

Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions

Steve Sheng,¹ Mandy Holbrook,¹ Ponnurangam Kumaraguru,² Lorrie Cranor,¹ Julie Downs¹

¹Carnegie Mellon University, ²Indraprastha Institute of Information Technology

xsheng@andrew.cmu.edu, holbrook@andrew.cmu.edu, pk@iiitd.ac.in,
lorrie@cmu.edu, downs@cmu.edu

ABSTRACT

In this paper we present the results of a roleplay survey instrument administered to 1001 online survey respondents to study both the relationship between demographics and phishing susceptibility and the effectiveness of several anti-phishing educational materials. Our results suggest that women are more susceptible than men to phishing and participants between the ages of 18 and 25 are more susceptible to phishing than other age groups. We explain these demographic factors through a mediation analysis. Educational materials reduced users' tendency to enter information into phishing webpages by 40% percent; however, some of the educational materials we tested also slightly decreased participants' tendency to click on legitimate links.

Author Keywords

Phishing, Social engineering, Roleplay, Survey, Mechanical Turk, User behavior

ACM Classification Keywords

J.4 [Social and Behavioral Sciences]: Psychology; H.1.2 [User/Machine Systems]: Software psychology; K.4.4 [Electronic Commerce]: Security

General Terms

Security; Human Factors; Experimentation; Measurement

INTRODUCTION

Phishing attacks, in which scammers send emails and other messages to con victims into providing their login credentials and personal information, snare millions of victims each year [12]. A variety of efforts aim to combat phishing through law enforcement, automated detection, and end-user education. Researchers have studied why people fall for phishing attacks; however, little research has been done to study demographic factors in susceptibility to

phishing. By determining which groups are most susceptible to phishing, we can determine how best to focus anti-phishing education.

In this paper, we present the results of our roleplay phishing study, administered to 1001 online survey respondents in order to study demographics and phishing susceptibility. The rest of the paper is organized as follows. In the next section, we present background and related work on why people fall for phishing. We then describe the design of our experiment and present the results of our study, identifying several important demographic factors that affect phishing susceptibility and describing the effects of education in bridging these gaps. Finally, we discuss the limitations of our study and the implications of our findings.

BACKGROUND AND RELATED WORK

Research has shown that people are vulnerable to phishing for several reasons. First, people tend to judge a website's legitimacy by its "look and feel," which attackers can easily replicate [2]. Second, many users do not understand or trust the security indicators in web browsers [31]. Third, although some consumers are aware of phishing, this awareness does not reduce their vulnerability or provide useful strategies for identifying phishing attacks [3]. Fourth, the perceived severity of the consequences of phishing does not predict users' behavior [4].

Demographics and Phishing Susceptibility

To the best of our knowledge, there has been no study dedicated to understanding what demographic factors correlate with falling for phishing and to what extent educational interventions have been effective in bridging the demographic divide. We highlight here a few studies that have measured susceptibility to specific types of phishing attacks or have studied the effectiveness of anti-phishing education while reporting at least some data on gender and other demographic factors.

Jagatic et al. performed a spear phishing experiment at Indiana University to quantify how reliable social context would increase the success of a phishing attack. They launched a phishing attack targeting college students aged 18–24 years old by using information harvested from social

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.

networking sites. In their study, they determined whether the 487 participants fell for the attack by observing whether participants provided their personal information to the phishing website. Female students fell for 77% of the spear phishing attacks, while male students fell for 65% [14].

In a related study, Kumaraguru et al. conducted a real-world phishing study with 515 participants to study the long-term retention of PhishGuru anti-phishing training [18]. They did not find significant differences based on gender, but did find that participants in the 18-25 age group were consistently most vulnerable to phishing attacks.

Finally, Kumaraguru et al. [16] conducted a study of 5,182 Internet users measuring the effectiveness of Anti-Phishing Phil, an interactive game that teaches people not to fall for phish. They found that men were more likely to correctly distinguish phishing and legitimate websites than women (75.5% correct vs. 64.4% correct). They collected only coarse-grained participant age data, but found that people under the age of 18 performed worse than those above 18.

Although past studies have found differences in phishing susceptibility based on gender and age, they generally did not collect enough information about study participants to isolate these variables from other potentially confounding factors. In addition, previous studies did not address why these demographic factors correlate with falling for phishing. In our paper, we address these research questions.

Susceptibility vs. Risk Behavior

The risk literature has shown reliable demographic differences in risk perceptions on various topics: relatively oppressed groups (e.g., women, racial and ethnic minorities, and less wealthy people) generally perceive the most risk in the world around them [11, 30]. Such perceptions may be linked to these groups' experiences of a riskier world, perhaps due to lower degrees of control over risky processes. Age has also been linked to risky behavior: on average, adolescents tend to engage in riskier behaviors, perhaps as a function of their ongoing learning about the world [27,6]. Because real-world risk behaviors are complex and subject to such varied predictors as knowledge, goals, and benefits associated with what is perceived to be risky behavior, there have been relatively few studies with the power to assess multiple mediators of demographic effects on risky behavior. The current paper takes a specific, well-defined behavior as a context in which to identify specific factors that may explain these effects.

Security User Education

Despite claims by some security and usability experts that user education about security does not work [13], there is evidence that well-designed user security education can be effective in the real world [18]. Web-based training materials, contextual training, embedded training, and interactive games have all been shown to improve users' ability to avoid phishing attacks [16, 17, 29].

A number of organizations have developed online training materials to educate users about phishing [7, 9, 22, 23]. In a previous study, Kumaraguru et al. tested the effectiveness of some of these online materials and found that, while these materials could be improved, they are surprisingly effective when users actually read them [16].

Several studies have adopted a contextual training approach in which users are sent simulated phishing emails by the experimenters to test vulnerability to phishing attacks. At the end of the study, users are given materials that inform them about phishing attacks. This approach has been used in studies involving Indiana University students [15], West Point cadets [10], and New York State employees [24].

A related approach, called embedded training, teaches users about phishing during their regular use of email. This trainer sends phishing email to users and, if users click on phishing links, immediately presents an intervention designed to train them to avoid falling for phishing attacks. Kumaraguru et al. created several intervention designs based on learning sciences, and found that these interventions were more effective than standard security notices that companies email to their customers [17]. The researchers continued to refine the most successful intervention, a comic strip featuring a character named PhishGuru. A follow-up study showed that people were able to retain what they learned from this training [18].

Finally, Sheng et al. designed Anti-Phishing Phil, an online game that teaches users good habits to help them avoid phishing attacks. The researchers used learning science principles to design and iteratively refine the game. Their evaluation showed that participants who played the game were better able to identify fraudulent websites compared to participants in other conditions [29].

We studied the effectiveness of several of these educational approaches in bridging the demographic divide, including a set of popular web-based training materials, Anti-Phishing Phil, a PhishGuru cartoon, and the combination of Anti-Phishing Phil and a PhishGuru cartoon.

STUDY DESIGN

In this online study participants answered survey questions to determine their background and assess their knowledge about phishing, and completed a roleplay task to assess their behavioral susceptibility to phishing prior to receiving one of several forms of training. Participants then completed a second roleplay task to assess reductions in phishing susceptibility as well as any changes in participants' tendencies to be suspicious of legitimate emails. Participants were assigned randomly to a control condition (no training) or one of four experimental conditions that varied based on the type of training to which participants were exposed. Half the participants completed the roleplay and then the survey questions; the other half completed the survey questions prior to the roleplay.

Email Subject	Legitimacy	Relevant features of email and websites
Earn Bonus Points #1	real	win a prize in an online scavenger hunt from BRU Information Security Office link: https://www.bru.edu/iso/aware/ncsam/hunt/bonus
Picture from last weekend's party	possible malware	impersonal greeting link: http://picasaweb.google.com/stevewulitzer/Partypics/ actual url: http://128.3.72.234/Partypics.jpg.exe
No obligation bankruptcy consultation	spam	text of link: "Apply online now" actual url: https://www.bankruptcylawyerfinder.com/freeconsultation.htm?...
Bandwidth Quota Offer	phishing	misspelling in url and .org domain link http://www.brubandwithamnesty.org/bandwidth/agree.htm actual url: same
eBay Accounts Security	phishing	threatens account suspension link: https://signin.eBay.com/ws/eBayISAPI.dll?SignIn&sid=verify ... actual url: http://www.security-validation-your-account.com/signin.ebay/...
Your Amazon.com Order (#103-0607555-6895008)	real	problem with shipping link: www.amazon.com/help/confirmation actual url: same
Your eBay item sold!	real	text of link: "Send Invoice Now" actual url: http://payments.ebay.com/eBayISAPI...

Table 1: A representative sample of emails in Pat's inbox from one of the roleplays.

Recruitment

Participants were recruited through Amazon.com's Mechanical Turk (mTurk), a marketplace for work requiring human intelligence. In this online environment, requesters post tasks known as HITs (Human Intelligence Tasks), and workers are paid for completing these HITs. We offered to pay participants four dollars for those that qualified and twenty cents to those who did not. In total, 1001 participants qualified and completed the entire study.

To disqualify people who were hoping to earn money for completing the study without actually paying attention to the study tasks, we asked all participants a series of questions about an email message that discussed an upcoming meeting. We used two of these questions, both of which could be answered correctly by a careful reading of the email, to screen out those participants who were not paying attention to the email content. We also asked basic demographic questions (such as questions about occupation and age) so that participants would not be able to easily identify qualifying questions [5].

Roleplay

Behavior was measured by performance in a roleplay task, with two equivalent exercises administered before and after training (the order of which was counterbalanced). This task is based on an established roleplay exercise that has been shown to have good internal and external validity [4]. The benefit of the roleplay is that it enables researchers to study phishing without conducting an actual phishing attack. Participants were told to assume the role of Pat Jones, who works at (fictitious) Baton Rouge University and uses the email address patjones@bru.edu for both work and personal emails. Each roleplay showed participants 14 images of emails along with context about Pat Jones that

may help them to interpret the emails. Images matched the participant's operating system and browser (e.g. Firefox on a Mac or Internet Explorer on a PC or other combinations) so that all images and cues would be familiar to the participant. Participants were asked to indicate how they would handle the emails if they received them in their own email inbox. Participants were asked to check boxes corresponding to all of the actions they would be likely to take from a list of responses generated through earlier qualitative work [3]:

- Reply by email
- Contact the sender by phone or in person
- Forward the email to someone else
- Delete the email
- Keep, save or archive the email
- Click on the selected link in the email (the one that the browser hand is pointing to)
- Copy and paste the selected URL (the www address) from the email into a web browser, if a URL is selected in this email
- Type the selected URL into a web browser, if a URL is selected in this email
- Click on a different link in the email (please specify which link(s) you would click on)
- Other (please specify)

The first email was created to familiarize the participant with the procedure. It was a short message from the same domain as Pat's email address. This message from the BRU Information Security Office announced a scavenger hunt for National Cyber Security month. The participants continued through the roleplay task by viewing a combination of real, phishing, malware and spam email

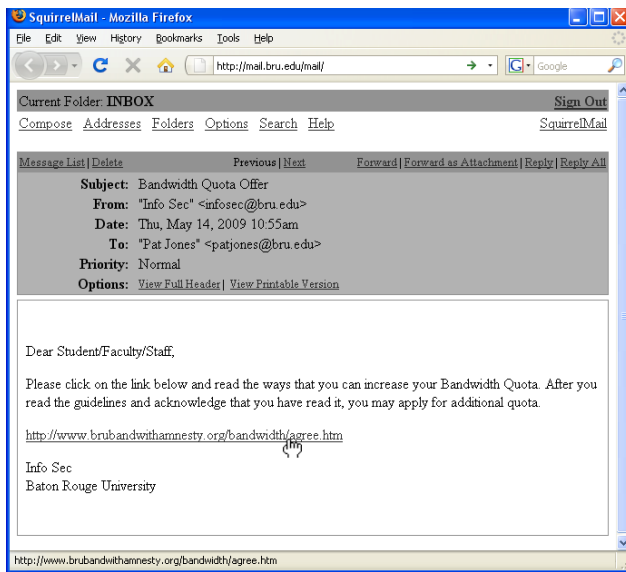


Figure 1: A phishing email used in our roleplay.

images. Table 1 lists a representative sample of the emails that Pat encounters in one of the roleplays. Each set of 14 emails included 6 phishing emails, 5 legitimate emails, 2 spam messages, and 1 possible malware email that contained links. Most of the links in these emails, including all bru.edu links and all of the phishing links, were fictitious links made up for this study.

As shown in the example email in Figure 1, each email contained a link to a web page shown with the mouse pointer positioned on the link and the actual URL destination displayed in the status bar as it would be if users prepared to click on the link on their own computer. For individuals who indicated that they would click on the link or otherwise end up at the web page, an image of that web page was displayed. Each web page requested information to be entered and participants were asked to indicate if they would click on a link on the page, enter the requested information, bookmark the page, visit another related web page, close the website, or take other action. No matter what other actions the user indicated, those who said that they would enter the requested information were coded as having fallen for phishing or complied with a legitimate email, corresponding to the legitimacy of the email in question.

Education Materials

Participants were randomly assigned to the control condition or to one of the four experimental conditions in which they were shown educational materials on ways to avoid falling for phishing attacks: a PhishGuru cartoon, Anti-Phishing Phil, several popular web-based training materials, and a combination of Anti-Phishing Phil plus a PhishGuru cartoon.

For popular web-based training, we selected three consumer oriented educational web pages from the first page of

Google search results using the search query ‘phishing’ — Microsoft Online safety [22], OnGuardOnline phishing tips [25], and National Consumer League Fraud tips [23]. In total, these materials have 3107 words, and would take roughly 15 minutes to complete reading with a scanning speed of 250 words per minute. After participants read each of the first two web pages they had a choice of reading more information or moving to the next part of the study. Participants in the popular training materials condition reviewed an average of 1.5 of the three websites and spent an average of 1.8 minutes on training.

In the Anti-Phishing Phil conditions, participants were taken through three levels of the game and allowed to exit at any point. On average participants completed 2.7 rounds of Phil in a little over 8 minutes.

The PhishGuru conditions provided participants with one page of materials and then participants moved on to the next part of the study. Participants in the PhishGuru condition spent an average of 0.5 minutes on training.

All participants who viewed any of the educational materials were asked how likely they would be to visit that specific educational tool again and how likely they would be to recommend it to someone else, on a scale ranging from 1 (not at all likely) to 7 (extremely likely).

Previous Experiences and Demographics

Along with asking participants extensive demographic questions, all participants were asked to complete a series of questions about their online experiences, including questions about their choice of websites for recent purchases, their use of online banking and their prior exposure to anti-phishing educational materials. Participants also indicated relevant negative experiences such as having information stolen or compromised in some way by entering it into a website. Table 2 presents basic demographics of the sample.

Knowledge and Technical Background

Knowledge questions prompted participants to choose the best definition for four terms related to computer security: ‘cookie,’ ‘phishing,’ ‘spyware,’ and ‘virus.’ Participants were given the same list of eight possible definitions to choose from for each, as well as choices to indicate lack of familiarity with the word. Each term had one correct answer on the list (shown here in parentheses). The options included:

- Something that protects your computer from unauthorized communication outside the network
- Something that watches your computer and send that information over the Internet (*spyware*)
- Something websites put on your computer so you don't have to type in the same information the next time you visit (*cookie*)
- Something put on your computer without your permission, that changes the way your computer works (*virus*)

Characteristics	Control	Popular training materials	Anti-Phishing Phil	PhishGuru Cartoon	Anti-Phishing Phil with PhishGuru
Sample Size	218	217	166	201	199
Gender					
Male	50%	48%	54%	45%	45%
Female	50%	52%	46%	55%	55%
Average age in years	30	30	29	30	31
Education					
High school or less	10%	8%	7%	7%	8%
Some college	33%	32%	37%	39%	36%
Completed 4-year college degree	29%	29%	30%	30%	27%
Some Post-graduate education	11%	12%	10%	6%	10%
Have master or PhD degree	17%	19%	16%	18%	17%
Percentage from US?	74%	71%	73%	78%	80%
Percentage student?	25%	26%	31%	20%	25%
Average years on the Internet	13	12	12	13	13
Average emails per day	44	44	32	57	43

Table 2: Participant demographics by conditions.

- Email trying to trick you into giving your sensitive information to thieves (*phishing*)
- Email trying to sell you something
- Other software that can protect your computer
- Other software that can hurt your computer
- I have seen this word before but I don't know what it means for computers
- I have never seen this word before
- Decline to answer
- Other (please specify)

To assess the level of their technology background, we asked participant if they had an Information Technology-related degree and any experience with programming languages, and they self-rated how technologically savvy they were on a scale ranging from 1 (not at all savvy) to 7 (extremely savvy).

Risk Perceptions

To evaluate participants' risk perceptions, we presented them with a series of statements taken from the Domain-Specific Risk-Taking scale of adult populations (DOSPERT) [1], drawing on the categories of financial risk and health and safety risk. These questions prompted participants to rate the risk associated with activities such as betting a day's income at the horse races and riding a motorcycle without a helmet, on a scale ranging from 1 (not at all risky) to 7 (extremely risky).

RESULTS

In this section we explain how we measured phishing susceptibility; describe our regression analysis; and then discuss the effects of gender, age, and anti-phishing education on phishing susceptibility.

Measuring Phishing Susceptibility

We measured participants' susceptibility to phishing by examining two kinds of errors in the roleplay scenarios before and after educational interventions: false positives and falling for phish. A false positive is when a user mistakenly judges a legitimate email or website as a phish and refuses to take the desired actions. Falling for phish occurs when a phishing email or website is incorrectly judged to be legitimate and users click on the email or submit information to the website. In some studies, falling for phish is determined based on whether users click on links in phishing emails; in other studies it is determined based on whether they submit information to phishing websites. In this study, similar to previous studies, we found that about 90% of participants who would click on a phishing link would go on to provide information to phishing websites [17, 18]. In this paper we focus our analysis on the stricter measure — giving information to phishing websites — as falling for phish.

Overall, prior to training, participants indicated they would click on 52% of phishing links and provide information to 47% of phishing websites. These results are similar to a previous real-world phishing study [16] in which 52.3% of participants clicked on the simulated spear phishing emails

Model Parameters	Standardized Coefficients
Prior exposure to anti-phishing training	0.19
Gender	0.14
Age	-0.12
Participants' technical knowledge	-0.10
Risk perception for financial investment	-0.08

Table 3: Regression analysis with parameters that are significant at $p < 0.01$

and subsequently 40.1% gave information to phishing sites. The similarity in our results suggested the validity of the roleplay survey instrument.

We found no significant differences in performance based on whether participants completed the survey before or after the first roleplay. Therefore we have collapsed those conditions and analyzed them together.

Regression Analysis

To explore factors that predict phishing susceptibility, we performed a multivariate linear regression. This section explains the steps we took to build the model and discusses the results from the linear regression.

We used factor analysis to reduce the dimensionality of our variables on participants' online experience (eight variables), participants' technical knowledge and experience (five variables), and participants' risk perception (12 variables). The factor analysis, using principle component and varimax rotation, reduced our list of variables from 40 to 22.

To study age groups and their vulnerability to phishing, we mapped age to the following categories: 18-25, 26-35, 36-45, 46-55, and >56.

We then ran the regression predicting falling for phish from the 22 variables. In Table 3 we report variables that are statistically significant at $p \leq 0.01$. Participants' degree of prior exposure to anti-phishing education significantly predicts their phishing susceptibility ($B = 0.189$, $p < 0.01$). Participants who had previous anti-phishing training (56.6% of total participants) fell for 40% of the phish in the roleplay, whereas those who had no previous anti-phishing training fell for 60% of phishing websites ($t(896) = -9.02$, $p < 0.001$). This factor had the most impact on phishing susceptibility, suggesting that exposure to education may play a larger role than other important factors.

Women fell for significantly more phish than men ($B = 0.140$, $t = 3.98$, $p < 0.01$), an average of 53.1%, compared to just 41% for men ($t(981) = -5.48$, $p < 0.001$). We explore

reasons for women's greater susceptibility in the next section.

Participants' age linearly predicts their susceptibility to phishing ($B = -0.116$, $p < 0.01$). An analysis of variance (ANOVA) comparing age groups found a significant overall effect ($F(4, 996) = 9.65$, $p < 0.001$) driven by participants aged 18 to 25 falling for phishing more than other age groups (all post-hoc tests comparing this group to other groups significant at $p < 0.01$; no other groups were significantly different from one another).

Participants' self-rated knowledge about technology also significantly predicts whether they will fall for phishing. For each standard deviation increase in their technical knowledge score, participants fell for 3.6% fewer phish.

Finally, participants' risk aversion, as measured by reactions to risks of financial investments, also predicts whether they will fall for phishing. The more risk-averse a participant is, the less likely he or she will fall for phish. For each standard deviation increase in their risk perception score, participants fell for 2.8% fewer phish.

Gender and Falling for Phish

In order to better understand why women appear to be more susceptible to phishing, we examined the effect of gender on clicking on phishing links, giving information to phishing websites, clicking on legitimate URLs, and giving information to legitimate websites.

We found that, before training, women were more likely than men to click on phishing links and enter information on phishing websites. On average, women clicked on 54.7% of phishing emails, compared to just 49% for men ($t(981) = 2.69$, $p < 0.01$). After clicking on a phishing link, women continued on to give information to the corresponding phishing website 97% of the time, compared to 84% for men ($t(981) = 5.42$, $p < 0.001$). This further exacerbates the gender differences in clicking on links.

In an attempt to explain these gender effects, we did a mediation analysis using all the key predictors as potential mediators. Mediation analysis explains "how" an effect occurred by hypothesizing a causal sequence. The basic mediation model is a causal sequence in which the independent variable (X) causes the mediator(s) (M) which in turn causes the dependent variable (Y), therefore explaining how X had its effect on Y [19, 20]. Mediation processes are common in basic and applied psychology.

We used the multiple mediator model developed by Preacher and Hayers [26] for our mediation analysis. For gender, we used technical knowledge and technical training as mediators; our hypothesis is that women have less technical experience than men and therefore fall for phishing more. Our results support this hypothesis. We report the mediation statistics in Table 4 and illustrate the results of the analysis graphically in Figure 2.

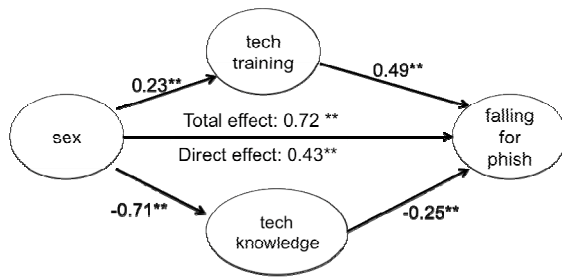


Figure 3: Mediation analysis for gender. Each path is quantified with unstandardized regression coefficients. The direct effect of gender on phishing susceptibility (measured by number of phishing websites participants’ giving information to) is calculated as total effect minus all the effect through each of the mediators, which is calculated as the product of coefficients in the paths.

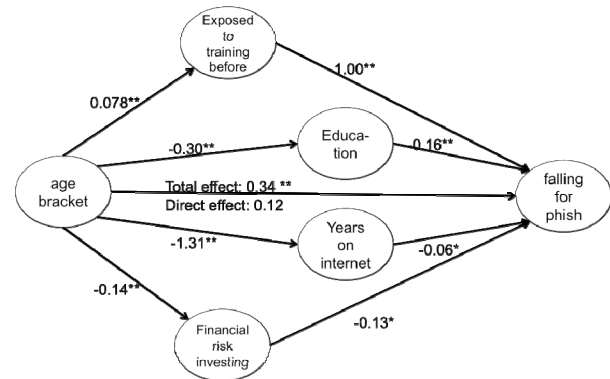


Figure 2: Mediating the effect of age with prior exposure to training, education, years on the Internet and risk perception for financial investment. Each of the paths is quantified with unstandardized regression coefficients.

	Point estimates	Percentile 95% CI	
		Lower	Upper
Total Effect of gender on falling for phishing	0.72		
Total effect of various mediators	0.29	0.18	0.42
Tech knowledge	0.17	0.10	0.27
Tech training	0.12	0.02	0.21

Table 4: Total effect of gender on falling for phishing and effect of various mediators that are significant at $p < 0.01$. The total effect is quantified with the unstandardized regression coefficient. The effect of various mediators is quantified as the product of coefficients in the paths shown in Figure 2.

	Point estimates	Percentile 95% CI	
		Lower	Upper
Total Effect of age on falling for phishing	0.34		
Total effect of various mediators	0.23	0.16	0.29
Exposed to training before	0.08	0.04	0.12
Years on Internet	0.08	0.03	0.13
Education	0.05	0.02	0.08
Risk financial investing	0.02	0.00	0.04

Table 5: Total effect of age on falling for phishing and the effect of various mediators that are statistically significant at $p < 0.01$.

As shown in Figure 2, the effect of being female on falling for phishing drops from a total effect of 0.72, $p < 0.01$, down to a direct effect of just 0.43, $p < 0.01$. The difference between these effects represents the total indirect effect through the two mediators, with a point estimate of 0.29, and a 95% CI of 0.18 to 0.42 (see Table 4). Thus, women in our study have less technical training and have less technical knowledge than men, which appears to partially account for their greater susceptibility to phishing.

The mediation relationship is only partial, as the direct effect is still statistically significant. This partiality suggests that there are other factors that are not captured by our survey instruments; these factors might be explored in future work.

We included several other predictors that did not mediate this relationship. For example, women may fall for phishing

more because they have fewer opportunities or are less motivated to learn about phishing. However, while women in our study had less technical training than men, more women than men claimed to have had prior exposure to anti-phishing education in particular. Thus, prior exposure to anti-phishing education did not turn out to be a significant mediator. Neither income nor education were significant mediators for the effect of gender on phishing susceptibility.

Other factors that we did not measure might potentially explain the remaining tendency for women in our study to be more susceptible to phishing than men. Factors that may be worth further exploration include differences in the way men and women use the Internet, differences in the way

men and women make trust decisions, and differences in the tendency of men and women to be cooperative or comply with instructions.

Age and Falling for Phish

As described above, people in the 18–25 age group were more likely to fall for phish than people of other ages. We used the multiple mediator model to determine why younger people are more susceptible to phishing. We report the mediation statistics in Table 5 and Figure 3.

Taken as a set, participants' prior exposure to phishing, number of years on the Internet, perception of financial risk, and education mediate the effect of age on falling for phishing. As can be seen in Figure 3, the total effect of age on falling for phishing fell from 0.34, $p < 0.01$, down to 0.12 (not significant). The difference between the total and direct effects is the total indirect effect through the four mediators, with a point estimate of 0.23, and a 95% CI of 0.16 to 0.29 (see Table 5). Because younger people have a lower level of education, fewer years of experience with the Internet, less exposure to training material, and less of an aversion to financial risks, they tend to be more susceptible to phishing.

Effects of Anti-Phishing Education

Before training, participants on average fell for 47% of phishing websites. After the training, this number reduced to 28%, a 40% improvement.

Table 6 summarizes the roleplay results by condition, before and after training. All training materials reduced participants' tendency to enter information into phishing webpages by about 34 to 44 percent, while there was no statistically significant improvement for the control group ($F(3,778) = 2.22$, $p = 0.84$).

Anti-Phishing Phil, the Phishguru cartoon and Anti-Phishing Phil with the Phishguru cartoon did not decrease participants' tendency to click on legitimate links and go to legitimate websites. However, in the popular training condition, participants' tendency to click on legitimate links was slightly reduced, ($t(216) = 2.01$, $p < 0.05$), suggesting that the participants may learn an avoidance strategy from popular training materials rather than strategies for better detection.

Since the various education materials perform similarly in reducing the number of people who fall for phishing, we combined all the training conditions together in order to study the effect of education in bridging demographic gaps.

We found that women in the training conditions learned more than men about avoiding phishing links ($t(767) = 5.63$, $p < 0.01$); after training, women and men performed equally well in not clicking on phishing links in emails ($t(767) = -0.05$, $p = 0.96$). However, women and men learned similarly about entering information into phishing websites ($t(767) = -1.51$, $p = 0.13$). Thus, both before and after training women were more likely than men to go on to enter

information into phishing websites ($t(767) = -4.22$, $p < 0.001$).

Finally, people of different age groups learned similarly from training, leaving no statistical difference between age groups' performance increase ($F(4,778) = 1.66$, $p = 0.16$).

Condition	Giving info to phishing sites		Clicking on legitimate websites	
	1st role play	2nd role play	1st role play	2nd role play
Control	50%	47%	70%	74%
Popular training	46%	26%	67%	61%
Anti-Phishing Phil	46%	29%	73%	73%
PhishGuru Cartoon	47%	31%	70%	64%
Anti-Phishing Phil with Phishguru cartoon	47%	26%	68%	59%

Table 6: Roleplay results by condition.

Participants between the ages of 18 and 25 were the most susceptible group in the first roleplay, and they remained more susceptible to phishing in the second roleplay. People in different education groups also learned similarly, ($F(5,763) = 1.4$, $p = 0.20$). We found no significant effect for education or race.

DISCUSSION

We conclude with a discussion of our study limitations and a summary of findings.

Limitations

There are several limitations to the current study. First, the sample was drawn from mTurk users and is not expected to be representative of the larger population of email users. Our sample of mTurk users tends to be younger, more educated, and more tech-savvy than the general public.

A second limitation of this study is the lack of direct consequences for user behavior. Participants might be more willing to engage in risky behavior in this roleplay if they feel immune to any negative outcomes that may ensue. Similarly, participants are not risking opportunity costs from being too conservative in their behavior. However, there is no reason to believe that the predictors described here should differ in their relationship to roleplay behavior compared to real-world behavior.

Summary of findings

Prior exposure to phishing education is associated with less susceptibility to phishing, suggesting that phishing

education may be an effective tool. Also, more risk-averse participants tended to fall for fewer phish.

Gender and age are two key demographics that predict phishing susceptibility. Specifically, women click on links in phishing emails more often than men do, and also are much more likely than men to continue on to give information to phishing websites. In part, this difference appears to be because women have less technical training and less technical knowledge than men. There is also a significant effect for age: participants aged between ages 18 and 25 are much more likely than others to fall for phishing (as seen by other researchers). This group appears to be more susceptible because participants in this age group have a lower level of education, fewer years on the Internet, less exposure to training materials, and less of an aversion to risks. Educators can bridge this gap by providing anti-phishing education to high school and college students.

All of the education materials in our study reduce users' tendency to enter information into phishing webpages by 40%. However, some education materials decreased participants' tendency to click on legitimate links; this finding suggests that educators need to do a better job of teaching people how to distinguish phish from non-phish so that they avoid false positives.

Demographics such as age, gender, race, and education do not affect the amount of learning, suggesting that good training materials can provide benefit for all groups. However, while the 40% reduction in phishing susceptibility after training is substantial, even after training participants fell for 28% of the phishing messages in our roleplay. This finding shows that education is effective and needed but is not a cure-all.

REFERENCES

- Blais, A. - R. and Weber, E. U. A domain-specific risk-taking (DOSPERT) scale for adult populations. *Judgment and Decision Making* 1, 1 (2006), 33–47 KW
- Dhamija, R., J. D. Tygar, and M. Hearst. 2006. Why phishing works. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM Press, New York, NY, 581-590.
- Downs, J., M. Holbrook and L. Cranor. 2006. Decision strategies and susceptibility to phishing. In Proceedings of the Second Symposium on Usable Privacy and Security (Pittsburgh, Pennsylvania, July 12 - 14, 2006). SOUPS '06, vol. 149. ACM Press, New York, NY, 79-90.
- Downs, J. S., Holbrook, M., and Cranor, L. F. 2007. Behavioral response to phishing risk. In Proceedings of the Anti-Phishing Working Groups 2nd Annual Ecrime Researchers Summit (Pittsburgh, Pennsylvania, October 04 - 05, 2007). eCrime '07, vol. 269. ACM, New York, NY, 37-44.
- Downs, J., M. Holbrook, S. Sheng and L. Cranor. 2009. Are Your Participants Gaming the System? Screening Mechanical Turk Workers. CHI 2010.
- Downs JS & Fischhoff B. (2009). Theories and models of adolescent decision making. In RJ DiClemente, JS Santelli & RA Crosby (Eds.) *Adolescent Health: Understanding and Preventing Risk Behaviors*, San Francisco, CA: Jossey-Bass.
- eBay. Spoof Email Tutorial. Retrieved March 7, 2006, <http://pages.ebay.com/education/spoofutorial/>.
- Evers, J. Security Expert: User education is pointless. Retrieved, Jan 13, 2007, http://news.com.com/2100-7350_3-6125213.html.
- Federal Trade Commission. An E-Card for You game. Retrieved Nov 7, 2006, <http://www.ftc.gov/bcp/online/ecards/phishing/index.html>.
- Ferguson, A. J. 2005. Fostering E-Mail Security Awareness: The West Point Carronade. *EDUCASE Quarterly*. 2005, 1. Retrieved March 22, 2006, <http://www.educause.edu/ir/library/pdf/eqm0517.pdf>.
- Flynn J, Slovic, P, and Mertz, C. K 1994. Gender, Race, and Perception of Environmental Health Risks. *Risk Analysis* 14(6): 1101-1108.
- Gartner Research. Gartner survey shows phishing attacks escalated in 2007. Press Release, 2007. <http://www.gartner.com/it/page.jsp?id=565125>
- Gorling, S. 2006. The myth of user education. In Proceedings of the 16th Virus Bulletin International Conference.
- Jagatic, T., N. Johnson, M. Jakobsson and F. Menczer. Social Phishing. Communications of the ACM. Retrieved March 7, 2006.
- Jakobsson, M. The Human Factor in Phishing. <http://www.informatics.indiana.edu/markus/papers/acidf, 2006>.
- Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L. F., and Hong, J. 2007. Teaching Johnny not to fall for phish. Tech. rep., Carnegie Mellon University. <http://www.cylab.cmu.edu/files/cmucylab07003.pdf>.
- Kumaraguru, P., Y. Rhee, A. Acquisti, L. Cranor, J. Hong and E. Nunge. 2007. Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System. In Proceedings of the 2007 Computer Human Interaction, CHI 2007.

18. Kumaraguru, P. Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M. A., and Pham, T. School of Phish: A Real-World Evaluation of Anti-Phishing Training. In the Proceedings On Usable Privacy and Security, 2009.
19. Mackinnon, D. P and Dwyer, J. H. Estimating Mediated Effects in Prevention Studies. *Eval Rev* 17, 2 (1993), 144–158.
20. Mackinnon D.P, Fairchild, A.J and Fritz, M. S. Mediation analysis. *Annual Review of Psychology* 58, 1 (12 2006), 593–614.
21. MessageLabs. Messagelabs Intelligence May 2009. Report, May 2009.
<http://www.messagelabs.com/intelligence.aspx>
22. Microsoft. Recognizing phishing scams and fraudulent emails. Retrieved Oct 15, 2006.
<http://www.microsoft.com/athome/security/email/phishing.msp>.
23. National Consumer League, Internet fraud tips: phishing. [url] <http://www.fraud.org/tips/internet/phishing.htm> . Retrieved Jan 1, 2009.
24. New York State Office of Cyber Security & Critical Infrastructure Coordination. Gone Phishing. A Briefing on the Anti-Phishing Exercise Initiative for New York State Government. Aggregate Exercise Results for public release. 2005
25. OnGuardOnline, How Not To Get Hooked by a "Phishing" Scam. [url]: <http://www.onguardonline.gov/topics/phishing.aspx>. Retrieved Jan 1, 2009.
26. Preacher K.J and Hayers A.F. Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior research methods* 40, 3 (Aug 2008), 879–91.
27. Reyna VF & Farley F. (2006). Risk and rationality in adolescent decision making: Implications for theory, practice, and public policy. *Psychological Science in the Public Interest*, 7, 1-44.
28. Schechter, S. E., Dhamija, R., Ozment, A., Fischer, I., 2007 The Emperor's New Security Indicators. *IEEE Symposium on Security and Privacy*, 20-23 May 2007.
29. Sheng, S., Magnien, B., Kumaraguru, P., Acquisti, A., Cranor, L. F., Hong, J., and Nunge, E. 2007. Anti-Phishing Phil: the design and evaluation of a game that teaches people not to fall for phish. In Proceedings of the 3rd Symposium on Usable Privacy and Security (Pittsburgh, Pennsylvania, July 18 - 20, 2007). SOUPS '07, vol. 229. ACM, New York, NY, 88-99.
30. Slovic, P 2000. *The Perception of Risk*, Sterling, VA: Earthscan Publications Ltd.
31. Wu, M., Miller, R. C., and Garfinkel, S. L. 2006. Do security toolbars actually prevent phishing attacks?. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006). R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, and G. Olson, Eds. CHI '06. ACM Press, New York, NY, 601-610.