The Role of Social Influence In Security Feature Adoption

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
Social influence is key in technology adoption, but its role in security-feature adoption is unique and remains unclear. Here, we analyzed how three Facebook security features—Login Approvals, Login Notifications, and Trusted Contacts—diffused through the social networks of 1.5 million people. Our results suggest that social influence affects one's likelihood to adopt a security feature, but its effect varies based on the observability of the feature, the current feature adoption rate among a potential adopter's friends, and the number of distinct social circles from which those feature-adopting friends originate. Curiously, there may be a threshold higher than which having more security-feature adopting friends predicts for higher adoption likelihood, but below which having more feature-adopting friends predicts for lower adoption likelihood. Furthermore, the magnitude of this threshold is modulated by the attributes of a feature—features that are more noticeable (Login Approvals, Trusted Contacts) have lower thresholds.

Author Keywords
Social Cybersecurity; Social Influence; Facebook; Security; Security Adoption

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

INTRODUCTION

In early 2013, the Associated Press's Twitter account was compromised through a password phishing scheme, and erroneously tweeted that President Obama was injured in a bombing [30]. In response, stock prices plummeted [19], adversely affecting thousands. Moreover, this incident could have been easily prevented with the use of two-factor authentication—a security feature, available at that time, that requires entry of a random code generated on one's phone in addition to a password when authenticating [16].

CSCW'15, March 14 - 18 2015, Vancouver, BC, Canada

ACM 978-1-4503-2922-4/15/03...$15.00
http://dx.doi.org/10.1145/2675133.2675225

This incident is just one example of how the underutilization of available security features remains a large, outstanding problem. Indeed, in our age of increasing connectivity, it is critically important for widespread awareness of and appropriate use of security features.

Recent work suggests that one promising approach to widespread security feature awareness and uptake is by understanding the social diffusion of security feature adoption, or how the feature propagates from person-to-person through a social network. Indeed, a recent retrospective interview study [11] outlines that security behavior changes are often the result of a social diffusion process. Many participants of this study reported adopting security features solely because their friends also used those features—in other words, their behavior was driven by *social proof* [10] (e.g., other people like me use this feature, so I should too). However, there is also evidence that social processes can work *against* security related behavior change (e.g., only paranoid people use complex security features, and I'm not paranoid) [11, 13].

Thus, while the idea that social influence affects security feature adoption is interesting, what little we know suggests that *social influence may uniquely affect security technology* because security features are preventative, intrusive, and can be associated with paranoia [11,13]. Indeed, if only 'experts' or people who are perceived as paranoid initially use a security feature, lay people might develop an illusory correlation [9] between using a security feature and paranoia that makes them avoid using the feature. Conversely, *social proof* can also be an effective motivator for security driven behavior change [11], especially when people can *observe* others like them behaving securely. In other words, social influence can be both a helpful and harmful force in security-feature adoption, but we do not yet fully understand the parameters under which it is helpful or harmful.

To address this gap in the literature, we analyzed whether and how three Facebook security features—Login Approvals, Login Notifications and Trusted Contacts—diffused through the social networks of 1.5 million people who use Facebook. Our results confirm that social influence does indeed play a role in security-feature adoption, but that the direction and strength of its effect seems to be moderated by the overall adoption of the feature among a potential adopter's friends, the number of distinct social circles from which those feature-adopting friends originate, and the individual attributes of the security feature.

BACKGROUND

Social Diffusion and Friend Diversity

Earlier work suggests that exposure to novel information on social networking sites increases information diffusion through social channels [3], but that these diffusion chains are most effective when the seed information is shared by many different sources [3,27], especially when the information is intended to enact behavior change [6,7]. Ugander and colleagues [29] extended this result, finding that people who were invited to join Facebook through e-mail recommendations from their friends were more likely to join if the recommenders were from distinct social contexts—i.e., receiving an invitation from a school friend and a family member was more convincing than receiving invitations from two different family members. Romero and colleagues [22] found that the "persistence" of the information being spread—or, the marginal likelihood that content will be re-shared after one more exposure—is also important in determining whether content will be diffused. Specifically, controversial topics—like information about security, say—require repeated exposure from many sources before they are diffused.

H1: *People with exposure to feature-adopting friends from many distinct social contexts will be more likely to use that feature than others with exposure to the same number of feature-adopting friends from fewer distinct social contexts.*

Social Diffusion and Observability

It is well established that not all behavior diffuses equally [6,7], and the adoption of technology is no different. Thus, efforts have been made to model the factors that influence the adoption of technology. Rogers [21], in his seminal work on the diffusion of innovations, argued that new technology gets widely adopted through a process by which it is communicated through members of a social network. Rogers argues that primarily subjective perceptions 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 features, that prevent undesirable outcomes from happening—typically have low adoption rates, in part because of their low *observability*, or the invisibility of their use and benefits. More recently, Das and colleagues [11] confirmed that the *observability* of security tools and behaviors was a key factor in driving the adoption of security tools. In fact, they found that of all social catalysts for behavior change, observing others use security tools was the most prevalent.

H2: *More observable security features will more effectively diffuse through social channels than less observable security features.*

Diffusing Behavior via Social Influence

Social influence, of course, affects not only our knowledge but our behavior as well. Christakis and Fowler [12] argue that our social connections can influence our health and weight, whereas Bond and colleagues [5] found that awareness of friends' voting activity was significantly effective at getting people who use Facebook to vote in the 2010 U.S. congressional elections. Kramer [17] found that emotions are contagious on Facebook as well: friends of those who shared emotional content were themselves more likely to share similar emotions. Furthermore, in the context of security, prior work suggests that we do learn about security from each other [20], though communication about security remains relatively rare [11].

Others have looked at the psychological mechanisms underlying social influence. Indeed, most of the previously mentioned work addresses the concept of *social proof*, or our tendency to look to others for examples of what to do when we are uncertain [10]. Milgram, Bickman, and Berkowitz [18] 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. Other work has shown how social interventions can be powerfully effective at eliciting behavior change: for example, at reducing household energy consumption by informing people about their neighbors' reduced energy consumption [25], reducing hotel guests' wasteful use of towels by telling them previous patrons chose to be less wasteful [14], and even in eliminating young children's phobia of dogs by showing them film clips of other children playing with dogs [4].

Taken together, it appears that the social diffusion of behavior and technology adoption is often driven by *social proof* [10]—or evidence of what to use and how to behave based on the actions of others. Furthermore, this social proof becomes more compelling as one finds *more* examples of others acting similarly, especially examples of others similar to oneself [8,10,18].

The Unique Effect of Social Influence on Security

Prior work in the psychology and application of social influence implies, thus, that if many of one's friends and acquaintances use a security feature, one should be more likely to use that security feature herself. Yet, we see some counter examples of this implication in the usable security literature. Indeed, Gaw and colleagues [13] found that many non-experts perceived others who used e-mail encryption as "paranoid", a perception that inhibited their own use of e-mail encryption. Das and colleagues [11] found that their non-expert participants were similarly aversive towards using security tools, and spoke of their security-expert friends as being "nutty" or going "above and beyond".

Thus, it appears that social proof does not always have the expected effect on security feature adoption. We believe, in fact, that *the use of security features may be uniquely affected by social proof* given that security feature usage is often invisible, rarely communicated, and generally undesired [15,24]. Indeed, prior work in usable privacy and security suggests that many security features remain unused because stringent security measures are often antagonistic towards the specific goal of the end user at any given

Demographic Variables	
Age	Age of the individual.
Gender	Self-reported gender: male or female.
Friend count	Count of the individuals number of friends with Facebook accounts.
Account length	Days that have passed since the individual activated his account.
Days active in last 30	Days the individual was active on Facebook in the past 30 days.
Social Network Variables	
Mean friend age	Average age of the individual's Facebook friends.
Friend age entropy	Shannon entropy of the individual's Facebook friends' ages.
Percent male friends	Percentage of the individual friends that are male.
Mean friends' account length	Average number of days an individual's Facebook friends have used Facebook.
Friend country entropy	Shannon entropy of countries from which the user has friends.
Mean number of friends among friends	Average number of Facebook friends among an individual's Facebook 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.
Likes	Number of likes given.
Friends Added	Number of friends added.
Friends Removed	Number of friends removed.
Photos Added	Number of photos added.
Videos Added	Number of videos added.
Social Proof Variables	
Percent of friends who use Login Approvals	Percent of friends who use the Login Approvals security feature.
Percent of friends who use Login Notifications	Percent of friends who use the Login Notifications security feature.
Percent of friends who use Trusted Contacts	Percent of friends who use the Trusted Contacts security feature.
Number of diverse social contexts	Number of social contexts from which friends who use security features originate.

Table 1. Collected feature descriptions and distributions. These variables were all collected or computed at an individual level.

moment [24]. For example, while a user might want to check her e-mail, a complex password that usually requires three attempts to get right *prevents* her from checking her e-mail. Thus, people often reject security features when they expect or experience them to be weighty [1].

Consequently, typically only people who are especially dedicated to protecting their information use interruptive security features, and we know from prior work that non-experts may perceive these early adopters as "paranoid" [13]. More formally, because early adopters of security features are likely to be perceived by others as behaviorally *different* (e.g., either paranoid, or in possession of expert knowledge), non-experts may perceive an illusory correlation [9], or an exaggerated relationship, between security feature usage and this behavioral difference. In turn, as non-experts consider themselves different from those who use security features, they may reject the use of security features. Moreover, this illusory correlation should only *strengthen* as more of these security-enthusiast early adopters use the feature because of the "availability heuristic"—a mental shortcut that biases people's judgments towards what is more frequently recalled [28].

The upshot is that the subjective perceptions of a security feature that propagates through social channels may be tainted into working *against* its adoption, at least until enough of a potential adopter's behaviorally similar friends start using the feature so that its use becomes normative.

In other words, there may be a non-linear relationship between one's exposure to feature-adopting friends and one's likelihood to adopt a security feature. Specifically, if a potential adopter is only exposed to *few*, early-adopter friends who use a security feature, it is possible that he might find social proof that a security feature should *not* be used (because of an illusory correlation), and the strength of this *negative* social proof should increase with the number of these feature-adopting friends (because of the availability heuristic). On the other hand, once a potential adopter is exposed to *many* feature-adopting friends, especially those that are similar to himself, he might find social proof that a security feature *should* be used (because of the positive effects of homophilous networks on technology adoption [8]), and the strength of this *positive* social proof should increase with the number of his feature-adopting friends.

H3: *When a potential adopter is exposed to many feature-adopting friends, he will be more likely to adopt a security feature than those with fewer feature-adopting friends.*

H4: *When a potential adopter is exposed to few feature-adopting friends, he will be less likely to adopt a security feature than those with even fewer feature-adopting friends.*

METHODOLOGY

To test our hypotheses, we monitored security feature adoptions for the following three Facebook security features: (1) **Login Approvals—**A feature that requires adopters to enter a separate code, usually generated on or sent to the adopter's smartphone, in addition to their

password when they attempt to authenticate; (2) **Login Notifications**—A feature that notifies adopters, via e-mail or SMS, when their account is accessed from previously unseen browsers and devices; and, (3) **Trusted Contacts**—A feature that allows an adopter to specify three to five friends who can verify her identity if she forgot her password and cannot access her e-mail. We investigated multiple security features to avoid drawing conclusions specific to any one feature, especially because the individual attributes of a security feature may play a role in its diffusion [21]. Furthermore, we chose these three features because of their diversity and colocation within the "security settings" page on Facebook.

For 12 days in late 2013, we collected data from a random subset of people who use Facebook and newly adopted one of the aforementioned security features: Login Approvals, Login Notifications, or Trusted Contacts. In total, we collected data from n=250,000 people per feature (750,000 adopters overall)—the positive examples of feature adopters in our dataset. Then, for each day and feature, we also obtained a random sample of an equal number of people who had not adopted that feature up to that day—negative examples of feature adopters. In total, we had *n*=1,500,000 people across all twelve days, three features (Login Approvals, Login Notifications, Trusted Contacts), and two feature usage states (i.e., uses or doesn't use).

For all people in our sample, we also collected a set of variables that we believed could have affected one's decision to adopt a security feature. These variables fell under four categories: *demographic variables* that described individual characteristics such as age and gender; *behavioral variables* that described activity on Facebook, such as posts shared and deleted; *network variables* that described one's social network, such as friends' average age and gender diversity; and, *social proof variables* that described how many and which of a person's friends had adopted any of the aforementioned security features up to the day during which the data was collected. In Table 1, we provide a full list of variables included in our analysis. All data was de-identified prior to our analysis.

We selected people who newly adopted security features because security feature adoptions were not time-stamped in our data, so it would be otherwise impossible to know who, between two people, adopted a security feature first. For someone who newly adopted a security feature on a given day, however, we knew that all friends of their friends who used that feature adopted it before that day.

Notably, we could not measure *how* security feature adoptions diffused—i.e., we did not alter the observability of security feature usage and initiation. Rather, we simply control for other factors that also affect security feature adoption, such that we can compare the feature adoption rate of two sub-populations that differ primarily in their exposure to friends who have adopted a security feature. We do not believe this limitation to be stifling—understanding the channels through which social diffusion occurs is separate from our goal of understanding its ultimate effect on security feature adoption.

Finally, all data collection complied with Facebook's terms of use and data use policy and was performed in aggregate so that we were not privy to any individual's information. Furthermore, as our data was observational, we believe our analysis constituted minimal risk to those in our sample.

REGRESSION ANALYSIS

To test **H1**—that people with security-feature adopting friends from many distinct social circles should be more likely to adopt a security feature than those with the same number of feature-adopting friends from few distinct social circles—we estimated a logistic regression model for each security feature. These regressions modeled the strength of the relationship between a person's likelihood to adopt a security feature and the number of distinct social contexts from which his feature-adopting friends originated. Note that we define a "distinct social context" as a distinct connected component in one's friend graph, following similar definitions used in prior work [29].

As linear regression analysis assumes independence in the response variable (in our case, whether or not someone in our sample adopted a security feature), we only included a balanced subset of our full sample into the regressions after eliminating people in our sample who happened to be Facebook friends with one another. This reduced sample consisted of n=65,000 positive and negative examples of feature adopters for each of our three features, resulting in n=130,000 people for each regression, all of whom were not friends with one another.

In running these regressions, we controlled for the demographic, social network and behavioral variables described in Table 1. In addition, we also controlled for the *number* of one's feature-adopting friends, so that the coefficient for the *number of distinct social contexts* variable can be interpreted after controlling for a potential adopter's number of feature using friends.

Results

The coefficient for the *number of distinct social contexts* variable for each logistic regression is shown in Figure 1, while the full regression table is shown in Appendix A. These coefficients represent a change in "log-odds", or $ln\frac{P}{1-P}$, where P represents the probability that an individual adopted the security feature. A positive coefficient implies that the log-odds ratio increases, or that an increase in the variable increases the likelihood that a person adopts the feature, *P*. A negative coefficient implies the opposite. Furthermore, each variable was centered and scaled, such that its coefficient represents the expected change in log-odds that a person uses a feature given a one standard deviation increase in the predictor variable, holding all other numerical variables at their means and categorical variables at their baselines. Additionally, larger absolute

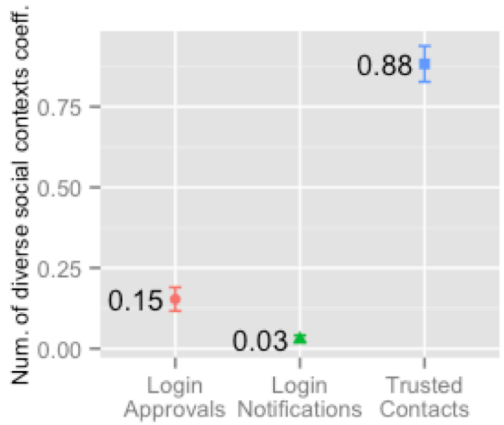

Figure 1. Coefficients for the three logistic regressions relating the number of diverse social contexts variable to use of each security feature, with 95% confidence intervals. All coefficients significant, p < 2e-16.

coefficient values imply a stronger relationship between the independent and dependent variables.

Thus, from Figure 1, we can see that the *number of diverse social contexts* variable positively correlated with the adoption of *every* security feature (b_{LA}=+0.15, p<2e-16; b_{LN}=+0.03, p<2e-16; b_{TC}=+0.88, p<2e-16). This finding offers support for **H1**—people with friends from more diverse social contexts (e.g., high school friends, college friends, family) who use a security feature should be more likely to adopt that feature themselves than those with feature-adopting friends from fewer distinct social contexts. In other words, it is not just the *number* of one's friends who use a security feature that matters in influencing one to adopt the feature himself; these friends should be independent of one another for the effect to be strongest.

In addition, the discrepancy of effect size across features offers some support for **H2**—that more observable security features will be more effectively diffused through social channels. Indeed, the absolute effect size of the *number of diverse social contexts* variable is largest, by far, for Trusted Contacts (the most observable feature, b_{TC}=+0.88), then for Login Approvals (the next most observable, b_{LA}=+0.15) and finally lowest for Login Notifications (the least observable feature, b_{LN}=+0.03).

Indeed, Login Notifications are private messages that are not very observable, and are thus difficult to passively diffuse via social channels. Thus, while having many different friends use Login Notifications may make for a more convincing case for a potential adopter to use the feature, the case is unlikely to be made. Login Approvals are more observable than Login Notifications in that friends who are collocated with an adopter will see the additional authentication step it requires, which in turn may passively provide these friends with social proof to use Login Approvals [11]. This modest increase in observability appears to correlate with a modest increase in the effect size of the *number of diverse social contexts* variable. Finally, the Trusted Contacts feature sends out a notification to each of one's friends who was specified as a Trusted Contact, thus substantially increasing its visibility in a direct way and, in turn, correlating with a substantial increase in effect size. It is also possible that the social nature of the feature—in enlisting friends to help recover one's account—lends itself to amplified social diffusion.

In summary, our regression analysis provides us with support for **H1** and limited support for **H2**, but we have yet to test **H3** and **H4**—that the feature-adoption rate of one's current set of friends will moderate whether the effect of social proof will be positive or negative on one's own likelihood to adopt that security feature. Unfortunately, linear regression analysis is limited in that it does not consider this form of non-linearity in the relationship between predictor and response. Furthermore, regression analysis confounds homophily-based diffusion with social-influence based diffusion [2,26]. In other words, because similar people cluster together as friends, we cannot tell if co-adoption of a feature is due to one friend influencing another or because both friends share an interest. Thus, to test **H3** and **H4**, we ran an adapted version of *matched propensity sampling* [2].

MATCHED PROPENSITY SAMPLING ANALYSIS

Matched propensity sampling is a form of causal inference that helps us differentiate feature adoption due to homophily from feature adoption due to social influence. It distinguishes between homophily and social influence by comparing the feature adoption rates of two sets of people who are *equally likely* to have a fixed proportion of friends who have adopted a security feature, where one set *actually does* have the fixed proportion of friends who have adopted this feature and the other set does not. People in the former set are "exposed" to their feature-adopting friends at this fixed rate, while those in the latter set are "unexposed."

Exposed and unexposed individuals are matched, in pairs, based on a "propensity score" computed from a set of covariates Z that are theorized to represent homophily-based diffusion [23]. We used a logistic regression to calculate the propensity score as suggested by prior work [2], and the covariates included in the model were the demographic, behavioral, and social network variables listed in Table 1. As we are not concerned about estimating exact coefficients and their variances with the logistic regressions in this analysis, we are able to break the independence assumption and include the full set of 1.5 million users in our sample.

Unfortunately, as we could not capture the security expertise of those in our sample, there remains some form of "latent homophily" for which we do not control. However, the demographic, behavioral, and social network variables for which we control likely predict security expertise, so we believe this limitation to be minimal.

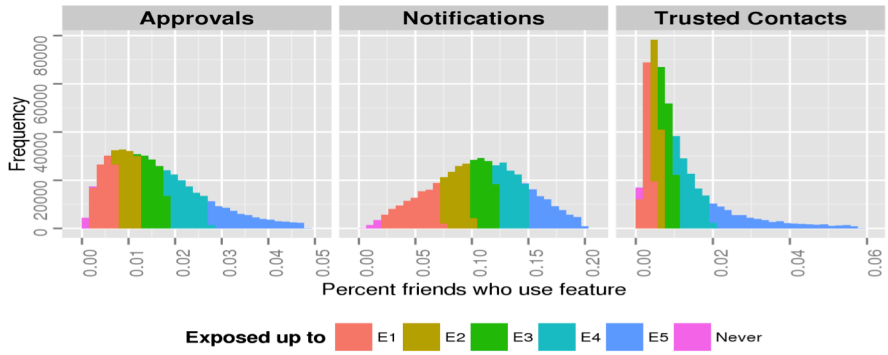

Figure 2. Histogram of percent of friends who use login approvals (left), login notifications (middle) and trusted contacts (right). Colors represent up to what exposed conditions users with x% of feature-adopting friends would be considered "exposed" in the matched propensity sampling analysis.

	Percentile	Approvals	Notifications	Trusted Contacts
E1	1st	0.2%	2.0%	0.1%
E2	21st	0.8%	7.3%	0.4%
E3	41st	1.3%	10.0%	0.7%
E4	61st	1.8%	12.3%	1.1%
E5	81st	2.7%	15.1%	2.0%

Table 2. Exposed condition prerequisites for each security feature. For example, if a user is "exposed" at E3 for login approvals, at least 1.3% of her friends must have adopted login approvals at the time of data collection.

By matching exposed and unexposed individuals who have the same likelihood of being exposed, we can take the difference in feature adoption rates between the exposed and the unexposed as evidence of the effect of social influence. Indeed, after the propensity matching process, the only theoretical difference between these two sets of people are that the exposed set has a certain proportion of friends who use a security feature and those in the unexposed set do not. If social influence has no effect, we should see the *same* rate of adoption for the exposed and unexposed, whereas if social influence has a positive or negative effect, we should see that exposed individuals adopt the feature at a *higher* or *lower* rate, respectively.

We specified five empirical exposure conditions for each security feature—Login Approvals, Login Notifications, and Trusted Contacts—with each exposure condition representing whether or not the user was at least in the 1st percentile, the 21st percentile, the 41st percentile, the 61st percentile, or the 81st percentile in the *percent of friends who use feature* variable, or the total percentage of their friends who used a security feature at the day of data collection. Notably, a potential adopter could count as "exposed" at some levels but not others.

Figure 2 depicts the values of the *percent of friends who use feature* variable that qualified for "exposure" under E1 through E5, with actualized values for these conditions shown in Table 2. Concretely, an individual is exposed in *E1* for Login Approvals if at least 0.2% of her friends adopted the feature, because that puts her at least at the 1st percentile of people whose friends have adopted the feature. Likewise, she is exposed in *E5* for Login Approvals if at least 2.7% of her friends adopted the feature.

We chose five exposure conditions uniformly spaced across the distribution of the *percent of friends who use feature* variable to get a detailed map of the relationship between exposure to friends who have adopted a security feature and one's own likelihood to adopt that feature at different levels of exposure. This map will help us evaluate both **H3** and **H4**—specifically, **H3** predicts a higher adoption rate for the *exposed* relative to the unexposed at high exposure conditions because of positive social proof, whereas **H4** predicts a higher adoption rate for the *unexposed* at low exposure conditions because of negative social proof.

Results

Figure 3 shows the rate of feature adoption for exposed and unexposed individuals for all three features across all five exposures. In interpreting the results of the matched propensity analysis in Figure 3, we note the following: (i) If social influence has *any* effect on the adoption of a security feature at a particular level of exposure, we should see a significant difference in the adoption rates of exposed and unexposed individuals; (ii) If social influence has a *positive* effect on the adoption of a security feature at a particular level of exposure, then we should see that exposed individuals have a significantly *higher* adoption rate than the unexposed; and, (iii) If social influence has a *negative* effect on the adoption of a security feature at a particular level of exposure, we should see that exposed individuals have a significantly *lower* adoption rate than the unexposed.

First, as we show in Table 4, *all* of the differences in adoption rate between the exposed and unexposed were significant, suggesting that irrespective of the security feature and level of exposure to friends who use that feature, social influence appears to have a significant effect on one's likelihood to adopt a security feature. This finding empirically supports some recent, smaller-scale qualitative results that surfaced social influence as a key factor in the adoption of security features [11].

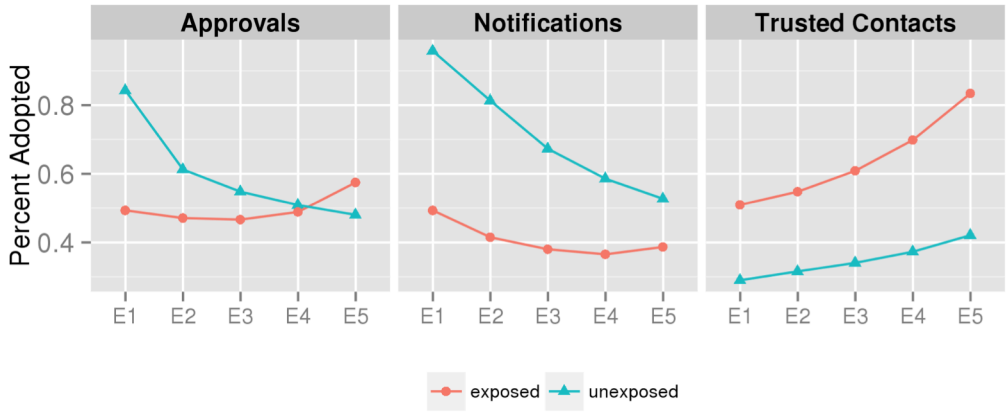

Figure 3. Feature adoption rates, plotted for each security feature for each exposure condition, for both exposed and unexposed individuals. Exposed feature adoption rates are plotted as red circles, and unexposed feature adoption rates are plotted as blue triangles. Every difference between the exposed and unexposed for all features was statistically significant.

	Approvals		Notifications		Trusted Contacts	
	N	χ^2, df=1	N	χ^2, df=1	N	χ^2, df=1
E1	5852	1553	25061	13743	4995	491
E2	122765	4994	518907	172603	105156	11742
E3	240061	3104	1014159	174619	205541	29775
E4	228905	140	963824	93771	196397	42022
E5	111092	1976	468147	18828	95393	34665

Table 3. Chi square significance tests for the difference in adoption rate between exposed and unexposed individuals across all exposure conditions and all security features. All differences significant, p < 2e-16.

For Login Notifications, we see that people who are *exposed* to a certain proportion of feature-using friends appear to be *less* likely to adopt those features than people who are *unexposed* for all levels of exposure we tested. Thus, in our sample, even people with a higher-than-average proportion of feature-adopting friends (i.e., those exposed at E4-E5 who are at least at the 61st percentile) were themselves *less* likely to use Login Notifications than people who had fewer friends who used those features. It appears, therefore, that exposure to friends who use Login Notifications *stifles* the adoption of Login Notifications, a finding that supports **H4**—that social influence will have a negative effect on feature adoption at low exposure levels—but conflicts with **H3**—that social influence will have a positive effect on feature adoption at high exposure levels.

We see just the opposite trend for Trusted Contacts, however: even at E1, the lowest level of exposure, exposed individuals are significantly *more* likely to adopt Trusted Contacts than the unexposed. In other words, it seems that *any* exposure to friends who use Trusted Contacts subsantially increases one's own likelihood to adopt that feature, a finding that supports **H3** but contradicts **H4**.

Finally, for Login Approvals, we see exactly the nuanced, thresholded relationship we predicted. At lower levels of exposure, unexposed individuals are more likely than the exposed to adopt the feature, but at the highest level of exposure, exposed individuals are more likely to adopt the feature—a finding that supports *both* **H3** and **H4**.

Thus, we have three security features for which adoption is significantly affected by social influence, but for which the effect of social influence appears to manifest differently. For Login Notifications, it appears that social influence is a categorically negative force on its adoption, for Trusted Contacts it is a categorically positive force, and for Login Approvals, the direction of its effect is based on a threshold level of exposure a potential adopter has to friends who already use that feature. What could explain the differences in the effect of social influence across these features?

Theoretical vs. Empirical Exposure Threshold

The matched propensity sampling analysis only reflects the effect of social influence on the adoption of a feature at its rate of adoption *at the time of data collection*. Indeed, our exposure conditions were based on an *empirical* division of the *percent of friends who use feature* variable; therefore, it is possible that there is a theoretical exposure *greater* than E5 where social influence could have a positive effect on the adoption of Login Notifications. Indeed, for Login Notifications, exposure at E5—at which about 15% of one's friends use Login Notifications—may not yet be at the threshold where **H3** predicts social influence should have a positive effect on its adoption.

To test this possibility, we must observe how the *adoption rate difference* between the exposed and unexposed varies across exposure conditions. We plot these differences in Figure 4, by subtracting the unexposed adoption rate from the exposed adoption rate. From this plot, we can understand the *marginal* effect of social influence on adoption at higher exposure conditions. In interpreting Figure 4, we note the following: (i) If unexposed individuals are more likely than the exposed to adopt a feature at a certain level of exposure, then the value of the difference will be negative, whereas it will be positive if exposed individuals are more likely to adopt the feature than the unexposed; and, (ii) If the value of the difference increases (moves up) at higher exposure conditions, then the marginal effect of having more friends who use a

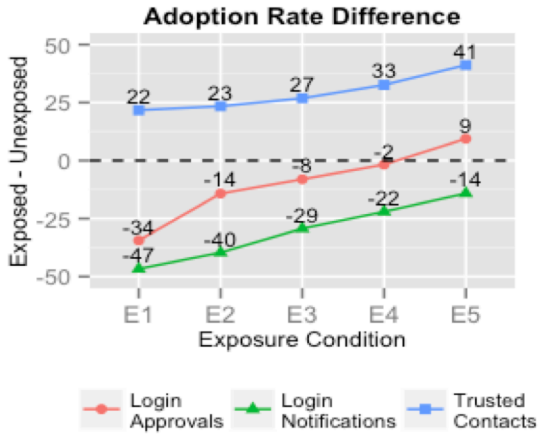

Figure 4. Differences in adoption rate between the exposed and unexposed for all three features across all exposure conditions. Values above the dashed horizontal line signify that those who were exposed had a higher adoption rate than the unexposed. All differences significant, p < 2e-16 (Table 4).

security feature on that feature's adoption is positive, whereas if the value of the difference decreases (moves down), then the marginal effect is negative.

From Figure 4, we see that the value of the difference between exposed and unexposed adoption rates *increases* (moves up) constantly, for all three features, from E1 to E5. For Login Notifications and Login Approvals, the initial adoption rate advantage of unexposed individuals gradually diminishes at higher levels of exposure. In fact, the advantage is ultimately in favor of exposed individuals for Login Approvals at E5, when the difference shifts from negative to positive. For Trusted Contacts, the advantage starts with exposed individuals and simply gets larger at higher levels of exposure. Thus, at higher levels of exposure, the likelihood for exposed individuals to adopt any of the security features grows at a rate faster than the unexposed. It seems likely, therefore, that there *is* a theoretical exposure higher than E5 where exposed individuals are more likely to adopt Login Notifications than the unexposed—as would be predicted by **H3**. Unfortunately, we did not have a large enough number of people at a high enough exposure to empirically confirm this prediction from the data in our random sample.

It is tempting to also apply this logic to entertain a theoretical exposure lower than E1 at which the effect of social influence is negative for Trusted Contacts. However, as the exposure threshold for E1 for Trusted Contacts is just 0.1%, the theoretical and empirical exposure lower bounds are essentially the same—i.e., having at least one friend who uses the feature. Thus, while it seems like **H3** may be true even for Login Notifications, it seems likely that **H4** may not be true for some features—social influence does not *have* to be a negative force at low exposure conditions.

Individual Feature Attributes

Another consideration in interpreting the differences in the effect of social influence across security features is the individual attributes of each feature. Specifically, as **H2** suggests, more *observable* security features should be more positively affected by social influence.

The threshold beyond which the effect of social influence toggles from negative to positive appears to be inversely proportional to the observability of the feature, lending further support for **H2**. Indeed, the threshold is "lowest" for Trusted Contacts in that the threshold seems to be at its theoretical lowest possible value of having just *one* friend who uses the feature. The threshold is next lowest for Login Approvals at E5—or when approximately 2.7% of ones friends use the security feature. Finally, the threshold is highest for Login Notifications at a level of exposure higher than E5, if such a threshold exists at all.

It makes intuitive sense that the threshold of friends required for *negative* social proof to be overcome by *positive* social proof should be lower for more observable features. If our reasoning for **H4** is correct, *negative* social proof is the result of stereotypes and generalizations that may be overcome if potential adopters can see, concretely, that security feature usage is not necessarily limited to those who they may consider "paranoid" or who have an unachievable level of specialized knowledge about security.

Summary

In summary, the results from our matched propensity sampling analysis lends additional support to **H2** and conditional support to **H3** and **H4**. Specifically, the prediction, of **H3** and **H4,** that the direction of the effect of social influence on a potential adopter's likelihood to adopt a security feature will shift at a threshold appears to be true for Login Approvals and is likely true for Login Notifications. For Trusted Contacts, however, it appears that social influence has a positive effect on its adoption, regardless of the level of exposure. Furthermore, the *observability* of a security feature appears to at least partially moderate the presence and value of this threshold.

DISCUSSION

We analyzed whether and how security feature adoptions diffused through the social networks of 1.5 million people who use Facebook. Our results provide large-scale empirical evidence that social processes do, indeed, affect the uptake of security features—for better *and* for worse. For Login Approvals, when a potential adopter is exposed to few feature-adopting friends, social proof appears to reduce her own likelihood of adopting the feature. But, when she is exposed to many feature-adopting friends, social proof appears to increase her own likelihood of adopting the feature. The same is likely true for Login Notifications, but we did not find conclusive evidence that social proof can have a positive effect on its adoption at a high enough level of exposure. For Trusted Contacts, though, social proof appears to increase a potential

adopter's likelihood of adopting a security feature with even just one feature-adopting friend. These results, taken together, provide us with conditional support for **H3** and **H4**—depending on the feature, social proof can be either helpful or harmful to security feature adoption, and there may be a threshold of exposure to feature-using friends at which the effect transitions from harmful to helpful.

In addition, if a potential adopter is exposed to feature-adopting friends from many, independent social contexts—for example, a tennis friend and a travel buddy as opposed to two tennis friends—she is more likely to adopt each of the three security features, a finding that supports **H1**.

Finally, the strength of these effects varied across the features we tested—features that are more observable more effectively spread through social networks, a finding that supports **H2**. Indeed, we found that the effect of the *percentage of friends who use a security feature* and the *number of distinct social contexts* from which those friends originated were highest for Trusted Contacts (the most observable), followed by Login Approvals (the next most observable), and then Login Notifications.

Practical Implications

Prior work has outlined the positive potential of social influence in driving security feature adoption [11]. While we agree that this approach shows potential, our results suggest that social proof can also have a negative effect on the adoption of some security features for people with low exposure to feature-adopting friends. These findings should be taken into consideration before attempting to leverage social proof to drive the adoption of security features.

Nevertheless, social proof *can* drive the adoption of security features once those features come into use by more than just early-adopters. Indeed, our findings indicate that people with exposure to many friends who use Trusted Contacts and Login Approvals are more likely to use those features themselves, even after controlling for co-adoption due to homophily. Furthermore, our results suggest that the same may be true for Login Notifications, but at overall adoption rates higher than they are presently. Thus, the question remains: how do we increase the adoption of security features to reach the point where social influence has a positive effect on their spread?

Our results suggest that having friends from many social circles that use a security feature is strong social proof that a security feature should be used. Thus, one method to maximize the social spread of security features may be to target people from distinct social contexts and offer them personalized incentives to try security features.

It also appears that individual attributes of a security feature can affect its spread. Therefore, security feature designers should be mindful of the fact that use of a security feature can have social consequences. If, as prior work suggests, use of security features can be seen as a sign of paranoia, then designing security features like Trusted Contacts that draw friends in, can be seen, and feel trusting and collaborative rather than isolating may make it easier to convince people to use security features.

Limitations and Future Work

Our data was observational, so it is harder to make causal claims about the effect of social influence on the adoption of security features. The matched propensity sampling method we employed was an attempt at causal inference, but we still cannot say with absolute certainty that individuals exposed to feature-adopting friends adopted a security feature *because* of their friends. There may be some forms of latent homophily [26] for which we did not control—for example, a user's technical proficiency—that might explain some of the variance in our data. Thus, one rich opportunity for future work is to experimentally test, with consent, whether higher security feature observability does indeed yield greater diffusion of security features.

Future work should also explore how *other* feature attributes affect a feature's diffusion through a social network. While *observability* appears to be important, it alone likely does not account for *all* of the differences in effect size we observed across the three security features. The *diffusion of innovations* [21] literature suggests four additional attributes of a security feature that may be fruitful to test: its (1) *relative advantage*, or how much of an advantage the innovation provides over what it is replacing; (2) *complexity*, or how difficult it is to use the innovation; (3) *compatibility*, or how well the innovation matches a user's needs, experiences and values; and, (4) *trialability*, or how easily the user can try the innovation before making a final decision about whether to adopt it.

CONCLUSION

Through an observational analysis of whether and how security features diffused through the social networks of 1.5 million people who use Facebook, we illuminated the unique effects of social influence on security feature adoption. Social influence appears to drive security feature adoption when people have examples of feature-adopting friends from multiple social circles—for example, a high school friend *and* a family member as opposed to just two high school friends. However, the relationship between the number of one's feature-adopting friends and her own likelihood to use the feature is more nuanced. Indeed, when one is exposed to many security feature-adopting friends, she is likely to find positive social proof that increases her own likelihood to adopt a security feature. Conversely, when one is exposed to just a few feature-adopting friends, she might find negative social proof that stifles her adoption of the security feature. Furthermore, these effects may vary substantially across security features. Specifically, security features that are more *observable*—i.e., features for which the use and benefits are easily noticeable—seem to be more effectively spread via social channels.

ACKNOWLEDGEMENTS

This work was generously supported, in part, by NSF Award #1347186 and the NDSEG Fellowship. We would also like to thank Tiffany Hyun-Jin Kim, Stuart Schechter, Sofus Macskássy and Lada Adamic for providing helpful feedback on this and related projects.

REFERENCES

1. Adams, A. and Sasse, M.A. Users are not the enemy. *CACM 42*, 12 (1999), 40–46.
2. Aral, S., Muchnik, L., and Sundararajan, A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *PNAS 106*, 51 (2009), 21544–9.
3. Bakshy, E., Rosenn, I., Marlow, C., and Adamic, L. The role of social networks in information diffusion. *Proc. WWW '12*, ACM Press (2012), 519–528.
4. 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.
5. 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.
6. Centola, D., Macy, M., and Macy, M. Complex Contagion and the Weakness of Long Ties. *American journal of Sociology 113*, 3 (2014), 702–734.
7. Centola, D. The spread of behavior in an online social network experiment. *Science (New York, N.Y.) 329*, 5996 (2010), 1194–7.
8. Centola, D. An experimental study of homophily in the adoption of health behavior. *Science 334*, 6060 (2011).
9. Chapman, L.J. Illusory correlation in observational report. *Journal of Verbal Learning and Verbal Behavior 6*, 1 (1967), 151–155.
10. Cialdini, R.B. *Influence*. Harper Collins, 2009.
11. Das, S., Kim, H.J., Dabbish, L.A., and Hong, J.I. The Effect of Social Influence on Security Sensitivity. *Proc. SOUPS'14*, (2014).
12. Fowler, J.H. and Christakis, N. a. Cooperative behavior cascades in human social networks. *PNAS 107*, 12 (2010), 5334–8.
13. Gaw, S., Felten, E.W., and Fernandez-Kelly, P. Secrecy, flagging, and paranoia. *Proc. CHI '06*, ACM Press (2006), 591–600.
14. 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.
15. Herley, C. So long, and no thanks for the externalities. *Proc. NSPW '09*, ACM Press (2009), 133–144.
16. 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.
17. Kramer, A.D.I. The spread of emotion via facebook. *Proc. CHI '12*, ACM Press (2012), 767–770.
18. Milgram, S., Bickman, L., and Berkowitz, L. Note on the drawing power of crowds of different size. *JPSP 13*, 2 (1969), 79–82.
19. 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.
20. Rader, E., Wash, R., and Brooks, B. Stories as informal lessons about security. *Proc. SOUPS '12*, ACM Press.
21. Rogers, E.M. *Diffusion of Innovations*. Free Press, 2003.
22. Romero, D.M., Meeder, B., and Kleinberg, J. Differences in the mechanics of information diffusion across topics. *Proc. WWW '11*, ACM Press (2011).
23. Rosenbaum, P.R. and Rubin, D.B. The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika 70*, 1 (1983), 41–55.
24. Sasse, M.A. Computer security: Anatomy of a Usability Disaster, and a Plan for Recovery. *Proc. CHI '03 Wkshp on HCI and Security Systems*.
25. 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.
26. Shalizi, C.R. and Thomas, A.C. Homophily and Contagion Are Generically Confounded in Observational Social Network Studies. *Sociological Methods and Research 40*, 2 (2011), 211–239.
27. Sun, E., Rosenn, I., Marlow, C.A., and Lento, T.M. Gesundheit! Modeling Contagion through Facebook News Feed. *Proc. ICWSM '09*, (2009).
28. Tversky, A. and Kahneman, D. Availability : A Heuristic for Judging Frequency. *Cog. Psych. 5*, 2 (1973), 207–232.
29. Ugander, J., Backstrom, L., Marlow, C., and Kleinberg, J. Structural diversity in social contagion. *PNAS 109*, 16 (2012), 5962–6.
30. 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.

APPENDIX A

Variable Name	Login Approvals		Login Notifications		Trusted Contacts	
Intercept	0.28	***	0.12	***	0.31	***
Age	-0.06	***	0.08	***	-0.01	
Gender: male (relative to female)	0.05	***	-0.06	***	-0.10	***
Days with active account	-0.04	**	-0.26	***	0.02	
Friend count	-0.07	***	-0.11	***	-0.03	***
Days active in past 30	0.62	***	0.58	***	0.50	***
Mean friend age	-0.47	***	-0.21	***	-0.49	***
Friend age entropy	-0.16	***	-0.36	***	-0.07	***
Percent male friends	0.36	***	0.34	***	0.43	***
Mean friends' days with active account	-0.84	***	-1.00	***	-1.04	***
Friend country entropy	0.32	***	0.21	***	0.29	***
Mean number of friends of friends	-0.08	***	-0.02	**	-0.14	***
Posts created	-0.20	***	0.19	***	-0.17	***
Posts deleted	0.27	***	0.20	***	0.15	***
Comments created	0.10	***	0.06	***	0.18	***
Comments deleted	0.18	***	0.17	***	0.16	***
Likes given	-0.07	***	-0.09	***	-0.01	
Friends added	1.81	***	2.37	***	1.36	***
Friends removed	0.57	***	0.49	***	0.50	***
Photos added	0.10	***	0.14	***	0.26	***
Videos added	-0.01	***	-0.02	**	-0.02	**
Percent of friends who use feature	**0.13**	***	**-0.12**	***	**0.29**	***
Number of diverse social contexts	**0.15**	***	**0.03**	***	**0.88**	***

Table A1. Coefficients for the three logistic regressions relating social proof variables (bolded, at the bottom), to use of login approvals (left), login notifications (middle) and trusted contacts (right). All coefficients are normalized.

*** p < 2e-16, ** p < 0.001