing viewers only that a positive number of their friends use security features.

The "Over", "Only", and "Raw" framing had two forms: a *number* form where the number of the viewer's security-feature using friends was rendered in the announcement, and a *percentage* form where the percentage of the viewer's security-feature using friends was rendered in the announcement. In total, thus, there was one control group, two "Over" framing groups, two "Only" framing groups, two "Raw" framing groups, and one "Some" framing group, for a total of 1+2+2+2+1=8 experimental groups. The eight experimental groups are summarized in Table 1, and a representative image of the announcements shown to our sample is shown in Figure 1.

Table 1. Prompt text in announcement across all 8 experimental groups. Some social groups have templates that are filled in with either the number or percentage of a user's security feature-using friends.

Group		Prompt Text
Control		You can use security settings to protect your account and make sure it can be recovered if you ever lose access.
Social conditions (Prefix + Control Prompt)		
Over	#/%	Over X of your friends use extra security settings. You can also protect your account and make sure it can be recovered if you ever lose access. *[Note: X rounded down to nearest 5th (e.g., 108 becomes 105)]*
Only	#/%	Only X of your friends use extra security settings. Be among the first to protect your account and make sure it can be recovered if you ever lose access.
Raw	#/%	X of your friends use extra security settings. You can also protect your account and make sure it can be recovered if you ever lose access.
Some		Some of your friends use extra security settings. You can also protect your account and make sure it can be recovered if you ever lose access.

3.1.2 Sample

We selected a random sample of n=50,000 people from the U.S. who used Facebook in English, were at least 18 years old, logged on to Facebook at least once in the month preceding the experiment, had at least 10 friends who enabled one of the promoted security features, and had not enabled any one of the security features we were promoting. We evenly assigned the n=50,000 people in our sample into one of the aforementioned eight experiment groups, amounting to 6,250 people per group. This assignment was mostly random, with the constraint that people assigned to the Over condition had to have at least 10% of their friends who enabled security features, and people assigned to the Only conditions had to have fewer than 10% of their friends who enabled security features.

Our participants were 40 years old on average (s.d., 16), and 68% were women, suggesting that our sampling criteria had a bias towards older females. Notably, our sampling criteria was also biased towards active, non-security experts, but we do not believe this to be a stifling limitation given that active, non-security experts are the intended target for interventions aiming to heighten security sensitivity, as these people potentially face the greatest risk of having their accounts compromised.

Finally, the n=50,000 sample size we selected for our experiment comfortably exceeded the 4,000 participant sample size suggested by a power analysis for generalized linear models [7], with 26 coefficients, a significance level of 0.001, a power of 0.999, and a very modest effect size of 0.02—i.e., a prediction that the best social announcement will only introduce 2% more clicks relative to the control condition. In practice, we expected the effect size to be greater than 2%, but we selected a low effect size for the power analysis to get an upper bound on the number of users we needed to obtain significant results for our experiment. The 26 coefficients in our model comprised of the 18 variables described in Table 2, in addition to seven categorical variables representing the experimental conditions, and one intercept variable.

3.1.3 Promoted Security Features

We decided to promote the following three security features in our initial campaign:

Login Notifications: A security feature that informs users, via text and/or e-mail, whenever their Facebook account is accessed under suspicious circumstances: e.g., from a city the person had not previously visited.

Login Approvals: A two-factor authentication security feature that requires users to enter a randomly generated security code (sent to or generated on their phone) in addition to their passwords in order to authenticate.

Trusted Contacts: A security feature that allows users to specify 3-5 friends who can vouch for the user's identity if she forgets her Facebook account password and cannot access her e-mail.

These three security features were all co-located within the "security settings" menu context in Facebook's user interface. We chose to promote three security features to avoid drawing conclusions specific to any single security feature, and because these features represented a wide range of definitions for "security features"—with Login Notifications simply informing people of potential breaches, Login Approvals adding an extra step to the authentication process, and Trusted Contacts asking people to draw in their friends to help protect their accounts.

3.1.4 Dataset

We measured click-through rate for each announcement, as well as the short-term and long-term adoption rate of the promoted security features up to a week and 5 months after running the experiment, respectively. We used click-through rate on the announcement as a proxy for raising *awareness* (as people who clicked on the announcement were taken to explore the promoted security features), and adoption rate as a proxy for raising *motivation* (as people who adopted security features must have gained the motivation to enact a behavior change). We could not measure the differential effects of the announcements on *knowledge*, however, as all announcements led viewers to the same interstitial with the same information.

In addition, we collected each viewer's number of security-feature-using Facebook friends, and a set of behavioral (e.g., frequency of posts and comments), demographic (e.g., age, gender) and social network descriptor (e.g., mean friend age, mean friend-of-friend count) control variables that we expected might affect click-through rate and security feature adoption among our sample. These variables are described in Table 2.

3.2 Hypotheses

Cialdini's [6] concept of *social proof* suggests that when we are confronted with making a decision where we are uncertain of the appropriate course of action—like adopting a security feature, say—we look to our friends and those around us for cues on how to act. Combined with Rogers [22] assertion that *observability*—or, the visibility of the use and benefits of an innovation—is critical to the widespread adoption of an innovation, and Das and colleagues' [9] confirmation that the observability of security feature usage is a major positive factor in security and privacy related behavior change, we predicted:

H1: Social announcements will have higher click-through rates than the non-social control.

Extending the idea that social proof is more convincing when people see larger groups conforming to an action [19], we also predicted:

H2a: People with more security-feature using friends will be more likely to click on the announcement.

Table 2. Collected feature descriptions and distributions for the n=50,000 people in our sample. † Approximate values.

Demographic Variables	
Age	Age of the user.
Gender	Self-reported gender: male or female.
Friend count	Count of the user's number of friends.
Account length	Days that have passed since the user activated his/her account.
Social Network Variables	
Mean friend age	Average age of friends.
Friend age entropy	Shannon entropy of friend ages.
Percent male friends	Percentage of friends that are male.
Mean friends' account length	Average number of days the user's friends have used Facebook.
Friend country entropy	Shannon entropy of countries from which the user has friends.
Mean friend of friend count	Average number of friends of friends.
Behavioral Variables (all aggregated across the week prior to data collection)	
Posts Created	Number of posts created.
Posts Deleted	Number of posts deleted.
Comments Created	Number of comments created.
Comments Deleted	Number of comments deleted.
Friends Added	Number of friends added.
Friends Removed	Number of friends removed.
Photos Added	Number of photos added.
Social Variables	
Feature-using friends	Number of friends who use security features.

H2b: People with more security-feature using friends will be more likely to adopt a security feature, both in the short and long-term.

Similarly, we predicted experiment groups that rendered higher values or otherwise suggested that *more* rather than *fewer* of the viewer's friends used security features would be more effective at getting users to click on the announcement and explore security features. Thus, we expected that "number" conditions would have higher click-through rates than their "percent" counterparts, as the former generally render higher numbers in the announcement (e.g., **20** friends vs. 20/400=**5%** of friends). Furthermore, as the "Raw" framing rendered the highest values, followed by the "Over" and then the "Only" framing, we expected that the click-through rates for these framings would fall in that order as well.

H3a. The "number (#)" context conditions will have higher click-through rates than their "percent (%)" counterparts.

H3b: The "Raw" framing will have the highest click-through rate, followed by the "Over" and then "Only" framings.

Next, as one of the driving forces for social proof is a search for a clear course of action in an unclear circumstance [6], we also suspected that clearer, more informative messages would be more effective at driving click-through rate.

H4: Less ambiguous social framings will have higher click-through rates. Thus, the "Some" context will have the lowest click-through rate.

For short-term adoptions, we expected that the effects of social conditions would be muted. Indeed, while it is cheap—in terms of time and effort—for people to explore and gather information about security features, it can be expensive for them to actually activate those features. For example, activating Login Approvals would require people to spend an extra few seconds every time they "logged in" to their Facebook accounts. Taken together with the previous finding that people generally only enact security and privacy related behavior change after personally experiencing or hearing about a threat [9], and Egelman and colleagues' finding that a "peer pressure" password meter did not raise people's motivation to create stronger passwords relative to a non-social password meter [14], we expected that, in the short term, there would be no difference in security feature adoption rate among those who view social and non-social announcements.

H5: The adoption rate for the promoted security features should be the same for those who view a social or a non-social announcement in the week following the experiment.

On the other hand, we expected that there *should* be a long-term increase in the overall security feature adoption among users in the social condition. While our experiment lacked a strong *catalyst* for security behavior change, we expected that people in the social conditions might more strongly retain the information that extra security features are available for when they *do* encounter a compelling catalyst (e.g., hearing about a security breach on the news or through a friend). As a number of highly publicized security vulnerabilities were surfaced in the five months following the experiment (including the widely publicized "Heartbleed" bug in OpenSSL [30]), we arrived at:

H6: The adoption rate for the promoted security features should be higher for those who view a social announcement compared to those who viewed a non-social announcement in the 5 months following the experiment.

3.3 Results

Out of the 50,000 people in our sample, 46,235 logged in to Facebook within the duration of the experiment and were shown an announcement. Across all conditions, 5971 (13%) people clicked on the announcement to explore the promoted security features, while 1873 (4%) people adopted one of the promoted security features within the following week, and 4555 (9.9%) within the following five months. In Table 3, we show an aggregated breakdown of clicks and adoptions across experiment groups. The raw data suggests that all social conditions had higher click-through rates than control, the best social announcements elicited higher adoption rates in the short and long term, and the "Raw #" announcement generally performed best of all.

To statistically test whether and how the existence of, specificity, and framing of the social cue in the announcement affected click-through rate and security feature adoption, we ran three logistic regressions for clicks, short-term adoptions, and long-term adoptions. The response variables for our three models were,

Table 3. Clicks and adoptions by experimental conditions. "N" represents the number of users who viewed the announcement. "ST" stands for short term, and "LT" stands for long term. These values are strictly *descriptive*. Statistical tests used and significance is mentioned where relevant in the text.

Group	N	Clicked	Adopted ST	Adopted LT
All Conditions				
Raw #	5862	846 (14.4%)	280 (4.8%)	623 (10.6%)
Some	5828	835 (14.3%)	243 (4.2%)	602 (10.3%)
Over #	5770	779 (13.5%)	248 (4.3%)	547 (9.5%)
Only #	5668	748 (13.2%)	225 (4.0%)	548 (9.7%)
Over %	5761	724 (12.6%)	223 (3.9%)	557 (9.7%)
Only %	5708	714 (12.5%)	221 (3.9%)	555 (9.7%)
Raw %	5953	730 (12.3%)	225 (3.8%)	573 (9.6%)
Control	5685	595 (10.5%)	208 (3.7%)	550 (9.7%)
Social vs. Non-Social				
Social	40550	4376 (13.3%)	1665 (4.1%)	4005 (9.9%)
Control	5685	595 (10.5%)	208 (3.7%)	550 (9.7%)
Social Number vs. Social Percent				
Number	17300	2373 (13.7%)	753 (4.4%)	1718 (9.9%)
Percent	17422	2168 (12.4%)	669 (3.8%)	1685 (9.7%)
Social Contexts				
Raw	11815	1576 (13.3%)	505 (4.3%)	1196 (10.1%)
Over	11531	1503 (13.0%)	471 (4.1%)	1104 (9.6%)
Only	11376	1462 (12.9%)	446 (3.9%)	1103 (9.7%)

respectively, binary values representing (i) whether or not an individual had clicked on the announcement they were shown, (ii) whether or not an individual had adopted any of the three promoted security features in the 7 days following our experiment, and (iii) whether or not an individual had adopted any of the three promoted security features in the 5 months following the experiment. Our independent variable was which of the eight social announcement an individual had seen, and we also included, as controls, the behavioral, demographic, and social network descriptor variables listed in Table 2. For the two adoption models, we included an additional control representing whether or not an individual had actually clicked on the announcement they were shown to "Improve Account Security".

In Table 4, we show the logistic regression coefficients for our independent variables predicting clicks, short-term adoptions and long-term adoptions. Appendix 1 contains the full regression table, including coefficients for the control variables in our model. Coefficients in Table 4 represent a change in "log-odds", or $ln\frac{P}{1-P}$, where P represents the probability that the user clicked on the announcement or adopted one of the three security features, depending on the model. A positive coefficient implies that the log-odds ratio increases, or that the variable for the coefficient increases the likelihood that the viewer clicked on the announcement or adopted a security feature. A negative coefficient implies the opposite. Furthermore, all variables are centered and scaled, such that the coefficient for each variable represents the expected change in log-odds that an individual uses a feature given a one standard deviation increase in the predictor variable, holding all other numeric variables at their means and categorical variables at their baselines. Additionally, larger absolute coefficient values imply a stronger relationship between the independent and dependent variables.

For example, the *feature-using friends* variable (i.e., the number of one's friends who use security features) coefficient for the "clicks" model is 0.09; thus, a one standard deviation increase in this variable increases the log-odds that a viewer clicks on the announcement by 0.09, and the actual odds by $e^{0.09} = 1.09$. More

Table 4. Coefficients for the three regressions predicting clicks (CL), feature adoptions up to a week after the experiment (A-ST), and feature adoption up to 5 months after the experiment (A-LT). A full regression table including coefficients for control variables are provided in Appendix 1.

	Variable Name	CL	A-ST	A-LT
†	Group: Over #	0.29 **	-0.07	-0.13
†	Group: Over %	0.21 **	-0.12	-0.06
†	Group: Only #	0.26 **	-0.16	-0.09
†	Group: Only %	0.19 **	-0.12	-0.05
†	Group: Raw #	0.36 **	-0.01	-0.001
†	Group: Raw %	0.17 **	-0.15	-0.06
†	Group: Some	0.35 **	-0.18	-0.03
	Feature-using friends	0.09 *	0.17 *	0.20 **
	Clicked on Announcement	N/A	4.38 *	1.94 *

*† Baseline: Control, * $p < 0.05$, ** $p < 0.001$*

concretely, our model predicts that someone with 80 security feature-using friends (one standard deviation above the mean) is 9% more likely to have clicked on the security announcement, compared to the average person in our sample.

From Table 4, we can see that, relative to the control condition, all social experiment conditions do elicit higher click-through rates for announcements, as evidenced by the positive *and* significant coefficients for every experiment condition coefficient. The "Raw #" (b_{CL}=0.36, p<0.001) condition had the highest click through rate, at 14.4%—a substantial 37% increase relative to control. Even the least effective social condition—the "Raw %" (b_{CL}=0.17, p<0.001) condition—significantly enhanced click-through rate relative to control, up to 12.3%. There does, therefore, appear to be strong evidence in favor of **H1—**that all social conditions will improve click-through rate relative to the control condition. The effect is both significant and substantial.

There is also support for both **H2a** and **H2b—**that people with more security-feature using friends will be more likely to click on the announcement and adopt a promoted security feature. The *feature-using friends* variable (b_{CL}=0.09, p<0.05; $b_{A\text{-}ST}$=0.17, p<0.05; $b_{A\text{-}LT}$=0.20, p<0.001) has a large and positive coefficient for all three models, suggesting that people with more security-feature-using friends are more likely to click on the announcement *and* actually adopt a security feature relative to the average person in our sample (with all numeric variables at the mean and categorical variables at the baseline).

The data, however, is not as clear in its support for the hypothesis that social framings that suggest *more* rather than *fewer* of a viewer's friends use security features will be more effective at driving click-through rate on the security announcement. There does appear to be support for **H3a—**that number conditions will outperform percent conditions in driving click-though rate on the security announcement. Indeed, all number conditions significantly outperformed all percent conditions, and, in aggregate, number conditions elicited 7% more clicks than percent conditions (χ^2(1, n=34,722)=12.3, p=0.0004). However, we found no support for **H3b**—that the "Raw" framing would outperform the "Over" framing, which, in turn, would out perform the "Only" framing in driving click-through rate. While the aggregated click-through rate of these framings do fall into the expected sequence (Raw=13.3%, Over=13.0%, Only=12.9%), the difference is not significant despite massive power (χ^2(2, n=34,722)=1.2, p=0.54).

Thus, while social announcements that suggest that *more* rather than *fewer* of a viewer's friends are currently using extra security

features can be more effective at getting people to click on the announcement, the specific framing of the social text does not appear to significantly impact its click-through rate.

Relatedly, there is *contrary* evidence for **H4**—that ambiguous social framings such as the "Some" framing will be less effective at driving click-through rate for the announcement. In fact, the "Some" (b_{CL}=0.35, p<0.001) framing is the second most effective experimental group in driving click-through rate, after the "Raw #" (b_{CL}=0.37, p<0.001) condition, with an overall click-through rate of 14.5%.

We derived **H4** from a simple understanding of social proof—if people look to their friends for cues on how to act during periods of uncertainty, then ambiguous cues are probably less effective than clear cues. However, in reality, the ambiguity appears to elicit more interest in the announcement than most of the more specific social framings. Perhaps this finding can be explained by the intuition that people may overestimate the number of their friends who use security features when it is left ambiguous. Future work can validate this hypothesis by looking at the discrepancy between people's perceptions of the number of their security-feature-using friends and the actual number of their security-feature-using friends.

Next, there appears to be support for **H5—**that social prompts will not be significantly more effective at driving feature adoption in the short-term than non-social prompts. Indeed, all of the coefficients for the social conditions are insignificant in the short-term adoptions model in Table 4. We expected this result for two reasons: (1) people usually only adopt security tools after experiencing a "catalyst" for security behavior change—for example, in the form of experiencing a security breach or hearing about a security breach [9], and (2) the social text is not reinforced in the security interstitial where people must actually make the decision to adopt a security feature—thus, as with Egelman and colleagues' study [14], potential adopters are *not* given enough social context at the moment of potential behavior change—for example, who among their friends use what security tools.

More surprising, however, is that this negative result holds even for long-term adoptions, disconfirming **H6—**that social announcements *will* be significantly more effective at driving security feature adoption in the long term relative to the non-social announcement. In the 5 months following the experiment, a number of widely publicized security vulnerabilities that could have served as catalysts for security behavior change were highly publicized (e.g., Heartbleed [30], the iOS SSL bug [31]). Nevertheless, there was no significant difference in adoption rate between those who saw the social and non-social announcements, perhaps because the social announcements were not more memorable. We also note, however, that H6 may in fact be valid, but only with respect to relevant security threats that are presented on time and in context: Activating Login Approvals would not have been a direct answer to Heartbleed or the iOS SSL bug, so the latter may not have easily triggered a memory of the former.

Importantly, the immediate *cascading* effects of raising people's awareness of security features should not be ignored. While there is no significant difference in the *rate* of feature adoption between people who *clicked on* either the social or non-social announcement, as significantly more people clicked on the social announcements, many more people who saw social announcements also actually *adopted* security features. Indeed, from Table 3, we can see that 280 of 5862 (4.8%) people shown the "Raw #" announcement adopted one of the promoted security features over the 7 days following the experiment, compared to just 208 of 5685 (3.7%) people shown the non-social announcement (χ^2(1, n=11,547)=8.7, p=0.003). In other words, significantly more people who saw a social announcement adopted the promoted security features because significantly more people *clicked* on the social announcements.

3.4 Discussion

We found that increasing the observability of security feature usage can be effectively used to increase both *awareness* of and *adoption* of available security features. Furthermore, this effect increases with the number of the viewer's friends who use security features. Furthermore, while neither the framing of a social cue nor its specificity appeared to have a large effect on raising click-through rate, social announcements that rendered the number of a viewer's friends that used security features, rather than the percent of the same, elicited higher click-through rates. On the other hand, we found no evidence that the aggregated social announcements were more effective than a non-social announcement at raising a viewer's *motivation* to use the promoted security features. Indeed, the rate of feature adoption among viewers who clicked on either a social or non-social announcement was the same.

In practice, we believe that social announcements should still be more effective than non-social announcements at raising people's motivation to use security tools, but they need to be presented in context at the time of potential behavior change, rather than at a time when people are not actively considering changing their security behavior. Furthermore, it should be noted that people who click on a *social* announcement are probably *less* likely to adopt security features than those who click on a *non-social* security announcement. Indeed, those who click on a *non-social* announcement are likely more intrinsically motivated to use security features, as they have no reason to click on the announcement but the desire to explore additional security features. Thus, the fact that there is *no* difference in adoption rate between these two groups may show that social announcements do in fact, raise lay people's motivation to use security features, because we would expect the more intrinsically motivated non-social announcement group to have a higher adoption rate than the more extrinsically motivated social announcement group. Future work can look at validating this hypothesis.

In summary, our results suggest that increasing the observability of security feature usage can substantially raise the exploration and adoption of available security features. However, our analysis from the first experiment remains primarily quantitative—we inferred an increase in awareness from people's actions to explore the interstitial we created. In a follow-up survey, we wanted to validate that people who clicked on our announcements actually had increased awareness of available security features, and we wanted to test whether the initial social information in the announcement had an effect on their subsequent exploration of available security features.

4. FOLLOWUP SURVEY

To more concretely measure whether our announcements increased people's awareness of available security features, we ran a second deployment of our best performing announcements from the initial experiment and collected survey responses.

Table 5. Number of survey responses per solicitation method (rows) and experimental group (columns).

	Holdout	Non-Social	Raw #	Some
Interstitial	0	498	226	254
Viewed Announcement	0	127	72	67
Holdout	788	322	214	246

4.1 Methodology

We re-ran a second campaign of our experiment with a separate set of n=50,000 people, randomly sampled among across users who used Facebook in English, logged in to Facebook at least once in the past month, and had at least 10 friends who used security features. People in our sample were shown one of three announcements mirroring the announcements in the previous experiment: the unambiguous "raw number" social condition, the ambiguous "some" social condition, and the non-social control condition—all exactly matching the corresponding condition from the initial experiment. All announcements were once again outfit with an "Improve Account Security" button that, when clicked, would navigate the clicker to an interstitial that explained the promoted security features, as well as allowed viewers to enable the same. The follow-up study ran between December 20th and December 22nd, 2013.

In this second campaign, we also asked people to complete a short survey with the following 3-point Likert-scale question: Facebook provides me with the necessary security settings to protect my account (i.e., the "Provides Features" statement). We decided to ask this question to test whether social information in the announcement influenced people's perceptions of the security features we promoted—namely, whether a viewer believed the features were sufficient to address their security concerns.

We had three methods to solicit survey responses. First, we surveyed people who fully navigated through the interstitial (i.e., the "interstitial" solicitation group). We separately sent the survey to people who saw an announcement but never clicked on it (i.e., the "viewed announcement" solicitation group), and also to a random sample of 80,000 people who used Facebook in English, logged in to Facebook at least once in the past month, and who never viewed any of our security announcements (i.e., the "holdout" solicitation group).

In total, we had 2814 responses to our survey. Table 5 shows a tabulation of the how many users per experimental condition and survey solicitation method.

4.2 Results

Table 6 shows the coefficients for a proportional-odds logistic regression [16] predicting the likelihood of an individual selecting a higher value of agreement with the "Provides Features" statement. Coefficients in Table 6, like those in Table 4, represent a change in "log-odds" that the user selected "neutral" over "disagree" or "agree" over "neutral" as a response to one of the questions. We included the viewer's experiment group as well how they were solicited to complete the survey as independent variables, and included the behavioral, demographic and social network descriptor variables described in Table 2 as controls.

Just as in the previous study, a positive coefficient implies that the log-odds ratio increases, or that the variable for the coefficient increases the likelihood that the user selected "neutral" over "disagree" or "agree" over "neutral". A negative coefficient implies the opposite. Furthermore, predictor variables were centered and scaled, such that each coefficient represents the expected change in log-odds that the user selected a higher value response given a one standard deviation increase in the predictor variable, holding all other numerical variables at their means and categorical variables at their baselines.

Table 6. Coefficients for the two proportional-odds logistic regressions predicting agreement with the trustworthy and protection statements. A full regression table including coefficients for control variables are provided in Appendix 1.

	Variable Name	Provides Features
†	Group: Non-Social	-0.08
†	Group: Raw #	-0.19
†	Group: Some	-0.16
Δ	Solicitation: Interstitial	1.04 *
Δ	Solicitation: Viewed Announcement	0.16

*† Baseline: Holdout; Δ Baseline: Holdout, * $p < 0.001$*

From Table 6, there appears to be no significant effect of viewing any of the security announcements on people's agreement with Facebook providing necessary security features, helping to explain why we saw the same adoption *rate* between those who saw social and non-social announcements in the previous experiment. Indeed, none of the coefficients for the "Group" variable were significant in either model.

On the other hand, people who actually clicked on the announcement and navigated through the security interstitial were significantly and substantially more likely to agree with the "Provides Features" statement (b=1.04, $p<0.001$) statement. Thus, while showing people security announcements with social information does not appear to directly affect people's sentiment towards Facebook's security tools, social announcements drive more people to the security interstitial and thus can at least *indirectly* raise their awareness or available security tools and their belief that those security tools are effective.

5. GENERAL DISCUSSION

In a nutshell, our results suggest that social proof is a promising approach to increase people's security sensitivity, but it is not a panacea. Indeed, showing people simple security announcements with social cues that increased the *observability* of security feature usage is a powerful, simple, and effective way to raise viewer's *awareness* of available security features, and thus, indirectly, their *adoption* of the same, and their *sentiment* towards the efficacy of the promoted features. However, the aggregated, impersonal social information we showed people only seemed to raise their interest in *exploring* security features—we did not find strong evidence that these announcements were more effective than a non-social announcement in increasing people's likelihood of actually adopting one of the promoted security features (though our results do not prove the opposite, either).

To summarize, through a large-scale experiment and survey on Facebook, we found that social announcements that increase the observability of security feature usage are more effective, than non-social announcements, at getting people to explore available security features—thus, the social prompts we tested seem to be very effective at raising security *awareness*. The positive effect of these social announcements on click-through rate is especially strong when viewers have many friends who use security features and when that information is rendered directly in the announcement, as with our "Raw #" announcement—a finding aligning with both the concept of *social proof* [6] and the *diffusion of innovations* [22]. This result suggests that the positive effect of these social cues will strengthen over time as more and more

people start using security features (and thus higher and higher numbers will be rendered in the announcement).

On the other hand, social announcements were no more effective than non-social announcements at getting people who clicked on the announcement to actually adopt a promoted security feature. Thus, used alone, the social announcements we tested appeared to be no better than a non-social announcement at raising users *motivation* to adopt the promoted security features. This finding holds true in both the short and long term, even through a number of widely publicized security vulnerabilities including Heartbleed [30] and the iOS SSL implementation bug [31] that could have been potential catalysts for security behavior change [9]. Nevertheless, as more people who saw a social announcement *clicked* on the announcement and explored the promoted security features, significantly more people who saw a social announcement *adopted* one of the promoted security features. There was, thus, an indirect increase in security feature uptake as a result of showing people a social announcement.

We also found evidence that social announcements *indirectly* appeared to increase viewers' belief that the security features they needed to secure their accounts were available. Indeed, people who viewed a social announcement were far more likely to click on the announcement and navigate through the resulting security interstitial, and people who navigated through the security interstitial were far more likely to agree that Facebook provided them with necessary security features.

Taken together, we have provided some experimental evidence that simple social proof cues *can* be used to raise peoples' security sensitivity—specifically, their awareness of available security tools. Furthermore, using these simple social cues may have the additional indirect benefits of raising security feature adoption and people's sentiment towards the promoted features, as well. Care should be taken, however, to sparingly surface these announcements so that people do not get desensitized to them. For example, to maximize the efficacy of a campaign to raise security sensitivity, social announcements should probably only be shown once every few months to people who already have many friends who use the security features promoted in a campaign.

Importantly, our findings do *not* suggest that social cues are ineffective at raising people's motivation to use security features. Rather, our null result at raising motivation was likely an artifact of the fact that the prompts we tested were aggregated, out of context and not very informative. For example, showing someone an announcement that 100 of her friends use security features does not inform her *why* those friends use security features, *which* security features are being used (or for what purpose), *who* among her friends are using those security features, and whether or not her friends would actually *recommend* using those features. In other words, our absence of results in raising motivation may be due to lack of compensation for an invalid *context*—i.e., asking people to consider extra security features when they are not really thinking about security. Accordingly, motivation to adopt security features might be best driven by a paired approach of security threat detection followed by a timely delivery of a security announcement with social cues.

5.1 Limitations and Future Work

As with any research endeavor, there are limitations to the present work that open up exciting avenues for future work. Below we list two of the most salient:

First, we did not test the many forms of social influence apart from *social proof*. The *reciprocity* principle of social influence, for example, suggests that people would be more likely to follow the advice of a friend who previously did them a favor [5], though examining this principle would require showing people more specific information about who amongst their friends had enabled a security feature. The *liking* principle suggests that people are more likely to follow advice from others who are similar to them—for example, a social cue that shows a female user that 100 of her *female* friends use security features may be more effective at driving her interest than a cue that just shows her that 100 of her friends use security features [6]. There are, therefore, many avenues for future work to explore the efficacy of other forms of social influence in raising security sensitivity, though preserving the privacy of individual security feature users is also important.

Second, our current methodology only tested the efficacy of simple social proof cues in raising peoples' *awareness* of security features and *motivation* to use those features. We did not test the efficacy of these cues in raising the third component of the security sensitivity—the *knowledge* of how and when to use a security feature. Future work can look at constructing social narratives to help viewers understand how and when to use promoted security features.

6. CONCLUSION

The need for increased security sensitivity remains one of the largest outstanding problems in computer security. In this work, we explored if it is possible to heighten security sensitivity by showing people *social proof* that their friends use security features. In a study of responses to security information, we found that showing people security announcements that showed them the number of their friends who use security features was significantly more effective at getting them to explore those security features than a non-social announcement. In turn, as more people who viewed social announcements explored the promoted security features, more people who viewed social announcements also adopted one of the promoted features. Furthermore, this positive effect strengthened for people who had more friends who already used security features, suggesting that these social announcements will only get more potent over time as more people start adopting security features. On the other hand, among those people who clicked on either announcement, the social announcements were no more effective than the non-social announcement at getting people to adopt the promoted features. Finally, in a follow-up survey, we also found that people who saw social announcements were also indirectly more likely to agree that they had available the security features necessary to protect their Facebook accounts. In all, we have provided among the first empirical evidence that social proof can be used to heighten security sensitivity, and believe that our work provides a solid foundation for further exploring the use of simple social cues to increase the awareness of, motivation to use, and knowledge of how to use security tools.

7. Acknowledgements

This work was generously supported, in part, by NSF Award #1347186 and the NDSEG Fellowship. We would also like to thank Melissa Luu-Van, Tam Nguyen and Peter Flemming of Facebook for assisting with ideation and execution of this research. We would also like to thank Tiffany Hyun-Jin Kim for helping with the refinement of this and related projects.

8. REFERENCES

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

[2] 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.

[3] Bond, R.M., Fariss, C.J., Jones, J.J., et al. A 61-million-person experiment in social influence and political mobilization. *Nature 489*, 7415 (2012), 295–8.

[4] Burke, M., Marlow, C., and Lento, T. Feed me: Motivating Newcomer Contribution in Social Network Sites. *Proc. CHI '09*, ACM Press (2009), 945–954.

[5] Cialdini, R.B., Vincent, J.E., Lewis, S.K., Catalan, J., Wheeler, D., and Darby, B.L. Reciprocal Concessions Procedure for Inducing Compliance: The Door-in-the-Face Technique. *JPSP 31*, 2 (1975), 206–215.

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

[7] Cohen, J. *Statistical Power Analysis for The Behavioral Sciences*. England: Lawrence Erlbaum Associates, Inc., Hillsdale, NJ, 1977.

[8] Das, S., Hayashi, E., and Hong, J. Exploring Capturable Everyday Memory for Autobiographical Authentication. *Proc. UbiComp'13*, (2013).

[9] Das, S., Kim, H.J., Dabbish, L.A., and Hong, J.I. The Effect of Social Influence on Security Sensitivity. *Proc. SOUPS'14*, (2014).

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

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

[12] 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.

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

[14] Egelman, S., Sotirakopoulos, A., Muslukhov, I., Beznosov, K., and Herley, C. Does my password go up to eleven? *Proc. CHI '13*, ACM Press (2013), 2379–2388.

[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] Hardin, J.W. and Hilbe, J.M. *Generalized Linear Models and Extensions*. Stata Press, College Station, Texas, 2007.

[17] Johnson, R.C. Cyber security solutions underused. *EE Times*, 2013. http://www.eetimes.com/author.asp?section_id=36&doc_id=1287251&page_number=1.

[18] Kramer, A.D.I. The spread of emotion via facebook. *Proc. CHI '12*, ACM Press (2012), 767–770.

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

[20] Moore, H. and Roberts, D. AP Twitter hack causes panic on Wall Street and sends Dow plunging. *The Guardian*, 2013. http://www.theguardian.com/business/2013/apr/23/ap-tweet-hack-wall-street-freefall.

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

[22] Rogers, E.M. *Diffusion of Innovations*. Free Press, 2003.

[23] 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).

[24] 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.

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

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

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

[28] Wyler, G. AP Twitter Hacked, Claims Barack Obama Injured In White House Explosions. *Business Insider*, 2013. http://www.businessinsider.com/ap-hacked-obama-injured-white-house-explosions-2013-4.

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

[30] Heartbleed. *Wikipedia*, 2014. http://en.wikipedia.org/wiki/Heartbleed.

[31] Vulnerabilities Summary for CVE-2014-1266. *National Vulnerabilities Database*, 2014.

Appendix A

Table A1. Coefficients for the two proportional-odds logistic regressions predicting agreement with the trustworthy and protection statements. Bolded coefficients are of interest.

	Variable Name	Provides Features
†	**Group: Non-Social**	**-0.08**
†	**Group: Raw #**	**-0.19**
†	**Group: Some**	**-0.16**
Δ	**Solicitation: Interstitial**	**1.04 ***
Δ	**Solicitation: Viewed Announcement**	**0.16**
	Feature-using friends	-0.13
	Age	0.04
	Gender: Male	0.15
	Account length	0.20 *
	Friend count	0.25 *
	Mean friend age	-0.14
	Friend age entropy	0.07
	Percent male	0.03
	Mean friends days since confirmed	-0.57 *
	Friend country entropy	0.005
	Mean number of friends of friends	-0.08
	Posts created	0.02
	Posts deleted	-0.07
	Comments created	-0.06
	Comments deleted	0.05
	Friends added	0.05
	Friends removed	-0.05
	Photos added	-0.001

*† Baseline: Holdout; Δ Baseline: Holdout, * $p < 0.05$*

Table A2. Coefficients for the three regressions predicting clicks (CL), feature adoptions up to a week after the experiment (A-ST), and feature adoption up to 5 months after the experiment (A-LT). Bolded coefficients are of interest (non-control).

	Variable Name	CL	A-ST	A-LT
†	**Group: At Least #**	**0.29 ***	**-0.07 ***	**-0.13**
†	**Group: At Least %**	**0.21 ***	**-0.12**	**-0.06**
†	**Group: Only #**	**0.26 ***	**-0.16**	**-0.09**
†	**Group: Only %**	**0.19 ***	**-0.12**	**-0.05**
†	**Group: Raw #**	**0.36 ***	**-0.01**	**-0.001**
†	**Group: Raw %**	**0.17 ***	**-0.15**	**-0.06**
†	**Group: Some**	**0.35 ***	**-0.18**	**-0.03**
	Feature-using friends	**0.09 ***	**0.17 ***	**0.20 ***
	Intercept	-2.16 *	-5.23 *	-2.62 *
	Age	-0.01	-0.19 *	-0.18 *
	Gender: Male	-0.03	-0.06	-0.13 *
	Account length	0.11 *	0.03	0.03
	Friend count	-0.16 *	-0.06	-0.15 *
	Mean friend age	0.14 *	-0.16	-0.24 *
	Friend age entropy	0.03	0.28 *	0.26 *
	Percent male	0.02	0.08	0.13 *
	Mean friends days since confirmed	0.007	0.003	-0.08 *
	Friend country entropy	0.04 *	0.04	0.03
	Mean number of friends of friends	-0.04	-0.09	-0.09 *
	Posts created	0.05	0.02	-0.02
	Posts deleted	-0.008	0.02	-0.002
	Comments created	0.09 *	0.07 *	0.10 *
	Comments deleted	0.07	-0.13	-0.01
	Friends added	-0.003	0.004	0.02
	Friends removed	-0.004	0.01	0.02
	Photos added	0.03 *	0.004	0.03
	Clicked on Announcement	N/A	4.38 *	1.94 *

*† Baseline: Control, * $p < 0.05$*