Dissecting Spear Phishing Emails for Older vs Young Adults: On the Interplay of Weapons of Influence and Life Domains in Predicting Susceptibility to Phishing

Daniela Oliveira Harold Rocha Huizi Yang Donovan Ellis Sandeep Dommaraju
Melis Muradoglu ‡ Devon Weir Adam Soliman Tian Lin Natalie Ebner
University of Florida New York University ‡
{danielaoliveira,hrocha,huiziyang,donovanmellis,sandom,solimada,lintian0527,natalie.ebner}@ufl.edu,
mdm546@nyu.edu

ABSTRACT
Spear phishing emails are key in many cyber attacks. Successful emails employ psychological weapons of influence and relevant life domains. This paper investigates spear phishing susceptibility as a function of Internet user age (old *vs* young), weapon of influence, and life domain. A 21-day study was conducted with 158 participants (younger and older Internet users). Data collection took place at the participants' homes to increase ecological validity. Our results show that older women were the most vulnerable group to phishing attacks. While younger adults were most susceptible to scarcity, older adults were most susceptible to reciprocation. Further, there was a discrepancy, particularly among older users, between self-reported susceptibility awareness and their behavior during the intervention. Our results show the need for demographic personalization for warnings, training and educational tools in targeting the specifics of the older adult population [1].

ACM Classification Keywords
H.1.2 User/Machine Systems: Human Factors; J.4 Social and Behavioral Sciences Psychology; K.4. Computers and Society Social Issues: Abuse and crime involving computers; K.6.5 Security and Protection

Author Keywords
Spear-phishing, weapons of influence, susceptibility, aging

INTRODUCTION
Spear phishing emails are used as a key component of many cyber attacks and, more recently, also as a first step in advanced persistent threats [52, 28, 59, 43]. These attacks are appealing because they are simple, are low cost, and complicate attribution [10, 14].

Successful spear phishing emails apply psychological principles of influence – authority, commitment, liking, perceptual contrast, reciprocation, scarcity and social proof [13]. For example, according to the scarcity principle, opportunities seem more valuable when their availability are limited. An adversary can leverage this principle by tricking an Internet user into clicking on a malicious link to avoid missing out on a "once-in-a-lifetime" opportunity.

These principles of influence (called *weapons* in this paper to emphasize their malicious use in the present context) exploit common human heuristics that are often beneficial in simplifying decision-making, but can also result in misrepresentation, and can lead to deception. The effectiveness of these weapons in spear phishing emails can be increased when the email places the weapon in a life domain context that is relevant for the user, such as the financial, health, ideological, legal, security, and social domains. For example, an email to a person deeply touched by gun violence inviting him to sign a petition for background checks before gun purchases pertains to the ideological domain.

We propose that older compared to younger Internet users constitute a particular at-risk population for spear phishing attacks. This hypothesis is based on evidence that general cognitive processing capacities, as well as sensitivity to deception decline with age, while self-reported trust increases [56, 39, 29, 38, 54]. Further, recent reports have shown that while scamming is prevalent at all ages, they are most effective in older adults [39, 21, 36].

Older adults' possible increased vulnerability to such attacks is relevant to cyber security. Older adults are the fastest growing segment of the U.S. population [12, 2]. Further, they often have accumulated financial assets over a lifetime and many occupy powerful positions in finances and politics.

These considerations raise the following research questions: (i) Do younger and older Internet users differ in their susceptibility to spear phishing email attacks? (ii) Which weapon(s) is/are particularly effective? (iii) Does effectiveness of weapons vary by age group? (iv) Which life domain(s) is/are particularly effective? (v) Does effectiveness of

[1]Rocha, Yang, Ellis, and Dommaraju contributed equally to this paper.

CHI 2017, May 6-11, 2017, Denver, CO, USA.

http://dx.doi.org/10.1145/10.1145/3025453.3025831

life domains vary by age group? (vi) Are younger and older Internet users aware of their susceptibility to spear phishing attacks?

To address these questions we conducted a user study over a period of 21 days [2]. The study took place in the participants' homes to increase ecological validity. During the study interval, we exposed 100 younger and 58 older Internet users to experimentally controlled spear phishing emails to determine attack susceptibility as a function of age, weapon of influence, and life domain.

At study start participants installed a browser extension on their personal computers that recorded all the URLs they visited during the study period. Without their knowledge, participants received 21 spear phishing emails over the course of the study. The emails systematically varied in the weapon of influence applied and the life domain they referred to.

Susceptibility to spear phishing emails constituted our central outcome variable. It was operationalized as clicking on the email link provided in each of the spear phishing emails, as indicating that a participant fell for the attack.

Our results show that Internet users are highly susceptible to spear phishing emails. More than 40% of our participants clicked on an email link at least once during the study period. Older women were the group most susceptible to the attacks. While scarcity was the most effective weapon for younger adults, reciprocation was the most effective for older adults. Authority was highly effective for all users. Interestingly, we found a discrepancy, in particular among older users, between high behavioral susceptibility and low self-reported susceptibility awareness.

Currently, still little is known about deception and aging, especially in the context of cyber security. Also, the majority of research on decision-making in aging pertains to the financial or health domains, and is conducted in lab or via surveys [41, 48, 58]. Further, current phishing research has largely focused on general education and anti-phishing exercises [9]. To our knowledge, our study is the first behavioral measurement of aging effects on decision-making related to cyber security in the real-life context.

This paper starts with a review of related work in the areas of aging, trust, deception and spear phishing. Then we provide the theoretical background on the study constructs: weapons of influence and life domains. Next, we outline the study methods and the design and implementation of the study framework. We also provide examples of the simulated spear phishing emails used in this study. We proceed with a detailed analysis of the results, and end with conclusions.

RELATED WORK

Our work intersects the areas of cyber security and spear phishing with trust and deception in aging. In this section we discuss work in these areas in relation to the present study context.

[2]This study was approved by the University IRB.

Spear Phishing Attacks

Spear phishing email attacks are a widespread threat that impacts the security of individual computer users as well as whole organizations. These attacks leverage emails, instant messaging, social media and QRishing [57]. Protection solutions either propose elimination of malicious emails before users even see them, or support of users in detecting phishing using email features and machine learning methods [22, 55, 61, 51].

While phishing web sites are beyond the scope of our work, it is relevant to discuss some of the tools available to counter these attacks. Many of these tools are in the form of web browser extensions that warn users when they are browsing a suspicious phishing website. Examples include the Microsoft built-in phishing filter of Internet Explorer 7, the Netcraft Anti-Phishing Toolbar [55], and SpoofGuard. These tools, however, do not effectively protect against all phishing attacks, as attackers and tool developers are engaged in a continuous arms race [61]. Blacklists are ineffective when protecting users initially, but are updated at different speeds and vary in coverage [51]. The source of phishing URLs and the freshness of the URLs tested can significantly impact the results of anti-phishing tool testing, and many of the tools are vulnerable to simple exploits [61]. Moreover, many passive indicators implemented in anti-phishing tools are failing users because users do not understand or believe them [20, 60]. For example, Aquarium [37] employs crowdsourcing to identify phishing websites. It first clusters similar phishing websites together, has participants vote on clusters, and then applies a vote weighting mechanism based on participants' prior performance. Dhamija *et al.* [15] investigated what makes a phishing website credible with a study in which participants were shown websites to determine which ones were bogus. Forty percent of the time participants made incorrect choices because they did not look at security indicators of a page.

Phishing education has been researched as a line of defense against phishing. Anti-Phishing Phil [50] teaches employees of organizations good habits to avoid phishing attacks. PhishGuru [31, 33, 34] was developed as an embedded training system for identification of phishing attacks in emails. A challenge of phishing training is that people forget the learned material, and then fall for the same attacks shortly after the training. Caputo *et al.* [9] investigated anti-phishing exercises in the context of a large organization. The training had no significant effect on the likelihood that users clicked on a subsequent spear phishing email. Rather, users typically either clicked all links or none independent of training.

With few exceptions, a shortcoming of previous approaches is that user demographics were not taken into account. Sheng *et al.* [49] showed that women were more susceptible than men to phishing, and users between the ages of 18 and 25 were more susceptible to phishing than other age groups. Similarly, Kumaraguru *et al.* [32] found that users aged 18-25 years were more vulnerable to phishing attacks than older participants. These past studies, however, considered middle-aged rather than older (60 years and older) Internet users. In fact, there is some indication that adults aged around 55 years

may actually experience what is called their "golden age" of decision-making, characterized by high levels of fluid intelligence (processing speed, working memory) combined with high levels of crystallized intelligence (experience) and thus may be an age group that is not particularly at risk for faulty decisions [47]. In contrast, our study included older adults (65 years and older).

Another shortcoming of previous work is that phishing features and analysis were handled in an ad hoc fashion. For example, Caputo *et al.* [9] used only three different email messages during their trials in a large organization. Our study used 21 spear phishing emails per user, with those emails systematically varied regarding weapons of influence and life domains.

A separate line of work investigated reasons why people fall for phishing. Downs *et al.* [17] interviewed non-expert computer users and found that they had difficulty in perceiving and acting on unfamiliar risks. A follow-up survey [16] asked participants to indicate how they would respond to five emails from a hypothetical person's mail box. Our work extends this previous work in that it not only takes a behavior-based approach on understanding phishing susceptibility, but also is the first to investigate susceptibility in the context of psychological principles of persuasion and common life domains.

Trust, Deception, and Aging

Trust is fundamental to satisfying social relationships [35, 42]. It is currently still a largely understudied aspect of social relationships in aging. There is some evidence that perceived trust increases, while sensitivity to untrustworthy information declines with age [11, 45, 18, 19]. For example, older compared to younger adults' show impaired ability in detecting lies [45] and discriminating between trustworthy and untrustworthy faces [11]. These findings are in line with a well-documented age-related positivity bias, and more specifically, reduced attention and memory to negative compared to positive information in aging [46, 1].

There is broad research addressing risk and information perception of non-expert Internet users [3, 25, 4, 5, 27, 24]. For example, Asghapours *et al.* [5] advocate the use of mental models of computer security risks for improvement of risk communication to non-expert end users. The authors leveraged five mental models, as simplified internal concepts of processes in reality: Physical Safety, Medical Infections, Criminal Behavior, Warfare, and Economic Failures. Their study showed that mental model of security risks correlated with level of expertise. Garg and Camp [24] adopted the classic Fischhoff's canonical nine dimensional model of offline risk perception [23] to better understand the dimensions of online risk perceptions of end users. Results obtained for online risks differed from the ones obtained for offline risks. In addition, the severity of a risk was the biggest factor in shaping risk perception. Particularly relevant in the context of our work, Garg and Camp [25] confirmed the risk perception model with respect to older adults' general security education using educational videos and texts.

Boothroyd [7] examined, in self-report, 15 older and 16 younger adults' risk perceptions related to e-mail and Facebook. In contrast, our study used an experimental manipulation in a larger sample of 100 young and 58 older adults to measure susceptibility behavior. Our self-report results corroborate Boothroyd's findings, in that older adults' self-reported susceptibility awareness was lower than that of younger adults.

Based on this previous evidence, we hypothesized that, compared to younger users, older Internet users would be particularly susceptible to weapons of influence embedded in certain life domains in the context of well-crafted spear phishing emails. In addition we proposed that older adults may have lower awareness of their susceptibility to these attacks than younger users. The relevance of testing these assumptions is supported by information released by the FBI that senior citizens are common targets of scams, likely because con artists are aware of both age-related cognitive decline and often exceptional credit increasingly managed online, among the elderly [44].

WEAPONS OF INFLUENCE AND LIFE DOMAINS

In this section, we discuss the central study constructs of weapons of influence and life domains, and how they relate to spear phishing email attacks.

Principles of Influence.

Cialdini [13] proposed that humans have fixed behavioral patterns that are triggered by certain events. The explanation for such patterns is that humans have been hard-wired through evolution to use heuristics, which are often beneficial in promoting fast and frugal actions [30]. However, heuristics can also be exploited by adversaries to lure an individual into a behavior that is counter to his best interests. Cialdini calls the events that trigger these behavioral patterns *psychological principles of influence*. He differentiates seven principles, or *weapons* when used by adversaries: authority, commitment, liking, perceptual contrast, reciprocation, scarcity, and social proof, as briefly summarized next.

According to the principle of **Scarcity**, opportunities seem more valuable when their availability are limited [13]. An adversary can leverage this principle by tricking an Internet user into clicking on a malicious link to avoid missing out on a "once-in-a-lifetime" opportunity.

The principle of **Authority** states that humans tend to comply with requests made by figures of authority, such as law enforcement personnel, lawyers, or politicians [40, 26]. Therefore, sending emails in the name of authorities might be effective in luring recipients into clicking on malicious links.

The **Commitment** principle proposes that once humans have taken a stand, they will feel pressured to behave in line with their commitment. For example, consider Bob, a dog lover who feels devastated when he learns about cases of animal abuse and is vocal on Facebook about this issue. An adversary can target Bob by sending him an email about a petition to end animal cruelty in makeup testing.

The **Liking** principle assumes that humans tend to comply with requests from people they like or with whom they share similarities. Consider Bob, an older adult active in his church. Bob will feel more at ease accepting a request coming from Dan, 68, member of the same church, than a request coming from Alice, 19, student at a local University.

The **Perceptual Contrast** principle refers to the way people perceive a difference between two things that are presented subsequently. When the second item is rather worse than the first, people tend to see the first as more attractive than it actually is. For example, an adversary exploits this principle when he sends an email to a person with a certain health condition contrasting what the person pays for his medication now with what he could be paying in the future if he switches to a "generic" brand advertised on a malicious website.

The **Reciprocation** principle is based on the notion that humans tend to repay, in kind, what another person has provided them. An adversary can use this principle to lure a user into installing malware on his computer by offering a free gift attached (e.g., a travel guide).

Finally, the **Social Proof** principle leverages the human tendency to avoid mistakes by acting according to the majority of other people. An adversary can exploit this principle by advertising via a malicious link an offer from a company "voted" as one of the top 10 in the country.

Life Domains.

To increase effectiveness, attacks can place a phishing email in a particular life domain context. Personal relevance of specific life domains differs interindividually, and varies by demographics, such as by gender and age. In our study we considered six life domains: financial, health, ideological, legal, security, and social [6].

Financial emails focus on money, discounts, or offers, *e.g.*, an email offering a discounted vacation package. **Health** emails focus on mental and physical wellness, including sports and medication, *e.g.*, an email reminding an older adult about an upcoming follow-up with his physician. **Ideological** emails focus on principles and beliefs, *e.g.*, an email inviting to sign a petition. **Legal** emails focus on the law, such as emails about being arrested or sued, or an invitation to appeal a parking ticket. **Security** emails focus on safety, such as neighborhood watch or cybersecurity, for example, an email invite to explore the website of a neighborhood watch organization. Finally, **Social** emails focus on social interaction, including dating or participating in an interest organization, such as an email inviting an individual to learn more about an upcoming local farmer's market.

In this study we systematically created a set of fake spear phishing emails with each email referring to one of these seven weapons and related to one of these six life domains.

METHODS

This section presents the scientific methods of our study: recruitment, participant management, and the study procedure. The recruitment phase started in spring 2015. A pilot study was conducted in summer 2015. The actual study lasted from August 2015 to November 2016.

Participants

The study comprised 158 Internet users (100 younger; M = 21.74 yrs, SD = 4.11, range: 18-37, 56% females; 58 older; M = 71.7 yrs, SD = 6.80, range: 62-89., 43% females), recruited from the North Central Florida area through fliers and handouts posted throughout town and the county, newspaper ads, a lab internal participant pool, the University Psychology Subject Pool, HealthStreet [3], and Internet sources such as Facebook, Craigslist, and Researchmatch.org. Younger participants who were recruited through the University Subject Pool were compensated with course credit for study completion; all other participants received $50.

The written consent form disclosed study procedures, the minimal study risk, and data protection mechanisms. Participants were not informed about receipt of emails to ensure natural behavior, and were fully debriefed at study closure. To be included in the final data analysis, participants had to complete 21 days of the study session (to ensure they received all 21 spear phishing emails) and at least 50% recorded daily email inbox checking activity (explained below). These criteria excluded from the analyses 21.8% (n = 44) of the total 202 enrolled participants. Participants who had dropped out were equally distributed across gender and age. Table 1 summarizes demographic, health, and Internet usage information by age group.

Procedure

The cover story given to participants was that the study examined patterns in daily internet use. The study started with a phone screening, a written informed consent, and an installation phase, followed by the 21-day intervention phase. The study concluded with a survey and debriefing. During the screening, older participants underwent the Telephone Interview for Cognitive Status [8], a brief, standardized screening tool to determine cognitive status (cut off for study inclusion was a score > 30).

The study took place in the participants' home. After informed consent, participants installed (with help via the phone or by using premade video tutorials) the web browser extension for the study (see section Study Framework for details) on their primary home computer. The 21-day study period began the following day. On Day 1, participants filled in a short demographics questionnaire.

On each day of the 21-day intervention phase, participants engaged in one hour of study-related internet browsing. To assure diverse activity and regular email inbox checking activity, Internet browsing was divided into the following four specific activities, approximately 15 minutes each: (i) reading an informative source; (ii) reading entertainment/social network sources; (iii) engaging in unstructured browsing time; (iv) checking/handling emails from the email account the user had registered for the study. We included these activities because a targeted behavior can only be observed if the study

[3]A university-affiliated community recruitment and outreach program.

	Younger Users ($n = 100$) M (SD)/%	Older Users ($n = 58$) M (SD)/%	Age-Group Differences
Years of Education (younger n = 81, older n = 33)			
	14.24 (3.57)	16.30 (2.79)	$t(112) = -2.97, p = .004$
Highest Degree Earned (younger n = 81, older n = 33)			
Highschool	59.8	15.1	
Associates	10.9	15.1	
Bachelor's	14.1	28.3	
Master's	13.0	26.4	
Doctorate	0.0	13.2	
Other professional degree	2.2	1.9	
			$\tilde{\chi}^2(6) = 35.11, p < .001$
Annual Income (younger n= 91, older n = 54; participants with negligible income/unemployed reported "N/A")			
$< \$40,000$	36.3	33.3	
\$40,000-\$70,000	0.0	31.5	
$> \$70,000$	0.0	18.5	
N/A	63.7	16.7	
			$\tilde{\chi}^2(3) = 61.83, p < .001$
Race/Ethnicity (younger n = 91, older n = 52)			
American Indian or Alaskan Native	0.0	5.8	
Asian	29.7	1.9	
Black/African American	8.8	9.6	
Native Hawaiian	0.0	1.9	
Hispanic	19.8	1.9	
White	39.6	76.9	
Other	2.2	1.9	
			$\tilde{\chi}^2(6) = 36.68, p < .001$
Marital Status (younger n = 91, older n = 53)			
Single	68.1	15.1	
In a relationship, but not married	26.4	15.1	
Married	4.4	47.2	
Divorced/Separated	1.1	15.1	
Widowed	0.0	7.5	
			$\tilde{\chi}^2(4) = 69.09, p < .001$
Physical Health (younger n = 91, older n = 52; Rating scale from 1 = "Poor" to 10 = "Excellent")			
	7.69 (1.56)	7.58 (1.95)	$t(141) = .39, p = .70$
Mental Health (younger n = 91, older n = 52; Rating scale from 1 = "Poor" to 10 = "Excellent")			
	8.02 (1.53)	8.54 (1.5)	$t(141) = .1.86, p = .05$
Internet Usage/Week (younger n = 81, older n = 49)			
≤ 4 hours	21.0	20.4	
4-9 hours	34.6	44.9	
$geq10$ hours	44.4	34.7	
			$\tilde{\chi}^2(2) = 1.56, p = .46$
Average Number of URLs Visited/Day (younger n = 99, older n = 57			
	256.46 (166.14)	181.52 (115.85)	$t(154) = 3.01, p = .0030$

Table 1. Demographic, health and internet usage information about study participants by age group.

allows for a sufficient base rate of that behavior. The required activities (i.e., browsing and email checking – representative of everyday computer use) ensured a sufficient base rate. Individuals who rarely or never used computers were not eligible for the study.

The browser extension recorded all user's web activity. To ensure data integrity, participants were instructed to only use the study-registered browser, on the computer they had registered for the study.

While active in the study, participants received, without their knowledge, 21 (one per day) simulated spear phishing emails in their registered email account. These emails were drafted by drawing from examples of real-life spam and spear phishing emails collected in the context of a pilot project. The emails were formatted to have similar word lengths (between 50 and 150 words) and follow a similar structure. The topics of the emails were not restricted to older adults – they covered various life domains based on the literature [6]. To increase believability, each email included information, events, and contexts related to the geographic area (*e.g.*, city, county) targeted by this study. Further, each email addressed participants by their names [4] and had a unique sender. Each email contained a link that directed the user to a harmless static web page designed by our research group. We purchased seven domains to host the 70 façade web pages created for this study. Clicking on an email link in one of our spear phishing emails constituted our central outcome variable *susceptibility to spear phishing emails*.

The phishing emails were sent from 22 fake email accounts (11 male and 11 female names) hosted by Gmail, AOL, and

[4]Participant's anonymity was maintained throughout the study. Data collected was de-identified via a code that reflected gender and age group of the participant. Our report of the results does not make any reference to individuals.

Outlook/Hotmail. To prevent our fake accounts from being blocked and the study e-mails from being flagged by the email providers, our research team accessed each account daily. Simulated spear phishing emails were not sent out until an account had two months of lab-maintained regular interpersonal communication (*e.g.*, with other lab email contacts). Once the accounts were actively used in the study, we continued to maintain them.

After debriefing, we asked participants (and compensated them for) to install an automated spam extraction tool to send the contents of their spam folders to a lab secure server for analysis. This allowed us to verify whether the study emails ended up in the participants' spam folder. We only noticed one case of this issue and promptly replaced the email.

For successful and consistent sendout of our phishing emails we used mainstream email domains (Gmail, Hotmail, AOL, and Outlook). Until recently, there was a propensity for spam/phishing to arrive from senders of unknown domains (i.e @hotfax.com). However, the new generation of phishing attacks is targeted and sophisticated. Also, more convincing characteristics are increasingly being used, such as better grammar, spelling, targeted interests, and established domains [53]. The goal of our study was to study susceptibility using the new generation of spear-phishing attacks.

Gender of sender and domain-weapon combination were counterbalanced across age groups and gender of participants. In particular, participants were assigned to one of four sets of 21 emails. These four sets were created as follows: every combination of the seven weapons and six life domains was utilized for a total of 84 spear phishing emails across the entire experiment.

Each email was reviewed by three writers to assure representation of each weapon-domain combination. This collection of emails was divided into two pseudo-random sets so that each set had three instances of each weapon and each domain. Each set was then duplicated by writing equivalent emails for each weapon-domain combination.

On Day 21, participants were asked to review 21 spear phishing emails from a complementary set and rate them in terms of how interested they were in the emails, how likely they would be to click on the link, and how convincing the argument of the email was. At the end of this final survey, participants were given a manipulation check and were debriefed about the true study purpose. Participants were then contacted one last time to uninstall the browser extension.

STUDY FRAMEWORK

This section details the design and implementation of our study framework. Figure 1 illustrates the high level view of our architecture.

The browser extension (Merlin) was installed by users prior to Day 1 of the study and tracked all the URLs participants visited during the study period together with the timestamps for access. Merlin communicated continuously with a module in our server called *Log manager* (see Figure 1), which recorded the data in log files. Merlin did not collect any identifiable or sensitive information from the participants, such as bytes transferred, keys typed, or files accessed. Participants were identified by a user ID that reflected their gender and age group.

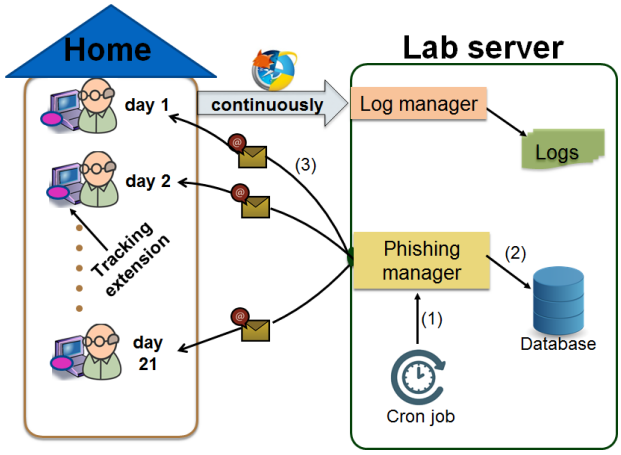

Figure 1. Study framework. Daily cron jobs invoked the phishing manager (1) to fetch participant and schedule information and phishing emails from the database and (2) to send phishing emails to the users (3). The Merlin extension continuously sent all URLs users visited to the log manager, which recorded the data in log files (see "continuously" in the figure).

The *Phishing manager* module was responsible for preprocessing and sending phishing emails to participants following a controlled timeline and the counterbalancing scheme. The phishing manager operated when invoked by a set of cron jobs (see [1] in Figure 1). Depending on the day of the week, participants received emails either at a random time during the day, or at a time as close as possible to the time they were engaged in study-related browsing activity. The phishing manager accessed the study database (see [2] in Figure 1) to: (i) retrieve users' study schedule and information (gender, first and last name, email address and county of residence), and (ii) pre-process and send spear phishing emails according to the counterbalancing scheme (see [3] in Figure 1).

The phishing email preprocessing allowed participants to receive targeted spear phishing emails. The phishing manager also documented and managed participants' activity, *e.g.*, reporting absence of study activity, sending reminder emails for study activities, and recording the day of study.

Implementation Details

Our study framework was deployed on a Linux machine with Intel Xeon CPU E5-2650 0 2.00GH and 4GB of memory. The server also held a MySQL database that included the set of phishing emails and participants' information. We used WordPress to host the façade web pages. Participants' log data occupied 3.04 MB; on average, 23.38 KB per user.

We developed Merlin extensions for the most popular browsers: Chrome, Firefox, Safari and Internet Explorer. Merlin created a unique ID for each user. A popup window appeared when the user clicked on the extension icon for the first time during the installation process asking for the participant's age group and email address. The email address

was required to allow personalization of the spear phishing emails. URL logs were associated with the user ID only. Merlin could not track web activity if the participant was browsing in private mode. Thus, a study inclusion criteria was no browsing in private mode.

Sending Spear Phishing Emails

Our framework was composed of five cron jobs (Cj1-Cj5) that invoked the phishing manager, implemented in Java. Cj1 ran every Sunday, Monday, Wednesday and Friday at 9am. It invoked the phishing manager to randomly select a time between 9am and 5pm, for sending to each active participant a personalized spear phishing email according to the counterbalancing schedule.

Cj2, ran every Tuesday, Thursday, and Saturday and invoked the phishing manager to track user web activity and send the personalized email as soon as it detected survey activity from this user. Cj2 invoked the *phishing manager* every hour from 1am to 10pm to perform the check. At 10:30pm, Cj3, invoked the phishing manager to send a spear phishing e-mail to all users who had not received a phishing email on that day because of lack of web activity.

Cj4 ran daily at 9am to invoke the phishing manager to scan the day of all active users in the study and then send a reminder email if users were on Day 1, 7, 14, or 21 of the study. This email reminded users about which day they were in the study, about how many days they had left, not to change browsers over the course of the study, and of the daily as well as upcoming final survey.

Finally, Cj5, ran every day at 11:59pm to invoke the phishing manager to check the participants' daily activities and generate a report. Checking these activities allowed our team to contact participants in a timely manner in case of lapsed study activity. We contacted users if their Internet activity was absent for three subsequent days or if they had not yet finished their daily survey.

EXAMPLES OF PHISHING EMAILS

In this section we provide examples of the body of some of the emails used in our study combining weapons and domains.

Authority and Legal: *"Our resources have indicated that you have a parking violation from 12/17/2015 at SW 89th Avenue at 3:34pm. Please go to our website to obtain more information about the violation and to pay your fine or refute your ticket: < link >"*

Commitment and Ideological: *"Would you like to put an end to animal abuse? 3 out of 10 domestic animals in the US are currently abused by their owners. That is why we started our non-profit organization, PetLove, to lobby for additional support for shelters that take in abused pets. Help us end this unfair treatment of our beloved pets. Last year we were able to collect just over 15,000 signatures, and this year we look forward to collecting more. Make your mark by signing the petition to help these animals have the future they deserve. Please click the link below to sign digitally! < link >"*

Liking and Security: *"My name is Dan. I don't know if I ever formally introduced myself to you, but I am your neighbor from down the street. I'm glad to finally be in touch with you! I am writing to ask you if you have seen any suspicious activity. I know there have been some issues of security breaches around town of people breaking into homes or cars in the past, and I just wanted to make sure everyone stays safe in our neighborhood. If you're interested, I made a blog with tips on how to keep your home and family safe that you can visit at < link > Please let me know if you have any questions or want to get together some time; I love spending time with my fellow neighbors."*

Perceptual Contrast and Health: *"How much are you paying for Nebivolol? Recent changes in the manufacturing process of Bystolic and Nebilong means that capsules are now being produced at a fraction of the original cost. If you are currently paying more than $40 for 10 mg (about 30 capsules) of Bystolic, Nebilong, Hypoloc, Lobivon, Nebilet, Nebilox, Nobiten, or Temerit, then you might be in need of crucial information that can save you $20 for 30 capsules. For a list of recommended pharmacies with reduced rates for Nebivolol, please visit the link below: < link >"*

Reciprocation and Social: *"Congratulations, you have won $20 dollars towards your next purchase of edible goods. This money has been donated by APPLES Co., a non profit organization founded to promote the purchase of organic foods. These $20 will be applicable in any local grocery supermarket. You will receive it in the form a gift card that will be sent to your mailing address. In the meantime, please click the link below to vote for APPLES Co. as the top 10 non-profit of the year in our region! < link >"*

Scarcity and Financial: *"You can save 20 percent on your next electric bill by filling out our online survey within the next three days. Your participation will allow us to provide Regional Utilities with accurate information as to how they can improve their services. Take advantage of this limited-time opportunity by clicking the link below: < link >"*

Social Proof and Social: *"I would like to invite you to join the thousands of other clients who have experienced our vacation excursions. We are currently serving ten clients in your neighborhood and this is how we got your contact information. Our company called Dunes was developed for clients to experience lavish vacation spots at an affordable price. We have earned the title of Top 10 Best Travel Agency from TripAdvisor. To learn more on how you can start planning your dream vacation, visit our website at < link >."*

DATA ANALYSIS AND RESULTS

This section presents the analyses and results. We used the statistical software package STATA 14.0 for data analysis.

Susceptibility to spear phishing emails was operationalized as clicking on the email link provided in each of the 21 emails each user received during the study period. There was no relationship between the participant's click on a link and the day in the study. A large number of users showed susceptibility to the phishing attacks – 43% clicked on at least one of the spear phishing email links, and 11.9% clicked

on more than one email link. These findings support the effectiveness of the spear phishing emails used in our study.

Research question (i): Do younger and older Internet users differ in their susceptibility to spear phishing email attacks?

To address this research question, we conducted a multilevel logistic regression, in which the click status of each email link (i.e., 0 = link was not clicked, 1 = link was clicked) was the dependent variable (dichotomous) and participant age (0 = young, 1 = older) and gender (0 = male, 1 = female) were the independent variables. Older adults' probability of clicking emails links (3.2%) was numerically higher than that of young users (2.9%), with a significant age group by gender interaction, in that older women were particularly susceptible to spear phishing attacks ($B = .98$, $z = 2.02$, $p = .04$).

Research question (ii): Which weapon(s) is/are particularly effective?

Among all the recorded clicks on spear phishing email links during the study interval, scarcity with 26.0%, and authority with 21.9% email link clicks were the most effective weapons. This was followed by perceptual contrast with 16.7%, liking with 11.5%, reciprocation with 10.4%, commitment with 7.3%, and social proof with 6.3% clicks.

To determine significant differences in the effectiveness of weapons, we conducted a multilevel logistic regression, in which the click status of each email link was the dependent variable and participant age, weapon, and their interaction were independent variables. The effect of weapon was significant ($B = -.34$, $z = -4.79$, $p < .001$). Simple effect analysis showed that scarcity (5.3%) was significantly more effective than reciprocation (2.3%), liking (2.1%), commitment (1.5%) and social proof (1.3%). Further, authority (4.4%) was significantly more effective than commitment, and social proof. Perceptual contrast (3.4%) was significantly more effective than social proof (Figure 2).

Research question (iii): Does effectiveness of weapons vary by age group?

The interaction between participant age and weapon was also significant ($B = .20$, $z = 2.03$, $p = .04$). As shown in Figure 2, younger, compared to older users were more likely to click on the links in emails that used scarcity. In contrast, older compared to younger users were more likely to click on links in reciprocation emails; and, to a somewhat lesser extent in liking emails.

Research question (iv): Which life domain(s) is/are particularly effective?

Among all the recorded clicks, legal with 38.5% and ideological with 18.8% email link clicks were the most effective domains. This was followed by health with 14.6%, social with 11.5%, security with 10.4%, and financial with 6.3% clicks.

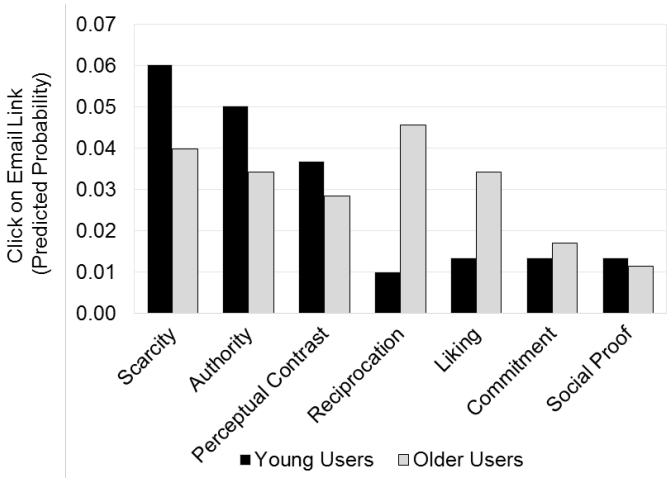

Figure 2. Susceptibility to spear phishing email attacks for various weapons by age group: Predicted probability of click on email link.

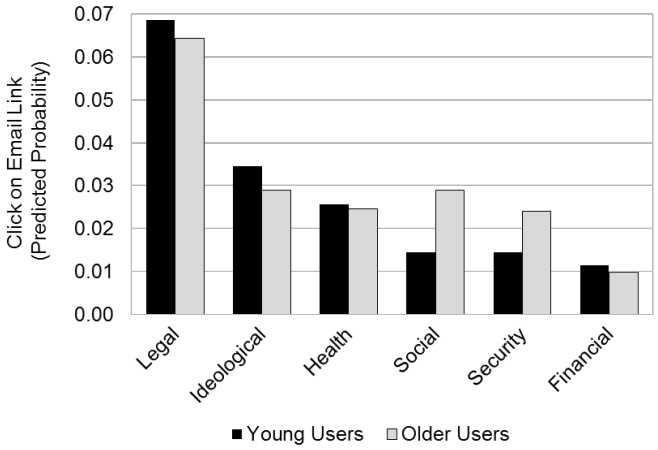

Figure 3. Susceptibility to spear phishing attack for various life domains by age group: Predicted probability of click on email link.

To determine significant differences in the effectiveness of life domains, we conducted a multilevel logistic regression, in which the click status of each email link was the dependent variable and participant age, life domain, and their interaction were independent variables. The effect of life domain was significant ($B = -.41$, $z = -4.91$, $p \leq .001$). Simple effect analysis showed that emails pertaining to the legal domain (6.7%) were significantly more effective than emails from all the other domains. Further, ideological emails (3.3%) were significantly more effective than financial emails (1.1%) (Figure 3).

Research question (v): Does effectiveness of life domains vary by age group?

The interaction between participant age and life domain was not significant ($B = .12$, $z = .93$, $p = .35$), suggesting that the effectiveness of life domains was comparable across age groups.

Research question (vi): Are younger and older Internet users aware of their susceptibility to spear phishing attacks?

Susceptibility awareness was operationalized as the mean of study users' self-reported likelihood of clicking on spear phishing email links. On Day 21, participants were asked to indicate for 21 spear phishing emails that they had not been exposed to during the study period how likely they were to click on a link in the email (response scale: 1 = not at all to 5 = very much). We had data from 155 users for this analysis [5] Overall, participants indicated a low likelihood of clicking on spear phishing email links ($M = 2.16$, $SD = 1.35$), i.e., they reported a low susceptibility awareness. These findings stand in contrast to the quite high behavioral susceptibility to spear phishing emails observed during the study period.

To compare susceptibility awareness between younger and older Internet users, we conducted a multilevel regression, in which the likelihood rating of clicking on the link in an email was the dependent variable (continuous) and participant age, gender, and their interaction were independent variables. The effect of age was significant ($B = -.78$, $z = -2.11$, $p = .035$): younger users ($M = 2.30$, $SD = .92$) reported higher susceptibility awareness than older users ($M = 1.96$, $SD = .93$). The interaction between participant age and gender was not significant ($B = .36$, $z = .57$, $p = .57$). In addition, the effect of gender was not significant ($B = -.11$, $z = -.36$, $p = .72$).

To explore variations in susceptibility awareness by weapon, thereby considering age group, we conducted a multilevel regression, in which the likelihood rating of clicking on the link in an each email was the dependent variable and participant age, weapon, and their interaction were independent variables. The effect of weapon was significant ($B = -.06$, $z = -4.91$, $p < .001$) (Figure 4). Simple effect analysis showed that susceptibility awareness was significantly higher for scarcity ($M = 2.27$, $SD = 1.44$) and authority ($M = 2.43$, $SD = 1.48$) than for perceptual contrast ($M = 2.06$, $SD = 1.36$), reciprocation ($M = 2.11$, $SD = 1.30$), and social proof ($M = 1.84$, $SD = 1.16$). In addition, susceptibility awareness for authority was higher than for liking ($M = 2.20$, $SD = 1.38$) and commitment ($M = 2.20$, $SD = 1.26$). The interaction between age group and weapon was not significant ($B = .01$, $z = .51$, $p = .61$), suggesting that the susceptibility awareness for various weapons was comparable across age groups. Interestingly, older, compared to young users, did not show a particular susceptibility awareness for emails using liking and reciprocation as the behavioral susceptibility data would have suggested.

To explore variations in susceptibility awareness across domains, thereby considering age group, we conducted a multilevel regression, in which the likelihood rating of clicking on the link of an email was the dependent variable and participant age, domain, and their interaction were independent variables. Neither the effect of domain ($B = -.03$, $z = -1.66$, $p = .097$) nor the interaction between age and domain ($B = -.02$, $z = -.61$, $p = .54$) were significant.

[5]Three young participants did not complete the final day survey.

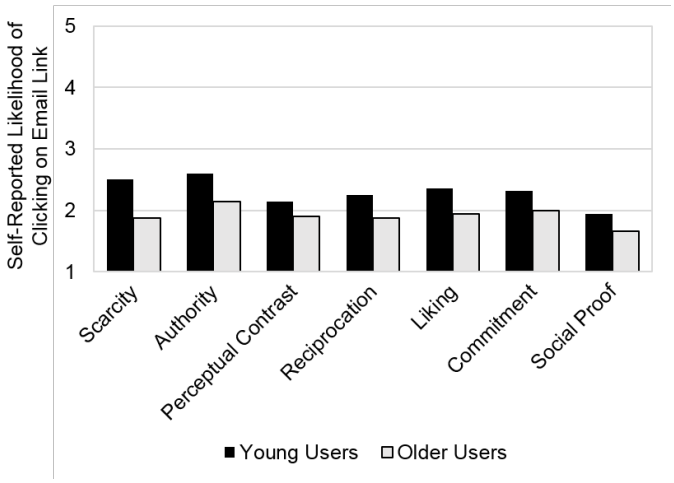

Figure 4. Susceptibility awareness for various weapons by age group: Mean likelihood rating of click on email link (1 = not at all, 5 = very much).

At study closure, we debriefed participants on the purpose of the study. Debriefing showed that only nine younger and three older adults (8% of the participants based on 143 collected responses pertaining to this item) were suspicious of the true study intent. Our results were comparable when data from these twelve participants was excluded.

Limitations of the Study Design

Participants knew they were being observed, as part of participation in a research study, even though they did not know the true purpose of the study. This may have affected their behavior, in causing lower/higher Internet usage or may have lowered their defenses. To counter this possible effect, in the attempt to not increase their Internet exposure because of study participation, we only recruited participants into the study who used the Internet (browsing and e-mail) daily. Also our emails controlled for a variety of confounding factors that could have influenced attack effectiveness – the emails limited both elements that may have contributed to their persuasiveness and elements that would have detracted from the attack, such as bad grammar and spelling, established online domains, and lack of specificity to the victim's interests. Nonetheless, we are cognizant that unobtrusive assessment would have been preferable and would have further optimized the ecological-validity of the study.

We propose for future research to correlate email volume (*e.g.*, average number of emails a user received per day) with susceptibility to phishing. This information was not assessed in our study and, thus, we were not able to submit it as covariate in our analysis.

Summary

Our study showed that Internet users were highly susceptible to spear phishing email attacks, as more than 40% of the participants in our study clicked on at least one email during the 21-day study period.

Older women were the most susceptible group to phishing attacks. Scarcity and authority were the most effective weapons of influence for all age groups. However, while younger users were most susceptible to scarcity, older adults were most susceptible to reciprocation.

All users were most susceptible to emails pertaining to the legal domain. Surprisingly, in our study the financial domain was not very effective. It is possible that educational efforts in everyday life outside of our study raise awareness about scams involving money, knowledge that our participants may have transferred into the study. It may also be that spear phishing in everyday life typically refers to the domain of finances and thus users have developed greater immunity to such attacks based on prior (negative) experience with such contents.

Also, the health life domain, even though increasingly important in old age [6], was not particularly effective in older adults. It is possible that older adults frequently hear from friends and family anecdotes about medication offers via email (*e.g.*, weight loss pills and Viagra), which may have contributed to the low appeal of this domain in our study.

An intriguing finding from our study is the great discrepancy, especially among older users, between actual behavioral susceptibility to spear phishing emails and participants' self-reported susceptibility awareness. This unawareness makes older users a particular at-risk population, which needs to be targeted in training and educational efforts.

Recommendations

Today's security training and warning, for phishing emails and websites, and SSL issues come as *one size fits all* solutions. However, our research shows that **one size does not fit all**. Rather, security communication with Internet users should be tailored to the specifics of their demographics. Further, training, warning, and security education for Internet users attempt to accomplish too much, and are not effective because people follow heuristics, forget, and habituate. Based on our findings, we advocate **targeted security interventions** (e.g., training and warnings), with a focus on the most relevant vulnerabilities of a given population of Internet users. For example, our research showed that, **for email phishing, younger users were most susceptible to scarcity, while older users were most susceptible to reciprocation**. We believe that an age-tailored training and prevention approach will increase the effectiveness of security measures because an age-targeted solution will impose less requirements on people and will match their specific vulnerabilities.

In this study we compared older and younger adults. However, teenagers and middle-age adults might show specific vulnerabilities, which need to be identified in future research. Further, a person background, such as education level, profession, or culture/country might influence user susceptibility to social engineering attacks. Our study is the first to point out these differential susceptibility patterns and suggests the need for targeted security warnings, training and education. We hope that our findings will spur future investigation pertaining to a variety of demographics.

In sum, the main contribution of this work is a better understanding of various factors, such as age, gender, weapons of influence and life domains involved in the effectiveness of spear-phishing attacks. Future work is warranted that integrates research in user interface (UI) design, artificial intelligence, human factors, and cognitive sciences to develop effective prevention and mitigation tools.

CONCLUSION

This paper investigated susceptibility to spear phishing email attacks as a function of Internet user age, gender, weapon of influence, and life domains these emails referred to. In addition to behavioral effects, we also examined self-reported susceptibility awareness. We adopted an ecologically valid research design, by examining younger and older Internet users in their homes during daily Internet activity. Over 21-days, 158 Internet users were exposed to experimentally controlled spear phishing emails. Our results suggest that **older women were the most susceptible group to spear phishing email attacks**. While **younger users were most susceptible to scarcity, older users were most susceptible to reciprocation; authority was highly effective in both age groups and both age groups**. Further, we found an intriguing **discrepancy between susceptibility shown behaviorally and self-reported susceptibility awareness**; this effect was present in the total sample but particularly pronounced in older users.

Results from this study highlight that training and warning solutions should be age-targeted, in accordance with specific vulnerabilities of individual's demographics. While our study demonstrates the extent to which younger and older adults differ in their susceptibility to spear phishing attacks, future studies are warranted to determine differential susceptibility patterns across various demographics, including different age groups, educational levels, and cultures. The take-home message of this research is that **one-size does not fit all**. We propose that applying our study's insights to the development of the next generation of defense solutions for Internet users will increase user compliance and the effectiveness of the security measures.

ACKNOWLEDGMENTS

We thank the CHI anonymous reviewers and shepherd for their insightful comments and guidance. We also thank Aliye Karakoyun, Dinia Salmeron, Marvis Cruz, Sebastian Marin, Andrew Varan, Robert Rainer, Cheyenne Reynolds, Hannah Burrichter, Paul Talty, Nicole Phillips, Sami Winder, and Andrea Alonzo for their contributions to this project. This research was supported by NSF grant SBE-1450624.

REFERENCES

1. L. Chan A. E. Reed and J. A. Mikels. 2014. Meta-Analysis of the Age-Related Positivity Effect: Age Differences in Preferences for Positive Over Negative Information. In *Psychology and Aging*. 1–15.

2. AARP. 2012. Healthy Aging Improving and Extending Quality of Life Among Older Americans. (2012). `http://www.aarp.org/livable-communities/learn/health-`

wellness/info-12-2012/Healthy-Aging-Improving-and-Extending-Quality-of-Life-Among-Older-Americans-2011.html

3. Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (1999).
4. Eirik Albrechtsen. 2007. A Qualitative Study of Users' View on Information Security. *Computers and Security* 26, 4 (2007).
5. Farzeneh Asgapour, Debin Liu, and L. Jean Camp. 2007. Mental Models of Computer Security Risks. *Financial Cryptography and Data Security Lecture Notes in Computer Science* 4886 (2007), 367–377.
6. Paul B. Baltes, Ulman Lindenberger, and Ursula M. Staudinger. 2007. Life Span Theory in Developmental Psychology. *Wiley Online Library* (2007).
7. Vanessa Boothroyd. 2014. *Older adults' Perception of Online Risk*. Master's thesis. Carleton University.
8. J. Brandt, M. Spencer, and D. R. Davies. 1988. The telephone interview for cognitive status. *Neuropsychiatry, Neuropsychology, & Behavioral Neurology* 1 (1988), 111–117.
9. D.D. Caputo, S.L. Pfleeger, J.D. Freeman, and M.E. Johnson. 2014. Going Spear Phishing: Exploring Embedded Training and Awareness. *IEEE Security & Privacy* 12, 1 (2014), 28–38.
10. Jeffrey Carr. 2011. *Cyber Warfare*. O'Reily.
11. E. Castle, N. I. Eisenberger, T. E. Seeman, W. G. Moons, I. A. Boggero, M. S. Grinblatt, and S. E. Taylor. 2012. Neural and Behavioral Bases of Age Differences in Perceptions of Trust. In *Proceedings of the National Academy of Sciences*, Vol. 109. 20848–20852.
12. US Census. 2010. United States Census 2010. (2010). http://www.census.gov/2010census/
13. Robert B. Cialdini. 2006. *Influence - The Psychology of Persuasion*. Collins Business Essentials.
14. Cisco. 2011. Email Attacks: This Time It's Personal . (2011). http://itknowledgeexchange.techtarget.com/security-detail/cisco-report-email-attacks-this-time-its-personal/
15. Rachna Dhamija, J. D. Tygar, and Marti Hearst. 2006. Why Phishing Works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 581–590. DOI: http://dx.doi.org/10.1145/1124772.1124861
16. Julie S. Downs, Mandy Holbrook, and Lorrie Faith Cranor. 2007. Behavioral Response to Phishing Risk. In *Proceedings of the Anti-phishing Working Groups 2Nd Annual eCrime Researchers Summit (eCrime '07)*. ACM, New York, NY, USA, 37–44. DOI:http://dx.doi.org/10.1145/1299015.1299019
17. Julie S. Downs, Mandy B. Holbrook, and Lorrie Faith Cranor. 2006. Decision Strategies and Susceptibility to Phishing. *Symposium on Usable Privacy and Security (SOUPS)* (2006).
18. N. C. Ebner, P. E. Bailey, M. Horta, and J. Joiner. 2015a. *Multidisciplinary Perspective on Prosociality in Aging. (Invited book chapter)*. Sommerville & J. Decety.
19. N. C. Ebner, P. E. Bailey, M. Horta, J. Joiner, and S. W. C. Chang. 2015b. *Multidisciplinary perspective on prosociality in aging. In (Eds.), in Social Cognition for the Frontiers in Developmental Science Series (Psychology)*. J. Sommerville & J. Decety.
20. Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1065–1074.
21. FBI. 2016. Fraud against Seniors. (2016). https://www.fbi.gov/scams-and-safety/common-fraud-schemes/seniors
22. Ian Fette, Norman Sadeh, and Anthony Tomasic. 2007. Learning to detect phishing emails. In *Proceedings of the 16th international conference on World Wide Web*. ACM, 649–656.
23. Baruch Fischhoff, Paul Slovic, Sarah Lichtenstein, and Barbara Combs Stephen Read. 1978. How Safe is Safe Enough? A Osychometric Study of Attitudes Towards Technological Risks and Benefits. *Policy Sciences* 9, 2 (1978).
24. V. Garg and L. Jean Camp. 2012a. End User Perception of Online Risk Under Uncertainty. *Hawaii International Conference On System Sciences* 4886 (2012).
25. V. Garg and L. Jean Camp. 2012b. Risk Communication Design for Older Adults. *Gerontechnology* 11, 2 (2012).
26. Christopher Hadnagy. 2010. *Social Engineering: The Art of Human Hacking*. Wiley.
27. Ding-Long Huang, Pei-Luen, Patrick Raua, Gavriel Salvendya, Fei Gaoa, and Jia Zhoua. 2011. Factors Affecting Perception of Information Security and Their Impacts on IT Adoption and Security Practices. *International Journal of Human-Computer Studies* 69, 12 (2011).
28. Lance James. 2006. *Phishing Exposed*. Syngress.
29. Mitzi Johnson. 1990. Age Differences in Decision Making: A Process Methodology for Examining Strategic Information Processing. *Journal of Gerontology: Psychological Sciences* 45, 2 (1990), 75–78.
30. Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
31. Ponnurangam Kumaraguru. 2009. *Phishguru: a system for educating users about semantic attacks*. ProQuest.

32. Ponnurangam Kumaraguru, Justin Cranshaw, Alessandro Acquisti, Lorrie Cranor, Jason Hong, Mary Ann Blair, and Theodore Pham. 2009. School of phish: a real-world evaluation of anti-phishing training. In *Proceedings of the 5th Symposium on Usable Privacy and Security*. ACM, 3.

33. Ponnurangam Kumaraguru, Yong Rhee, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 2007. Protecting people from phishing: the design and evaluation of an embedded training email system. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 905–914.

34. Ponnurangam Kumaraguru, Steve Sheng, Alessandro Acquisti, Lorrie Faith Cranor, and Jason Hong. 2010. Teaching Johnny not to fall for phish. *ACM Transactions on Internet Technology (TOIT)* 10, 2 (2010), 7.

35. P. A. M. Van Lange. 2015. Generalized Trust: Four Lessons From Genetics and Culture. *Current Directions in Psychological Science* 24, 1 (2015), 71–76.

36. T. Lauricella. 2014. If you're over 50, you're a scam target. The Wall Street Journal `http://www.wsj.com/articles/if-youre-over-50-youre-a-scam-target-1412467756`. (2014). `http://www.wsj.com/articles/if-youre-over-50-youre-a-scam-target-1412467756`

37. Gang Liu, Guang Xiang, Bryan A Pendleton, Jason I Hong, and Wenyin Liu. 2011. Smartening the crowds: computational techniques for improving human verification to fight phishing scams. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*. ACM, 8.

38. Rui Mata, Anika Josef, Gregory Samanez-Larkin, and Ralph Hertwig. 2011. Age Differences in Risky Choice: A Meta-Analysis. *New York Academy of Sciences* (2011).

39. Mara Mather. 2006. *When I'm 64 - A Review of Decision-Making Processes: Weighing the Risks and Benefits of Aging*. The National Academies Press.

40. Kevin Mitnick, W. L. Simonand, and S. Wozniak. 2002. *The Art of Deception: Controlling the Human Element of Security*. Wiley.

41. E. Peters, M. A. Diefenbach, T. M. Hess, and D. Vastfjall. 2008. Age Differences in Dual Information-Processing Modes: Implications for Cancer Decision Making. *Cancer* 113 (2008), 12.

42. R. Petrican, T. English, J. J. Gross, C. Grady, T. Hai, and M. Moscovitch. 2013. Friend or foe? Age Moderates Time-Course Specific Responsiveness to Trustworthiness Cues. *The Journals of Gerontology Series B, Psychological Sciences and Social Sciences and Social Sciences* 68, 2 (2013), 215–223.

43. Threat Post. 2011. RSA: SecurID Attack Was Phishing Via an Excel Spreadsheet . (2011). `https://threatpost.com/rsa-securid-attack-was-phishing-excel-spreadsheet-040111/75099/`

44. Consumer Reports. 2015. Lies, Secrets, and Scams: How to Prevent Elder Abuse . (2015). `http://www.consumerreports.org/cro/consumer-protection/preventing-elder-abuse`

45. T. Ruffman, J. Murray, J. Halberstadt, and T. Vater. 2012. Age-related Differences in Deception. In *Psychology and Aging*, Vol. 27. 543–549.

46. T. Ruffman, S. Sullivan, and N. Edge. 2006. Differences in the Way Older and Younger Adults Rate Threat in Faces But Not Situations. In *The Journals of Gerontology Series B, Psychological Sciences and Social Sciences and Social Sciences*, Vol. 61. 187–194.

47. Gregory R Samanez-Larkin. 2013. Financial decision making and the aging brain. *APS observer* 26, 5 (2013), 30.

48. Gregory R Samanez-Larkin and Brian Knutson. 2015. Decision Making In The Ageing Brain: Changes In Affective And Motivational Circuits. *Nature reviews. Neuroscience* (2015).

49. Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. 2010. Who falls for phish?: a demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 373–382.

50. Steve Sheng, Bryant Magnien, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. 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*. ACM, 88–99.

51. Steve Sheng, Brad Wardman, Gary Warner, Lorrie Faith Cranor, Jason Hong, and Chengshan Zhang. 2009. An empirical analysis of phishing blacklists. (2009).

52. P.W. Singer and Allan Friedman. 2014. *Cybersecurity and Cyberwar: What Everyone Needs to Know*. Oxford University Press.

53. Symantec. 2016. Symantec Internet Security Threat Report. (2016). `https://www.symantec.com/security-center/threat-report`

54. Katya Tentoria, Daniel Oshersonb, Lynn Hasherc, and Cynthia May. 2001. Wisdom and Aging: Irrational Preferences in College Students But Not Older Adults. *Elsevier Science* (2001).

55. Netcraft Toolbar. 2009. Netcraft, Ltd. (2009).

56. Paul Verhaeghen and Timothy A. Salthouse. 1997. Meta-Analyses of Age-Cognition Relations in Adulthood: Estimates of Linear and Nonlinear Age Effects and Structural Models. *Psychological Bulletin* 122, 3 (1997), 231–249.

57. Tim Vidas, Emmanuel Owusu, Shuai Wang, Cheng Zen, and Lorrie Faith Cranor. 2012. QRishing: The Susceptibility of Smartphone Users to QR Code Phishing Attacks. *CMU-CyLab-12-022* (2012).

58. S. J. Westerman and D. R. Davies. 2000. Acquisition and Application of New Technology Skills: The Influence of Age. *Occup. Med.* 50 (2000), 1.

59. Tyler Wrightson. 2014. *Advanced Persistent Threat Hacking: The Art and Science of Hacking Any Organization*. McGraw-Hill Education.

60. Min Wu, Robert C Miller, and Simson L Garfinkel. 2006. Do security toolbars actually prevent phishing attacks?. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 601–610.

61. Yue Zhang, Serge Egelman, Lorrie Cranor, and Jason Hong. 2006. Phinding phish: Evaluating anti-phishing tools. ISOC.