Crime Science

RESEARCH

Open Access

Empirical analysis of weapons of influence, life domains, and demographic-targeting in modern spam: an age-comparative perspective

Daniela Seabra Oliveira[1*], Tian Lin[2], Harold Rocha[2], Donovan Ellis[2], Sandeep Dommaraju[3], Huizi Yang[1], Devon Weir[2], Sebastian Marin[2] and Natalie C. Ebner[2]

Abstract

Spam has been increasingly used for malware distribution. This paper analyzed modern spam from an age-comparative perspective to (i) discover the extent to which psychological weapons of influence and life domains were represented in today's spam emails and (ii) clarify variations in the use of these weapons and life domains by user demographics. Thirty five young and 32 older participants forwarded 18,605 emails from their spam folder to our study email account. A random set of 961 emails were submitted to qualitative content coding and quantitative statistical analysis. Reciprocation was the most prevalent weapon; financial, leisure, and independence the most prevalent life domains. Older adults received health and independence-related spam emails more frequently, while young adults received leisure and occupation-related spam emails more often. These age differences show a level of targeting by user demographics in current spam campaigns. This targeting shows the need for age-tailored demographic warnings highlighting the presence of influence and pretexting (life domains) for suspicious emails for improved response to cyber-attacks that could result from spam distribution. The insights from this study and the produced labeled dataset of spam messages can inform the development of the next generation of such solutions, especially those based on machine learning.

Keywords: Spam, Influence, Life domains, Older adults, Young adults, Targeting

Introduction

The classic definition of spam is unsolicited and undesired email messages to advertise products (Kanich et al. 2008, 2011; Stone-Gross et al. 2011; Stringhini et al. 2014). However, recent security reports have documented that spam has been increasingly used to distribute malware (e.g., ransomware) or to attempt to lure Internet users into falling for scams (Wong and Solon 2017; Symantec 2017). Even though a great number of spam is blocked by filters implemented by email providers and institutions, a number of messages evade detection on a daily basis. For example, recent reports document that Internet users receive, on average, 117 emails per day and that 53% of such emails are spam (Symantec 2017). Thus, even considering current spam filters' blocking rate of over 90%, end users will still experience at least a few spam emails reaching their inbox on any day. This is the case, because email filters are usually based on machine learning classification, which has limitations, such as their high dependency on good and up-to-date training sets. Of note, it only takes one click from a user on a malicious link in a message for their machine to be compromised. This can cause immense negative consequences for the individual, such as his credentials being stolen or malware being installed on his computer. Also, if such infection happens in a corporate

*Correspondence: daniela@ece.ufl.edu
[1] Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL, USA
Full list of author information is available at the end of the article

environment, it could represent the infiltration stage of an APT attack.

There has been a wealth of research on various aspects of spam, from detection/mitigation via machine learning (Drucker et al. 1999; Meyer and Whateley 2005; Sculley and Wachman 2007; Hao et al. 2009; Ramachandran et al. 2007; Stringhini et al. 2011) to analysis of spam botnets (Stringhini et al. 2011; Kanich et al. 2008) and the spam economical ecosystem (Kanich et al. 2011, 2008; Stone-Gross et al. 2011; Stringhini et al. 2014). However, as spam has evolved over the years as a mechanism for malware distribution, novel research questions have arisen. Among those questions are: (i) to what extent are psychological weapons of influence and specific life domains, as techniques to lure users into reading the spam email and/or clicking on its URLs or downloading attachments, represented in today's spam emails?; (ii) which weapons of influence and life domains are most popular?; and (iii) does the use of weapons of influence and life domains vary by user age group (young vs older)?

In an attempt to answer these questions we conducted a user study with 35 young (18–32 years) and 32 older (61–88 years) men and women who regularly use the Internet. We adopted an extreme-group cross-sectional design by contrasting young and older adults, a parsimonious methodological approach often applied in aging research (Verhaeghen 2003; Mata et al. 2011a; Reed et al. 2014). Study participants were asked to forward to our study email account the entire content of their spam/junk folders. In total, 18,605 emails (10,213 from young and 8392 from older users) were collected. A random set of 961 emails (514 from young users and 447 from older users) was selected for manual qualitative content coding by trained, independent coders, and submitted to subsequent quantitative statistical analysis.

Our approach extends previous work (Stringhini et al. 2014; Kanich et al. 2008; Edwards et al. 2015) in its adoption of an analysis of spam from an age-comparative perspective. This perspective allowed us to determine the extent to which spammers in today's spam emails target young vs older users differently, based on their particular vulnerabilities (Oliveira et al. 2017).

Investigation of older Internet users is an emerging topic with growing high relevance from a security standpoint, given that this age group controls over half of the US financial wealth and occupy many positions of power in politics, business, and finance. Older adults constitute a particular at-risk population for email-based attacks (Oliveira et al. 2017). This particular vulnerability may be a consequence of general deficits in cognitive processing capacities and reduced sensitivity to deception in advanced age (Verhaeghen and Salthouse 1997; Mather 2006; Johnson 1990; Mata et al. 2011b; Tentoria et al. 2001), (https://www.fbi.gov/scams-and-safety/common-fraud-schemes/seniors; http://www.wsj.com/articles/if-youre-over-50-youre-a-scam-target-1412467756).

The manual qualitative content coding process considered (i) seven weapons of influence (reciprocation, liking, scarcity, social proof, authority, commitment, perceptual contrast) (Cialdini 2006; Hadnagy 2010) and (ii) 16 life domains under two categories: six incentivizing domains (financial, health, ideological, social, legal, security) (Hadnagy 2010) and ten contextual domains (health, cognitive, independence, family, friends, life reflections, death, occupation, leisure, sexuality) (Schindler et al. 2006).

Quantitative statistical analysis of the content-coded emails showed that all seven weapons of influence, all six incentivizing life domains, and all ten contextual life domains were represented in current spam emails, suggesting that spammers currently use these techniques. Reciprocation was the most frequently used weapon, financial the most prevalent incentivizing domain, and leisure and independence the most popular contextual domains in today's spam.

Older adults were more likely to receive spam emails related to health and independence, while young adults were more likely to receive spam emails related to leisure and occupation. We found no age difference, however, regarding the use of specific weapons of influence.

Our study suggests a level of (age-specific) targeting in today's spam campaigns. This knowledge is crucial for the development of the next generation of spam mitigation solutions, such as regarding use of machine learning-based content analysis, detection of influence in text, and warning solutions that consider demographic-specific susceptibilities.

In summary, in this paper we empirically analyzed the extent to which Internet users from two distinct age groups (i.e., young vs older adults) are currently targeted in a demographic-specific fashion by spam campaigns. In this analysis we focused on weapons of influence used in emails and life domains emails refer to, serving as pretext of the message. This analysis is important for the following reasons: First, it unveils characteristics of current spam campaigns and provides insights about whether and how spammers are targeting spam recipients by their demographics (e.g., age) and, thus, advances scientific knowledge about spam. Second, our focus on distinct age groups is of particular relevance and innovation, given the parsimonious previous research on usable security for older adults, a vulnerable and important segment of the world's population. Studying older adults in the context of computer security is currently still a research niche and results will importantly qualify knowledge and advance the field. Furthermore, results from our study can inform the design and implementation of the next generation of warning tools and awareness/educational training programs and guidelines. In particular, these future tools and training could more effectively target vulnerabilities of particular demographic groups. Also, our findings and our labeled dataset of spam messages have the potential to advance the development of natural language processing models to detect influence and life domains (pretexts) in text with the goal to build effective warning solutions for the Internet user.

This paper is organized as follows. Section "Related work" discusses related work. Section "Background" provides a background regarding the psychological weapons of influence and life domains employed in spam emails. Section "Methods" describes the study methodology. Section "Statistical analysis" details and discusses the results of the qualitative content coding and quantitative statistical analysis. Section "Conclusions" concludes the paper.

Related work

In this section we discuss related work on spam mitigation and analysis and on phishing, as spam has been increasingly used for malware distribution (Symantec 2017).

Spam detection and analysis

The current literature has largely focused on determining whether a piece of email is spam or "ham" (benign email) (Drucker et al. 1999; Meyer and Whateley 2005; Sculley and Wachman 2007; Hao et al. 2009; Ramachandran et al. 2007; Stringhini et al. 2011; Xie et al. 2008; Stringhini et al. 2012; Schwartz 2004; Taylor 2006; Pitsillidis et al. 2010). The main approach is to analyze the content of emails using machine learning (Drucker et al. 1999; Meyer and Whateley 2005; Sculley and Wachman 2007), heuristics, and/or blacklists—for example, analyzing email sender IP addresses (Hao et al. 2009; Ramachandran et al. 2007; Stringhini et al. 2011), URLs used in the email (Xie et al. 2008), or network features (Stringhini et al. 2012). While such approaches are crucial for mitigating spam, and many of them are applied in commercial email servers (Schwartz 2004), they have limitations, such as performance requirements and false positives and negatives given the dynamic nature of spam.

Spam is usually sent by botnets, hired by spammers for their spam campaigns. Thus, many studies have focused on analyzing (i) the operation of such botnets (Stringhini et al. 2011; Stone-Gross et al. 2011); (ii) how the botnets automatically generate content for spam campaigns (Kanich et al. 2008); and (iii) the statistical features of large corpuses of spam (Edwards et al. 2015).

There is also a large and profitable underground economy fueled by spam. Research, therefore, has focused on economic aspects of the spam ecosystem from the financial conversion of spam (Kanich et al. 2008) over the spam product life-cycle (Kanich et al. 2011), to the relationships between actors (email harvesters, spam content generators and botnet masters) in this ecosystem (Stringhini et al. 2014).

Most closely related to our work is a study by (Redmiles et al. 2018), which investigated what drives users to click on social media spam. The study found that the spam topic was one of the most significant features in driving click behaviors, and women compared to men were more likely to click on social media spam.

Our work set out to analyze modern spam from the perspective of the Internet user to discover how today's spam received by Internet users in their everyday life targets end users, and specifically how it targets young and older users differentially.

Phishing

The current literature on phishing has focused on understanding what makes phishing attractive, why people fall for phishing, and on protecting users against phishing attacks (e.g., detecting phishing and educating users) Fette et al. (2007), (https://toolbar.netcraft.com/), Zhang et al. (2006), Sheng et al. (2009), Downs et al. (2006), Ferreira and Lenzini (2015), Uebelacker and Quiel (2014), Oliveira et al. (2017), Benenson et al. (2017). Automatic detection of phishing is challenging because phishing has become more targeted, thus creating difficulties in determining good features for machine learning classification. User education initiatives (Sheng et al. 2007; Kumaraguru 2009; Kumaraguru et al. 2007, 2010) are also challenging because people tend to forget what they learned after some time and fall for the same attacks shortly after training (Caputo et al. 2014).

Our work complements prior research on spam and phishing in that it takes a closer look at specific content of today's spam emails, with a particular focus on use of weapons of influence and life domains (Hadnagy 2010; Schindler et al. 2006). In this context, we did not analyze email messages that traversed an ISP domain, a honeypot, or a botnet, but rather those spam messages that Internet users had actually received in their spam/junk folders as part of their everyday Internet activities. Our analysis specifically considered users age (young vs older), to determine age-specific targeting in modern spam.

Principles of influence in email

Workman conducted an early empirical study of weapons of influence in social engineering (Workman 2007). His framework categorized weapons of influence as (i) commitment, reciprocation, and social proof, (ii) likability and trust, and (ii) authority, scarcity, and fear.

His framework stressed the potential impact of cultural biases on the relevance of these principles to Internet users. For example, authority might be perceived differently in different countries that vary in social norms (e.g., in Japan, old age implies an aura of authority).

Observing the behavior of street hustlers, Stajano and Wilson (2011) extended and modified Cialdini's framework and proposed nine principles of influence: distraction, social compliance (authority), herd (social proof), dishonesty, kindness, need and greed (e.g., visceral triggers), scarcity (time), commitment, and consistency. Uebelacker and Quiel (2014) analyzed 207 phishing emails following Cialdini's framework and constructed relations between personality traits of the Five-Factor Model (Big 5) and the principles of influence (Cialdini 2006). Ferreira and Lenzini (2015) studied the relationships and similarities between Cialdini (2006), Gragg (2003), and Stajano and Wilson (2011) frameworks and consolidated the principles of influence into five categories: (i) authority, (ii) social proof, (iii) liking/similarity/deception, (iv) commitment/consistency, and (v) distraction.

Akbar (2014) performed a quantitative analysis on suspected phishing emails collected from an institution in The Netherlands and found that authority and scarcity were disproportionately the most employed principles, followed by liking, consistency, reciprocation, and social proof. Considering different demographics (i.e., age and gender), Oliveira et al. (2017) conducted an empirical study comparing young vs older adults' susceptibility to Cialdini's principles of influence. Older women were the most susceptible group, and while younger adults were most susceptible to scarcity, older adults were most susceptible to reciprocation, and both groups were highly susceptible to authority.

Ortiz (2010) studied machine learning techniques to detect persuasion in negotiation transcripts. His classifier considered two classes for the dialogues: persuasive or not persuasive. Ortiz reports that his results provide a weak indication that these two classes can be distinguished. Moving forward, we plan to extend Ortiz's approach by distinguishing, via machine learning methods (including Natural Language Processing), each of Cialdini's principles of persuasion.

Background

Psychological principles of influence (called weapons in this study to emphasize their deceptive usage) are persuasive arguments used to compel recipients to perform an action that benefits the persuasive party. Cialdini (2006) described six such weapons of influence: reciprocation, liking, scarcity, social proof, authority, and commitment. A seventh weapon, perceptual contrast, was added based on Hadnagy (2010).

According to the Reciprocation principle people tend to repay, in kind, what another person has provided them. As an example, a spam message can convince a user into clicking on a link or replying to a message by offering the user a free gift attached to the email (e.g., the pdf of a travel guide to France). The travel guide might be of relevance to the user, who might feel indebted to the sender and think that the least he could do is to open the pdf, which can be malicious and infect his computer. The Liking principle is based on people's tendency to comply with requests from people they like or with whom they share similarities. The Scarcity principle is based on people perceiving opportunities as more valuable when their availability is limited. The Social proof principle states that people tend to avoid mistakes by acting like others. According to the Authority principle, people tend to feel at ease complying with requests made by "figures of authority", e.g., law enforcement personnel and lawyers (Hadnagy 2010; Mitnick et al. 2002). The Commitment principle states that people feel pressured to behave in line with their commitments. The Perceptual contrast principle refers to humans noticing a drastic difference between two situations or offers. When the second offer/situation is rather worse than the first, people tend to perceive the first as much more appealing. Spam emails can use these weapons as techniques to lure users into clicking on embedded malicious links or opening malicious attachments.

Spam emails can also refer to particular life domains [incentivizing (Hadnagy 2010) and contextual (Schindler et al. 2006)] to increase their appeal.

Incentivizing life domains refer to a category of information that might motivate users to attend to the spam email because they find it relevant and potentially beneficial to a particular aspect of their lives. We considered the following incentivizing life domains: financial, health, ideological, social, legal, and security. Financial emails focus predominantly on money, rebates, or offers. Health emails address mental and physical wellness, e.g., medication offers. Ideological emails relate to code of ethics and principles, e.g., an invitations to support a social cause. Social emails focus on interpersonal interactions, e.g., community events. Legal emails refer to the law, such as emails discussing a potential infraction. Security emails relate to physical or cyber safety, e.g., antiviruses offers.

Contextual life domains represent essential benefits, explicitly stated or more implicitly implied in an email that could prompt or enforce social behaviors by the recipient. They represent general life themes or directives. We considered the following contextual life domains: health, cognitive, independence, family, friends, life reflections, death, occupation, leisure, and sexuality.

Health emails relate to physical fitness, e.g., advertisement of workout routines. Cognitive emails focus on the recipient's capabilities regarding life skills, attention, and memory, e.g., brain training offers. Independence emails relate to the recipient's life autonomy, e.g., an advertisement of emergency buttons for the elderly. Family emails address relationships to relatives, e.g., information about detecting mental issues in family members. Friends emails relate to meaningful social connections with non-family members. Life reflections emails refer to the recipient's personal narrative and ability to engage in meaningful pursuits, e.g., emails about finding meaning in life after retirement. Death emails relates to the recipient's mortality, such as life insurance. Occupation emails target recipient's profession. Leisure emails relate to recipient's hobbies and free time, e.g., an email about a dog training club. Sexuality emails address sexual identity and romantic relationships, e.g., dating websites.

Although incentivizing and contextual domains were both subsumed under the broad category of life domains, they are distinct in that contextual domains apply even where no persuasive benefit (i.e., incentive) is present.

Methods

Our study proposed to analyze current spam from two perspectives: that of the spammer (offender) and that of the victim. The offender perspective is captured in our focus on email characteristics (i.e., weapons of influence and life domains) as tools to lure the user into clicking on potentially malicious links because the techniques are effective in getting individuals to act upon the request and because the life domains relate to interests and motivations the individual may have. The victim perspective is captured by considering computer user characteristics (i.e., age) that are targeted in a specific manner by the offenders as key variables in our analysis.

In the current study we took an ecologically valid approach by acquiring spam from real-life Internet users. The particular focus of our study was on analyzing how Internet users of different ages are currently targeted by spam attacks. Thus, even though the messages we analyzed had been classified as spam and had been blocked by spam filters, they had been sent out by spammers to target users in specific ways. Notice that this is different from analyzing susceptibility to malicious messages, which is beyond the scope of this study and already covered in the literature (Oliveira et al. 2017).

The study comprised men and women from two age groups. Young participants [$n = 35$ (60% females; 40% White), $M = 21.09$ years ($SD = 3.34$; range = 18–32)] were undergraduates from the University of Florida and other young adults residents of Alachua county (where the university is located). Older participants [$n = 32$ (50% females; 88% White), $M = 69.51$ years ($SD = 6.82$; range = 61–88)] were residents of Alachua county. Participants were recruited through the university Psychology Subject Pool, HealthStreet[1], fliers disseminated online, throughout the community, and university- and lab-internal participant registries. Young participants who were recruited through the subject pool were compensated with course credit; all other participants were financially compensated (see details below). Young participants reported a mean of 14.07 ($SD = 3.88$) years of education and older participants a mean of 16.06 ($SD = 2.86$).[2] Table 1 details demographic information of the participants. Data collection occurred in the spring of 2015. All emails covered approximately one month of participants' spam emails on the months of April and May 2015.

Participants were instructed to forward to the study team all spam emails they had in their current spam folder (from their primary personal email account) at the time of study enrollment. We did not put restrictions on the email provider to avoid influencing the type of spam we received or introducing bias regarding provider-specific anti-spam techniques. We focused on spam received by the users. To increase ecological validity, we did not try to control for the type of spam filter/email reader users adopted. Many providers, such as Gmail, already attempt to classify spam into separate folders. The goal of this study was not to investigate the effectiveness of anti-spam mechanisms, but to increase understanding of the extent to which spam targeting varies by user age.

Procedure

Researchers obtained informed consent from all participants prior to enrollment. Following consent, participants were provided with a demographic survey and instructions on how to submit their spam emails.

Spam emails were collected from participants' spam/junk folders. Each participant had the option to either manually forward spam emails to our study email account or use a web-based extraction tool we had developed that gathered spam emails automatically using OAuth 2.0. Research staff informed participants, that their email inbox would not be compromised by permitting the research team access to extract the contents of their spam/junk inbox. Participants were compensated with $0.20 for every email the research team received, for up to a total of $20 in the form of a prepaid VISA card. To be eligible for study compensation and inclusion in the analysis, participants were required to

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

[2] Two young and one older participant did not indicate gender, race, and years of education.

Table 1 Demographic information by age group

	Young users (n = 35)	**Older users (n = 32)**	**Age differences**
	M (SD)/%	**M (SD)/%**	
Age (years)—(young n = 34, older n = 29)			
	21.82 (5.14)	69.52 (6.82)	$t(61) = 4.38, p = .001$
Gender			
Male	41.2.8	48.3	$\tilde{\chi}^2(1) = 32, p < .57$
Female	58.8	51.7	
Years of education (young n = 34, older n = 29)			
	14.07 (3.89)	16.07 (2.87)	$t(61) = 0.84, p = .03$
Highest degree earned (young n = 29, older n = 19)			
High school	65.5	31.6	$\tilde{\chi}^2(4) = 6.28, p < .18$
Associates	6.9	10.5	
Bachelor's	17.2	36.8	
Master's	10.3	15.8	
Doctorate	0.0	0.0	
PhD	0.0	5.3	
Annual income (young n = 32, older n = 28)			
< $40,000	87.5	57.1	$\tilde{\chi}^2(2) = 7.04, p < .03$
$40,000–$70,000	6.3	21.4	
> $70,000			
Race/ethnicity (young n = 32, older n = 28)			
American Indian or Alaskan Native	0.0	0.0	$\tilde{\chi}^2(4) = 20.76, p < .001$
Asian	21.9	0.0	
Black/African American	6.3	0.0	
Native Hawaiian	0.0	0.0	
Hispanic	18.8	0.0	
White	46.9	100	
Other	6.3	0.0	
Marital status (young n = 32, older n = 28)			
Single	56.3	14.3	$\tilde{\chi}^2(4) = 43.08, p < .001$
In a relationship, but not married	40.6	0.0	
Married	3.1	53.6	
Divorced/separated	0.0	25	
Widowed	0.0	7.1	
Preferred internet browser (young n = 32, older n = 28)			
Internet explorer	9.4	25	$\tilde{\chi}^2(4) = 9.12, p < .06$
Google chrome	68.8	42.9	
Safari	12.5	3.6	
Mozilla firefox	9.4	21.4	
Other	0.0	7.1	

Sample size varied across demographic variables as some data were missing/not reported for some participants

submit a minimum of 40 emails. We determined 40 as the minimum number of emails that needed to be sent by the participants to our team to justify study compensation and to allow a randomized selection process for the subset of manually coded emails. The larger set of emails we collected will be leveraged in future analyses using machine learning to identify weapons of influence in text. A total of 18,605 spam emails were collected: 10,213 emails from young and 8392 from older users.[3] The average number of emails forwarded to our research team per participant was 275, the maximum number of emails sent by a participant was 1680, and the minimum (required for inclusion in the study)

[3] These numbers only include emails from eligible participants.

was 40. The large majority of our participants used our tool (for non-Gmail providers) or the Gmail folder archive tool option to forward their spam. The forwarded emails were the most recent in the participants' spam folders at time study participation.

Spam emails were converted into HTML files, which were stored in a secured database on the study server to assure confidentiality. Prior to content coding, we converted emails back to their original state, complete with images, text, and formatting. Each spam email was randomly assigned an identification number to keep content coders blind to the identity of participants and their age and gender.

Coding manual and procedure

Our coding manual was developed from the literature. In particular, for the category of 'weapons of influence', we based our manual on Cialdini's six principles of influence (Cialdini 2006); a seventh weapon, perceptual contrast, was added from Hadnagy (Hadnagy 2010). For the category 'life domains', the coding manual leveraged work on six incentivizing (Hadnagy 2010) and ten contextual (Schindler et al. 2006) life domains. The coding manual was then further refined after coding of 100 sample emails from our set of collected spam emails. The final coding manual was composed of (i) comprehensive definitions and examples of each weapon of influence and each life domain (see "Related work" section for a summary) and (ii) a set of key words obtained during the sample coding process and based on the literature (Cialdini 2006; Hadnagy 2010; Schindler et al. 2006) to allow coders to determine the content of the emails. For example, our coding instructions for the weapon Authority were as follows:

1 Definition: The principle of Authority states that humans tend to comply with requests made by figures of authority or reputable entities.
2 Example: A municipal parking authority sending an email about a traffic violation and inviting the victim to refute the claim online.
3 Key words: IT and HR Department, Loan offices, IRS, a government body, a parking or municipal authority, a Professor, a medical doctor, violation, fee, etc.

We developed a Qualtrics application/interface for the coders to perform their qualitative content analysis. Coders were instructed to follow the process below.

1 Enter e-mail ID in the Qualtrics interface.
2 Enter coder-ID in the Qualtrics interface.
3 Read the entire email.
4 Identify key words or phrases which fall into one of the categories of weapon influence (or life domain) defined in coding manual.
5 Add a primary and a secondary weapon of influence (or life domain) using a scale from 0 to 10 (0 signifying no presence to 10 indicating a perfect example of the chosen category).
6 Add a justification for the categorization and ratings in the comment field specifying words or phrases present within the email.
7 Click submit button to enter the coded information.

The coding procedure allowed for selection of a primary and a secondary weapon of influence and life domain. The qualitative content coding applied in this paper was done manually via trained human coders and lasted from August 2015 to September 2016. This approach is very time and human resource intensive. Therefore, we limited the number of coded emails to a manageable number of messages. To assure comparable representation of emails from all of our participants for the manual content coding process, we randomly selected up to 20 emails ($M = 14.31$, $SD = 3.78$) from each participant for a total of 961 emails [514 from young (57% female) and 447 from older (48% female)] users. In other words, for each participant, a random number from (1, 20) was drawn, representing the number of emails we would consider from this participant's set of forwarded emails. That is, in spite of our large dataset of collected emails, we limited the amount of emails that were coded because of time constraints. Our selection process of up to 20 emails per participant assured randomization and representation of emails from all participants.

Training and calibration of the six coders took two months (prior to actual coding) and resulted in good inter-rater reliability (Cohen's Kappa > .80).

Coders assigned the specified categories (which weapons of influence and life domains) to each email and rated each category's salience on an 11-point scale (0 signifying no presence to 10 indicating a perfect example of the chosen category). Coders justified their categorization and ratings by reporting in a comment field specific words or phrases present within the email. Sixty-three (7%) emails for weapons, 43 (5%) emails for incentivizing life domains, and 49 (5%) emails for contextual life domains did not fall under any of the categories and were excluded from the analysis (e.g., email was blank or contained just an image with no text).

Consider the spam email illustrated in Fig. 1. The subject of this email reads *Prize Notification*, signaling that the recipient had unexpectedly won a prize. In the way the email is constructed, the recipient was supposed to feel indebted to *Mrs. Miriam Inaki* and her organization

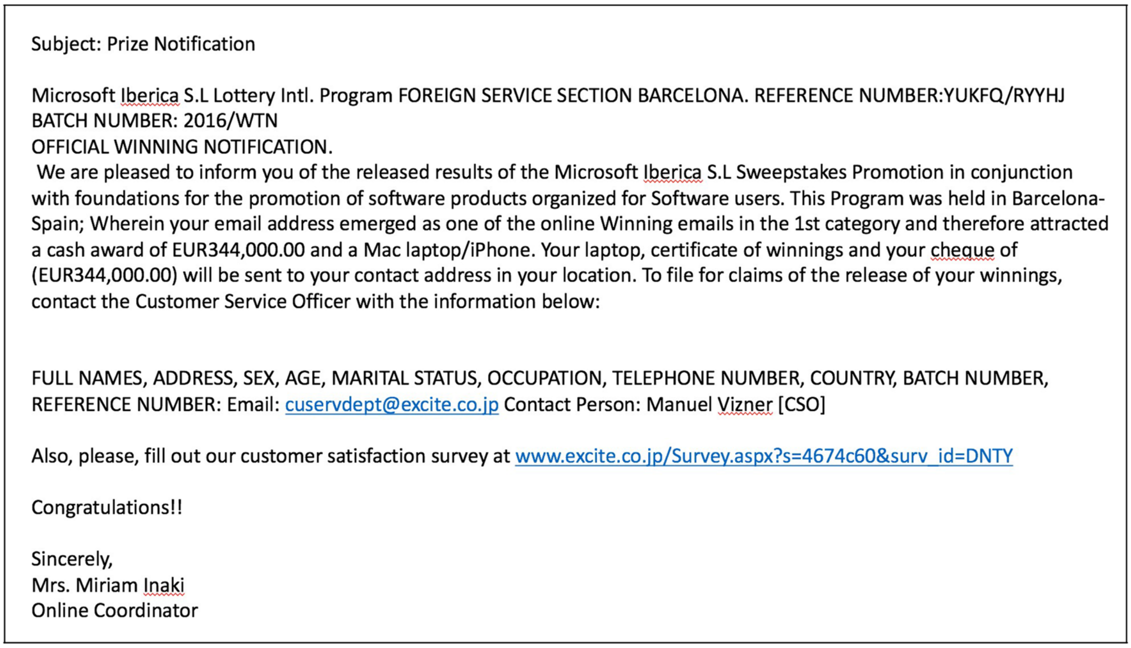

Subject: Prize Notification

Microsoft Iberica S.L Lottery Intl. Program FOREIGN SERVICE SECTION BARCELONA. REFERENCE NUMBER:YUKFQ/RYYHJ
BATCH NUMBER: 2016/WTN
OFFICIAL WINNING NOTIFICATION.

We are pleased to inform you of the released results of the Microsoft Iberica S.L Sweepstakes Promotion in conjunction with foundations for the promotion of software products organized for Software users. This Program was held in Barcelona-Spain; Wherein your email address emerged as one of the online Winning emails in the 1st category and therefore attracted a cash award of EUR344,000.00 and a Mac laptop/iPhone. Your laptop, certificate of winnings and your cheque of (EUR344,000.00) will be sent to your contact address in your location. To file for claims of the release of your winnings, contact the Customer Service Officer with the information below:

FULL NAMES, ADDRESS, SEX, AGE, MARITAL STATUS, OCCUPATION, TELEPHONE NUMBER, COUNTRY, BATCH NUMBER, REFERENCE NUMBER: Email: cuservdept@excite.co.jp Contact Person: Manuel Vizner [CSO]

Also, please, fill out our customer satisfaction survey at www.excite.co.jp/Survey.aspx?s=4674c60&surv_id=DNTY

Congratulations!!

Sincerely,
Mrs. Miriam Inaki
Online Coordinator

Fig. 1 Spam email example taken from the current study. This email applied reciprocation as weapon of influence (reflected in key words, such as prize, winning, cash award, or Please, fill out) and finances as life domain (reflected in EUR amount)

and fill out the survey in gratitude. In this case, the coders selected reciprocation as the weapon of influence being present. Coders justified their choice by inputting words signaling reciprocation, such as *prize*, *winning*, *cash award*, or *Please, fill out*.

Coders met once a week to resolve potential discrepancies with the goal to maintain a high intercoder agreement rate. To determine interrater reliability, 10% of the emails (randomly selected and assigned to coders) were independently coded by two coders, thereby following an often applied approach in qualitative content coding to not double code 100% of the content, but a small percentage (in our case 10%) (Saldana 2012).

Dichotomous variables were created for each category of weapons of influence, incentivizing life domain, and contextualizing life domains, respectively, based on the salience ratings given during the coding process. In particular, the value assigned was 1 if the corresponding salience rating was 5 or higher, indicating the presence of this category in the email; otherwise the value assigned was 0. These dichotomous variables were then used to determine the prevalence of each category in the email content (i.e., the frequency of use, expressed as a proportion, of a given category in an email relative to all emails collected from a given participant. For example, among 20 emails from a participant, seven emails had salience ratings on scarcity higher than 5. The prevalence of scarcity category of this participants was 35%. This prevalence of each category was used in the subsequent quantitative statistical analyses.

Based on the dichotomous variables created for each category of weapons, incentivizing life domains, and contextualizing life domains, a new set of categorical variables was created to indicate which type(s) in each category each email belonged to. For instance, an email belonged to a given category if the corresponding dichotomous variable was coded as 1. Thus, the Cohen's Kappa was calculated based on the categorical variables for all three dimensions respectively (weapons of influence = .78, incentivizing life domains = .90, and contextual life domains = .87), suggesting good to excellent interrater reliability. As the interrater reliability was calculated based on the category variable and there was one categorical variable for each rating dimension, there was only one Cohen's Kappa for each rating dimension.

Statistical analysis

This section presents the quantitative statistical analysis conducted on the content coded spam emails and results pertaining to our research questions.

(1) To what extent were weapons of influence and specific life domains represented in spam emails, and which weapons and life domains were most popular?

The respective prevalence of each weapon of influence (seven categories), incentivizing life domains (six categories), and contextual life domains (ten categories)

Table 2 Prevalence of each category of weapons of influence

	Mean	SD	Median	Min	Max	Skewness	Kurtosis
Reciprocation	0.75	0.18	0.78	0.00	1.00	− 1.35	3.34
Liking	0.17	0.16	0.14	0.00	1.00	2.39	9.98
Scarcity	0.15	0.14	0.13	0.00	0.46	0.64	− 0.62
Social proof	0.04	0.07	0.00	0.00	0.29	2.03	3.82
Authority	0.18	0.14	0.14	0.00	0.91	2.44	10.58
Commitment	0.08	0.10	0.06	0.00	0.33	1.09	0.35
Perceptual contrast	0.05	0.07	0.00	0.00	0.25	1.26	0.61

Table 3 Prevalence of each incentivizing life domain

	Mean	SD	Median	Min	Max	Skewness	Kurtosis
Financial	0.69	0.21	0.75	0.00	1.00	− 1.32	2.04
Health	0.12	0.15	0.09	0.00	1.00	3.54	19.46
Ideological	0.09	0.11	0.07	0.00	0.56	1.81	3.97
Social	0.17	0.18	0.13	0.00	0.91	1.60	3.88
Legal	0.02	0.06	0.00	0.00	0.33	3.37	13.59
Security	0.04	0.07	0.00	0.00	0.40	2.47	8.26

Table 4 Prevalence of each category of contextual life domain

	Mean	SD	Median	Min	Max	Skewness	Kurtosis
Health	0.14	0.14	0.12	0.00	1.00	3.60	20.21
Cognitive	0.02	0.04	0.00	0.00	0.15	2.36	4.62
Independence	0.27	0.20	0.22	0.00	0.86	0.69	0.08
Family	0.02	0.04	0.00	0.00	0.18	1.95	3.75
Friends	0.03	0.07	0.00	0.00	0.50	4.77	28.80
Life reflection	0.11	0.11	0.08	0.00	0.44	0.90	0.30
Death	0.01	0.02	0.00	0.00	0.10	3.63	12.48
Occupation	0.09	0.10	0.07	0.00	0.33	0.89	− 0.20
Leisure	0.41	0.30	0.35	0.00	1.00	0.41	− 0.98
Sexuality	0.08	0.12	0.00	0.00	0.64	2.59	8.07

were calculated. Descriptive statistics are presented in Tables 2, 3, and 4.

To test significant differences in the prevalence of specific weapons of influence and life domains in the content-coded spam emails, given the non-normal distribution of the data and the nested data structure (i.e., weapons of influence and life domains were nested in each participant; repeated measures), we conducted three separate Friedman's analyses of variance (ANOVAs) (Gravetter and Wallnau 2009). Category of a given dimension (weapons of influence and life domains, respectively) constituted the within-subject variable. For significant dimensions in the Friedman's ANOVA we followed up with simple effect analysis, which consisted of pairwise comparisons between this dimension and all other dimensions (e.g., reciprocation vs. liking). We used Wilcoxon signed ranks test for these follow up analyses. For these pairwise comparisons, Bonferroni correction was applied for determination of the statistical threshold (*p*-value), thus accounting for the type-I error inflation rate due to multiple comparisons. Bonferroni correction was based on the number of categories within each dimension. In particular, for weapons of influence the corrected *p*-value was 0.001, for incentivizing life domains it was 0.003, and for contextual life domains it was 0.002.

Table 5 Pairwise comparisons among weapons of influence

	Liking	Scarcity	Social proof	Authority	Commitment	Perceptual contrast
Reciprocation	*– 6.95*[a]	*– 7.06*[a]	*– 7.10*[a]	*– 6.73*[a]	*– 7.01*[a]	*– 7.06*[a]
Liking		– 0.30	*– 5.62*[a]	– 0.78	*– 3.60*[a]	*– 5.28*[a]
Scarcity			*– 5.07*[a]	– 1.05	*– 3.03*[a]	*– 4.36*[a]
Social proof				*– 6.10*[b]	*– 3.22*[b]	– 1.32
Authority					*– 4.75*[a]	*– 5.71*[a]
Commitment						– 2.29

[a] Indicates that the prevalence of weapon of influence shown in the row was higher than the prevalence of weapon of influence shown in the corresponding column.[b] Indicates that the prevalence of weapon of influence shown in the row was lower than the prevalence of weapon of influence shown in the corresponding column. Italic indicates significant difference between prevalence of weapons of influence at the Bonferroni adjusted level of $p = .001$

There was a significant difference in the prevalence of specific weapons of influence ($\tilde{\chi}^2(6) = 225.48$, $p < .001$). Reciprocation was more prevalent than any other weapon of influence (Table 5). Liking, scarcity, and authority were more prevalent than social proof, commitment, and perceptual contrast. In addition, commitment was more prevalent than social proof. There were no differences between social proof and perceptual contrast or between commitment and perceptual contrast.

There was also a significant difference in the prevalence of specific incentivizing life domains ($\tilde{\chi}^2(5) = 181.00$, $p < .001$). Financial emails were more prevalent than all other incentivizing life domains (Table 6). While social and health emails were as prevalent as ideological emails, they were more prevalent than legal and security emails. Ideological emails were equally prevalent as security emails, but they were more prevalent than legal emails. Legal and security emails did not differ in prevalence.

There was a significant difference in the prevalence of contextual life domains ($\tilde{\chi}^2(9) = 296.15$, $p < .001$). Independence and leisure emails were equally prevalent and were more prevalent than all other contextual life domains (Table 7). Emails related to health, life reflections, and occupations were comparable in prevalence, but were more prevalent than cognitive, family, friends, and death related emails. Emails pertaining to health and life reflections showed higher prevalence than emails pertaining to sexuality. However, prevalence of occupation and sexuality emails were not different from each other. Finally, emails related to cognitive, family, friends, and death did not differ in prevalence.

(2) Did the use of weapons of influence and life domains vary by user age group?

To test for age differences in the content of the spam emails, accommodating for the non-normal distribution of our data, separate Mann-Whitney U tests were conducted on the prevalence of each weapon of influence, incentivizing life domain, and contextual life domain, respectively. For weapons of influence, there were no significant age differences in the prevalence of any of the categories.

For incentivizing life domains (Fig. 2), health showed a significant age difference ($U = 347.50$, $p = .007$, $r = .33$). In particular, older users received health-related spam emails more frequently than young users ($Mdn_{young} = .05$, $Mdn_{older} = .13$). There were no age differences for financial, ideological, social, legal, and security.

For contextual life domains (Fig. 3), there were significant age differences for health ($U = 343.50$, $p = .006$, $r = .33$), independence ($U = 336.50$, $p = .005$, $r = .34$), occupation ($U = 384.50$, $p = .021$, $r = .28$), and leisure ($U = 361.50$, $p = .013$, $r = .30$). In particular, older users were more likely to receive spam emails relevant to health ($Mdn_{young} = .10$, $Mdn_{older} = .14$) and independence ($Mdn_{young} = .18$, $Mdn_{older} = .36$), whereas young users were more likely to receive spam emails relevant to occupation ($Mdn_{young} = .11$, $Mdn_{older} = 0$) and leisure ($Mdn_{young} = .58$, $Mdn_{older} = .23$). There were no significant age differences for cognitive, family, life reflections, and sexuality.

Table 6 Pairwise comparisons among incentivizing life domains

	Health	Ideological	Social	Legal	Security
Financial	*– 6.59*[a]	*– 6.96*[a]	*– 6.50*[a]	*– 6.96*[a]	*– 7.01*[a]
Health		– 0.86	– 2.13	*– 4.66*[a]	*– 3.99*[a]
Ideological			– 2.83	*– 4.04*[a]	*– 2.46*
Social				*– 5.29*[a]	*– 4.72*[a]
Legal					– 2.37

[a] Indicates that the prevalence of incentivizing life domain shown in the row was higher than the prevalence of incentivizing life domain shown in the corresponding column. Italic indicates significant difference between prevalence of incentivizing life do mains at the Bonferroni adjusted level of $p = .003$

Table 7 Pairwise comparisons among contextual life domains

	Cognitive	Independence	Family	Friends	Life reflections	Death	Occupation	Leisure	Sexuality
Health	*− 6.01*[a]	*− 4.75*[b]	*− 6.21*[a]	*− 5.32*[a]	− 0.49	*− 6.40*[a]	− 2.42	*− 5.28*[b]	− 2.83
Cognitive		*− 6.59*[b]	− 0.46	− 1.07	*− 5.28*[b]	− 1.93	*− 4.55*[b]	*− 6.79*[b]	*− 3.75*[b]
Independence			*− 6.58*[a]	*− 6.31*[a]	*− 4.87*[a]	*− 6.79*[a]	*− 5.73*[a]	− 2.24	*− 5.59*[a]
Family				− 0.68	*− 5.28*[b]	− 3.08	*− 4.26*[b]	*− 6.80*[b]	*− 3.61*[b]
Friends					*− 4.58*[b]	− 2.94	*− 3.87*[b]	*− 6.56*[b]	− 3.07
Life reflections						*− 5.78*[a]	− 1.54	*− 5.52*[b]	− 2.09
Death							*− 5.07*[b]	*− 6.85*[b]	*− 4.50*[b]
Occupation								*− 5.95*[b]	− 0.80
Leisure									*− 5.74*[a]

[a] Indicates that the prevalence of contextual life domain shown in the row was higher than the prevalence of contextual life domain shown in the corresponding column.[b] Indicates that the prevalence of contextual life domain shown in the row was lower than the prevalence of contextual life domain shown in the corresponding column. Italic indicates significant difference between prevalence of contextual life domains at the Bonferroni adjusted level of $p = .002$

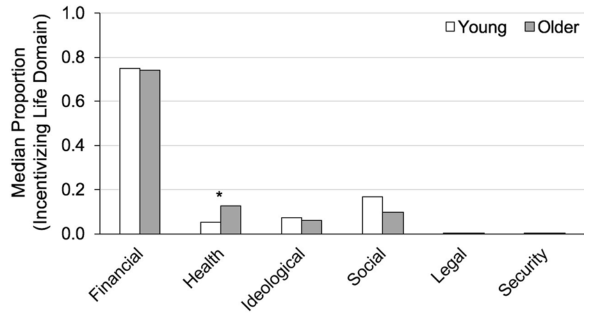

Fig. 2 Median prevalence of each incentivizing life domain in spam emails received by young and older users. The *prevalence* of each category of incentivizing life domains was represented by the proportion of emails in a given category relative to all emails collected from given participant. *Indicates significant age difference at $p < .05$

Discussion Our analysis showed that all weapons of influence and life domains were represented in young vs older user's spam emails, with the weapon of reciprocation, the incentivizing domain of financial, and the contextual domains of leisure and independence particularly prevalent. While all categories were represented, the prevalence of some (i.e., Death) was quite low. This finding overall supports our conceptualization of weapons of influence and life domains as techniques to lure Internet users into fall for (potentially malicious) spam.

Reciprocation, followed by authority, liking, and scarcity were the most frequently utilized weapons in today's spam emails. This finding is relevant in light of recent research on phishing susceptibility (Oliveira et al. 2017), which suggests that specific weapons of influence are particularly effective in luring users into clicking on links in emails. The present analysis suggests that spammers are employing effective weapons and relate email content to relevant life domains in their campaigns.

Our data also showed that emails pertaining to the financial incentivizing life domain were the most prevalent ones, followed by social and health emails. This finding is interesting in the context of recent research (Oliveira et al. 2017) showing that users do not have a particular susceptibility to financial phishing emails, but rather are highly susceptible to legal phishing emails.

An email selection from a larger and more representative group of Internet users would have increased generalizability of our findings and would have allowed additional subgroup analyses (e.g., comparing young and older women and men). Larger longitudinal studies are warranted to confirm the observed age effects and for a comprehensive developmental analysis of the content and dynamics of spam campaigns against individuals of

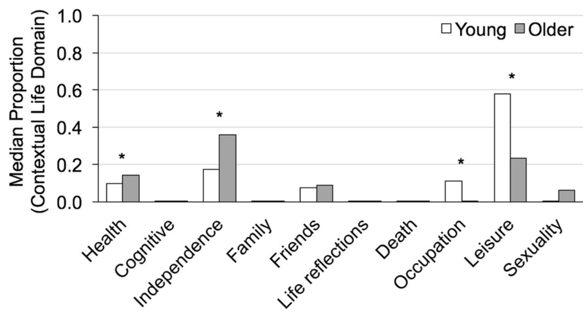

Fig. 3 Median prevalence of each contextual life domain in spam emails received by young and older users. The *prevalence* of each category of contextual life domains was represented by the proportion of emails in a given category relative to all emails collected from a given participant. *Indicates significant age difference at $p < .05$

different ages. Future studies should also include additional demographics in their analysis, including race/ethnicity, level of education, and socioeconomic background, as factors that need consideration in tailored implementation of future defense solutions.

Our findings were based on non-parametric analyses, given the non-normal distribution of our data and the relatively small sample size. While non-parametric analysis does not make stringent assumptions about the distribution of the data, it possesses reduced analytical power compared to parametric testing.

Our coding process was limited to emails that were already caught by spam filters. In the future, it would be interesting to study verifiably malicious emails that successfully passed through these filters, towards improvement of security measures.

In spite of these limitations, our study provides intriguing first evidence suggesting that current spam emails target age groups differently regarding life domains. While older users were more likely to receive spam emails relevant to health and independence, young users were more likely to receive spam emails relevant to leisure and occupation. In contrast, we found no age-differential targeting regarding use of different weapons. Recent research on phishing susceptibility (Oliveira et al. 2017) showed that young and older users significantly differ in their susceptibility to different weapons.

Taken together, our findings suggest that spam could potentially be used in a more effective way in the future, if spammers targeted users according to particularly effective weapons and life domains and in line with age-specific vulnerabilities, rendering development of the new-generation of effective detection and warning solutions even more relevant. In fact, Hadnagy (2010) discusses that prior intelligence-gathering about targets in social engineering attacks (including innocuous advertisements) is the first step conducted by professional social engineers, scammers, and even advertisers. Hadnagy also discusses how data originating from data breaches and available in black markets can streamline Internet user's targeting. The labeled dataset on influence and life domains in spam that we have created in this study and which we plan to make available to the research community, can be leveraged for the development of machine learning models for the detection of the use of influence in email text. The identification of influence in text can be a game-changer for the next generation of tools to detect spam and phishing by warning users about potentially deception cues in text.

Conclusions

This paper presented an analysis of modern spam from an age-comparative user perspective, integrating manual qualitative content coding and quantitative statistics. We aimed to clarify (i) the extent to which weapons of influence and life domains were represented in young vs older users' spam emails and (ii) variations of the prevalence of weapons of influence and life domains by age demographic. Our study demonstrated the presence of some level of age-specific targeting in current spam campaigns. This knowledge is crucial in its potential for integration in the development of future spam mitigation solutions, capable of detecting influence in emails and warning users in an demographic-targeted fashion such as by considering age-specific vulnerabilities. Moving forward, we plan to leverage this manually labeled dataset of emails to develop machine learning classifiers that can detect influence in text.

Abbreviations

APT: advanced persistent threat; US: United States; IP: Internet Protocol; URL: Uniform Resource Locator; ISP: Internet Service Provider; HTML: Hypertext Markup Language; IT: Information Technology; HR: Human Resources; IRS: Internal Revenue Service; ID: identification number; ANOVA: analysis of variance.

Authors' contributions

DSO and NCE: supervision and drafting of manuscript. TL: data analysis. HR, SM: coding and coding supervision. SD, HY: data collection infrastructure. DW: recruitment. All authors read and approved the final manuscript.

Author details

[1] Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL, USA. [2] Department of Psychology, University of Florida, Gainesville, FL, USA. [3] Department of Computer & Information Science & Engineering, Gainesville, FL, USA.

Acknowledgements

We thank the anonymous reviewers for their insightful comments and guidance. We would like to thank our undergraduate research assistants and coders: Cheyenne Reynolds, Nicole Phillips, Hannah Burrichter, Nicole Phillips, Paul Talty, and Melis Muradoglu.

Competing interests

The authors declare that they have no competing interests.

Availability of data and materials

Data code used in this paper can be retrieved from https://osf.io/8npjg.

Funding

This research was supported by NSF Grant SBE-1450624 and NIH/NIA Grant R01AG057764.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 1 June 2018 Accepted: 15 March 2019
Published online: 01 April 2019

References

Akbar, N. (2014). Analysing Persuasion principles in phishing emails. Master's thesis, University of Twente.

Benenson, Z., Gassmann, F., & Landwirth, R. (2017). Unpacking spear phishing susceptibility. Financial Cryptography and Data Security Workshops, 1–17.

Caputo, D. D., Pfleeger, S. L., Freeman, J. D., & Johnson, M. E. (2014). Going spear phishing: Exploring embedded training and awareness. *IEEE Security & Privacy*, *12*(1), 28–38.

Cialdini, R. B. (2006). *Influence—The psychology of Persuasion* (1st ed.). New York: Harper Business.

Downs, J.S., Holbrook, M.B., & Cranor, L.F. (2006). Decision strategies and susceptibility to phishing. Symposium on Usable Privacy and Security (SOUPS).

Drucker, H., Wu, D., & Vapnik, V. N. (1999). Support vector machines for spam categorization. *IEEE Transactions on Neural Networks*, *10*(5), 1048–1054. https://doi.org/10.1109/72.788645.

Edwards, B., Hofmeyr, S., Forrest, S., & van Eeten, M. (2015). Analyzing and modeling longitudinal security data: Promise and pitfalls. In: *Proceedings of the 31st annual computer security applications conference. ACSAC 2015* (pp. 391–400). New York, NY, USA: ACM. https://doi.org/10.1145/2818000.2818010.

FBI. Fraud against Seniors. https://www.fbi.gov/scams-and-safety/common-fraud-schemes/seniors.

Ferreira, A., & Lenzini, G. (2015). An analysis of social engineering principles in effective phishing. Workshop on Socio-Technical Aspects in Security and Trust (STAST).

Fette, I., Sadeh, N., & Tomasic, A. (2007). Learning to detect phishing emails. In: *Proceedings of the 16th international conference on World Wide Web* (pp. 649–656). ACM.

Gragg, D. (2003). A multi-level defense against social engineering. *SANS Reading Room*, 13.

Gravetter, F., & Wallnau, L. (2009). *Statistics for the behavioral sciences* (8th ed.). Independence: Cengage Learning.

Hadnagy, C. (2010). *Social engineering: The art of human hacking*. Hoboken: Wiley.

Hao, S., Syed, N. A., Feamster, N., Gray, A. G., & Krasser, S. (2009). Detecting spammers with snare: Spatio-temporal network-level automatic reputation engine. In: *Proceedings of the 18th conference on USENIX security symposium. SSYM'09* (pp. 101–118). Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=1855768.1855775.

Johnson, M. (1990). Age differences in decision making: A process methodology for examining strategic information processing. *Journal of Gerontology: Psychological Sciences*, *45*(2), 75–78.

Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G.M., Paxson, V., & Savage, S. (2008). Spamalytics: An empirical analysis of spam marketing conversion. In: *Proceedings of the 15th ACM conference on computer and communications security. CCS '08* (pp. 3–14). New York, NY, USA: ACM. https://doi.org/10.1145/1455770.1455774.

Kanich, C., Weavery, N., McCoy, D., Halvorson, T., Kreibichy, C., Levchenko, K., Paxson, V., Voelker, G.M., & Savage, S. (2011). Show me the money: Characterizing spam-advertised revenue. In: *Proceedings of the 20th USENIX conference on security. SEC'11* (pp. 15–15). Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=2028067.2028082.

Kumaraguru, P. (2009). Phishguru: A system for educating users about semantic attacks (PhD Thesis), Carnegie Mellon University.

Kumaraguru, P., Rhee, Y., Acquisti, A., Cranor, L.F., Hong, J., & Nunge, E. (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* (pp. 905–914). ACM.

Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L. F., & Hong, J. (2010). Teaching johnny not to fall for phish. *ACM Transactions on Internet Technology (TOIT)*, *10*(2), 7.

Lauricella, T.: 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.

Mata, R., Josef, A. K., Samanez-Larkin, G. R., & Hertwig, R. (2011a). Age differences in risky choice: A meta-analysis. *Annals of the New York Academy of Sciences*, *1235*, 18–29.

Mata, R., Josef, A., Samanez-Larkin, G., & Hertwig, R. (2011b). *Age differences in risky choice: A meta-analysis*. New York: New York Academy of Sciences.

Mather, M. (2006). *When I'm 64—A review of decision-making processes: Weighing the risks and benefits of aging*. Washington, DC: The National Academies Press.

Meyer, T. A., & Whateley, B. (2005). Spambayes: Effective open-source, bayesian based, email classification system. In: *Proceedings of the first conference on email and anti-spam. CEAS*.

Mitnick, K., Simonand, W. L., & Wozniak, S. (2002). *The art of deception: Controlling the human element of security*. Hoboken: Wiley.

Oliveira, D., Rocha, H., Yang, H., Ellis, D., Dommaraju, S., Muradoglu, M., Weir, D., Soliman, A., Lin, T., & Ebner, N. (2017). Dissecting spear phishing emails for older vs young adults: On the interplay of weapons of influence and life domains in predicting susceptibility to phishing. In: *Proceedings of the 2017 CHI conference on human factors in computing systems. CHI '17* (pp. 6412–6424). New York, NY, USA: ACM. https://doi.org/10.1145/3025453.3025831.

Ortiz, P. (2010). Machine learning techniques for persuasion detection in conversation. Master's thesis, Naval Post Graduate School.

Pitsillidis, A., Levchenko, K., Kreibich, C., Kanich, C., Voelker, G.M., Paxson, V., Weaver, N., & Savage, S. (2010). Botnet judo: Fighting spam with itself. In: *Proceedings of the network and distributed system security symposium, NDSS 2010*. San Diego, California, USA, 28th February–3rd March 2010.

Ramachandran, A., Feamster, N., & Vempala, S. (2007). Filtering spam with behavioral blacklisting. In: *Proceedings of the 14th ACM conference on computer and communications security. CCS '07* (pp. 342–351). New York, NY, USA: ACM. https://doi.org/10.1145/1315245.1315288.

Redmiles, E.M., Chachra, N., & Waismeyer, B. (2018). Examining the demand for spam: Who clicks? In: *Proceedings of the 2018 CHI conference on human factors in computing systems. CHI'18*.

Reed, A. E., Chan, L., & Mikels, J. A. (2014). Meta-analysis of the age-related positivity effect: Age differences in preferences for positive over negative information. *Psychology and Aging*, *1*, 1–15.

Saldana, J. (2012). *The coding manual for qualitative researchers*. Thousand Oaks: SAGE Publications.

Schindler, I., Staudinger, U. M., & Nesselroade, J. R. (2006). Development and structural dynamics of personal life investment in old age. *Psychology and Aging*, *21*, 737–753.

Schwartz, A. (2004). *SpamAssassin*. Sebastopol: O'Reilly Media, Inc.

Sculley, D., & Wachman, G. M. (2007). Relaxed online svms for spam filtering. In: *Proceedings of the 30th annual international ACM SIGIR conference on research and development in information retrieval. SIGIR '07* (pp. 415–422). New York, NY, USA: ACM. https://doi.org/10.1145/1277741.1277813.

Sheng, S., Wardman, B., Warner, G., Cranor, L.F., Hong, J., & Zhang, C. (2009). An empirical analysis of phishing blacklists. CEAS—Sixth conference on email and anti-Spam.

Sheng, S., Magnien, B., Kumaraguru, P., Acquisti, A., Cranor, L.F., Hong, J., & 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* (pp. 88–99). ACM.

Stajano, F., & Wilson, P. (2011). Understanding scam victims: Seven principles for systems security. *Communications of ACM*, *54*(3), 70–75.

Stone-Gross, B., Holz, T., Stringhini, G., & Vigna, G. (2011). The underground economy of spam: A botmaster's perspective of coordinating large-scale spam campaigns. In: *Proceedings of the 4th USENIX conference on large-scale exploits and emergent threats. LEET'11* (pp. 4–4). Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=1972441.1972447.

Stringhini, G., Holz, T., Stone-Gross, B., Kruegel, C., & Vigna, G. (2011). Botmagnifier: Locating spambots on the internet. In: *Proceedings of the 20th USENIX conference on security. SEC'11* (pp. 28–28). Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=2028067.2028095.

Stringhini, G., Egele, M., Zarras, A., Holz, T., Kruegel, C., & Vigna, G. (2012). B@bel: Leveraging email delivery for spam mitigation. In: *Presented as Part of the 21st USENIX Security Symposium (USENIX Security 12)* (pp. 16–32). Bellevue, WA: USENIX. https://www.usenix.org/conference/usenixsecurity12/technical-sessions/presentation/stringhini.

Stringhini, G., Hohlfeld, O., Kruegel, C., & Vigna, G. (2014). The harvester, the botmaster, and the spammer: On the relations between the different actors in the spam landscape. In: *Proceedings of the 9th ACM symposium on information, computer and communications security. ASIA CCS '14* (pp. 353–364). New York, NY, USA: ACM. https://doi.org/10.1145/2590296.2590302.

Symantec (2017). Symantec security report 2017. Retrieved June, 05, 2017, from https://www.symantec.com/security-center/threat-report.

Taylor, B. (2006). Sender reputation in a large webmail service. In: *Conference on email and anti-SPAM. CEAS*.

Tentoria, K., Oshersonb, D., Hasherc, L., & May, C. (2001). *Wisdom and aging: Irrational preferences in college students but not older adults*. Amsterdam: Elsevier Science.

Toolbar, N.: Netcraft Ltd. https://toolbar.netcraft.com/.

Uebelacker, S., & Quiel, S. (2014). The social engineering personality framework. Workshop on Socio-Technical Aspects in Security and Trust (STAST).

Verhaeghen, P. (2003). Aging and vocabulary score: A meta-analysis. *Psychology and Aging*, *18*(2), 332–339.

Verhaeghen, P., & Salthouse, T. A. (1997). Meta-analyses of age-cognition relations in adulthood: Estimates of linear and nonlinear age effects and structural models. *Psychological Bulletin*, *122*(3), 231–249.

Wong, J. C., & Solon, O. (2017). Massive ransomware cyber-attack hits nearly 100 countries around the world. Retrieved June, 05, 2017, from https://www.theguardian.com/technology/2017/may/12/global-cyber-attack-ransomware-nsa-uk-nhs.

Workman, M. (2007). Gaining access with social engineering: An empirical study of the threat. *Information Systems Security*, *16*(6), 315–331.

Xie, Y., Yu, F., Achan, K., Panigrahy, R., Hulten, G., & Osipkov, I. (2008). Spamming botnets: Signatures and characteristics. *ACM SIGCOMM Computer Communication Review*, *38*(4), 171–182. https://doi.org/10.1145/1402946.1402979.

Zhang, Y., Egelman, S., Cranor, L., & Hong, J. (2006). *Phinding phish: Evaluating anti-phishing tools*. Reston: ISOC.