

2-2019

Phishing and Cybercrime Risks in a University Student Community

cybercrime, risks, university, phishing, social engineering

Follow this and additional works at: <https://vc.bridgew.edu/ijcic>



Part of the [Criminology Commons](#), [Criminology and Criminal Justice Commons](#), and the [Information Security Commons](#)

Recommended Citation

Broadhurst, Roderic; Skinner, Katie; Sifniotis, Nicholas; Matamoros-Macias, Bryan; and Ipsen, Yuguang (2019) Phishing and Cybercrime Risks in a University Student Community, *International Journal of Cybersecurity Intelligence & Cybercrime*: 2(1), 4-23. <https://www.doi.org/10.52306/02010219RZEX445>

This item is available as part of Virtual Commons, the open-access institutional repository of Bridgewater State University, Bridgewater, Massachusetts.

Copyright © 2-2019 Roderic Broadhurst, Katie Skinner, Nicholas Sifniotis, Bryan Matamoros-Macias, and Yuguang Ipsen

Broadhurst, R., Skinner, K., Sifniotis, N., Matamoros-Macias, B., & Ipsen, Y. (2019). *International Journal of Cybersecurity Intelligence and Cybercrime*, 2 (1), 4-23.

Phishing and Cybercrime Risks in a University Student Community

Roderic Broadhurst*, Cybercrime Observatory, Australian National University

Katie Skinner, Cybercrime Observatory, Australian National University

Nicholas Sifniotis, Cybercrime Observatory, Australian National University

Bryan Matamoros-Macias, Cybercrime Observatory, Australian National University

Yuguang Ipsen, Research School of Finance, Actuarial Studies and Statistics, ANU

Key Words; cybercrime, phishing, social engineering

Abstract:

In an exploratory quasi-experimental observational study, 138 participants recruited during a university orientation week were exposed to social engineering directives in the form of fake email or phishing attacks over several months in 2017. These email attacks attempted to elicit personal information from participants or entice them into clicking links which may have been compromised in a real-world setting. The study aimed to determine the risks of cybercrime for students by observing their responses to social engineering and exploring attitudes to cybercrime risks before and after the phishing phase. Three types of scam emails were distributed that varied in the degree of individualization: generic, tailored, and targeted or 'spear.' To differentiate participants on the basis of cybercrime awareness, participants in a 'Hunter' condition were primed throughout the study to remain vigilant to all scams, while participants in a 'Passive' condition received no such instruction. The study explored the influence of scam type, cybercrime awareness, gender, IT competence, and perceived Internet safety on susceptibility to email scams. Contrary to the hypotheses, none of these factors were associated with scam susceptibility. Although, tailored and individually crafted email scams were more likely to induce engagement than generic scams. Analysis of all the variables showed that international students and first year students were deceived by significantly more scams than domestic students and later year students. A Generalized Linear Model (GLM) analysis was undertaken to further explore the role of all the variables of interest and the results were consistent with the descriptive findings showing that student status (domestic compared to international) and year of study (first year student compared to students in second, third and later years of study) had a higher association to the risk of scam deception. Implications and future research directions are discussed.

*Corresponding author

Roderic Broadhurst, Professor of Criminology, School of Regulation and Global Governance, Fellow Research School of Asia and the Pacific, Australian National University (ANU)

Email: roderic.broadhurst@anu.edu.au

Reproduction, posting, transmission or other distribution or use of the article or any material therein, in any medium as permitted by written agreement of the International Journal of Cybersecurity Intelligence and Cybercrime, requires credit to the Journal as follows: "This Article originally appeared in International Journal of Cybersecurity Intelligence and Cybercrime (IJCIC), 2019 Vol. 2, Iss. 1, pp. 4-23" and notify the Journal of such publication.

© 2019 IJCIC 2578-3289/2019/02

Introduction

As individuals become increasingly connected to the virtual world, the avenues through which cybercriminals may exploit them likewise increase. Although developments in technology have attempted to mitigate these risks, human error continues to be the ‘weakest link’ in cyber security (Mayhorn, Welk, Zielinska, & Murphy-Hill, 2015). When cybercriminals employ ‘spam,’ ‘phishing,’ or ‘spear phishing’ methods in their attempts to hack, distribute malware, or steal personal information, they target their victim’s judgment rather than their virtual security measures (Gratian, Bandi, Cukier, Dykstra, & Ginther, 2018; Alazab & Broadhurst, 2016).

A common vector for distributing malware is spam email. Spam can involve harmless advertising through unsolicited emails, SMS texts, or social network messages, but spam may also contain viruses or malware designed to exploit personal or sensitive information from its recipients. Though spam may seem insignificant at the individual level, estimates indicate that the world-wide average daily volume of spam was approximately 422 billion in January 2018, constituting about 85 percent of all daily global email traffic (Talos, 2018).

The widespread and ubiquitous threat of malware-borne spam poses significant economic and social consequences, however, experiences of victimization and susceptibility are not universal. Previous studies have suggested that gender (Sun, Yu, Lin, & Tseng, 2016), age (Gavett et al., 2017), and technical experience (Pattinson, Jerram, Parsons, McCormac, & Butavicius, 2012) influence an individual’s susceptibility to spam and phishing attempts. This study was designed to further explore how these factors and others – namely the type of scam, cybercrime awareness, gender, and IT competence – influenced cybercrime susceptibility amongst a sample of university students. To accomplish this, research participants were exposed to various fake emails sent from the research team’s web server, and their interactions with these scams were observed. Engaging university students offered the advantage of using a single institutional Internet service, which could also monitor real spam events, and reduced ethical concerns associated with an open or public sample.

Literature: Factors Influencing Susceptibility

Broadly speaking, ‘spam’ encompasses all unsolicited electronic messages that are usually, but not always, sent in bulk transmission. Composers of scam messages combine technology with social engineering techniques in order to lure and deceive their victims into giving up sensitive information. In short, these offenders engage in a ‘phishing’ deception by enticing a response through email. While the purposes of phishing vary, it is often used to deliver malware, or ransomware, or to obtain personal information from the recipient for the purpose of identity theft.

Chaudhry, Chaudhry, and Rittenhouse (2016) suggest a typical phishing attack is comprised of three elements: a lure, a hook, and catch. The lure often involves an email message appearing to be from a legitimate person or organization, the reliability of which is strengthened through the exploitation of:

- Curiosity: such as emails containing compromised links which appear to lead to videos of recent news or events;
- Fear: such as emails from the ‘bank’ urging users to validate their information due to account breaches;
- Empathy: such as emails impersonating a friend or relative who is in need of financial assistance or personal information.

This list is not exhaustive and can be augmented by appeals to other emotions such as greed (e.g., a winning lottery ticket), lust or vanity (e.g. an adoring admirer, a prestigious job opportunity). De Kimpe, Walrave, Hardyns, Pauwels, and Ponnet (2018) list characteristics that can either facilitate or hinder the success of phishing e-mails (e.g., the presence of spelling and design/format errors, monetary prize offers). Once receivers are convinced the mail is authentic, the next stage is to convince the recipients to divulge sensitive information. Various social manipulators such as liking or trusting the email source; implicating reciprocity (e.g., returning favors) or ‘social proof’ (i.e., others are participating); creating a sense of scarcity or evoking an authoritative source will help the deception to succeed.

When phishing emails make use of personalized data in their lures, they become examples of ‘spear phishing.’ Spear phishing is contextual, with emails often containing specific information that would be familiar or important to specific recipients (De Kempe, Walrave, Hardyns, Pauwels & Ponnet, 2018). In order to obtain such information, attackers spend time obtaining private information relevant to particular users, and then use this information to craft fake emails (Caputo Pfleeger, Freeman, & Johnson, 2014). These emails tend to impersonate well-known companies, trusted relationships or contexts that have personal relevance to the individual (De Kempe et al., 2018).

The success of any phishing or spear phishing email is linked to how well it is able to deceive its recipient. While the research literature has focused on phishing email structure (e.g., use of visual cues, presence of misspelling or attachments; Parsons, McCormac, Pattinson, Butavicius, & Jerram, 2015), this study explores email contextualization and personalization. Phishing emails containing personalized information relevant to its recipients have been shown to be effective in deceiving their targets (Benenson, Gassman & Landwirth, 2016).

Butavicius, Parsons, Pattinson, and McCormac (2015) tested the effect of different social engineering strategies by sending a series of genuine, phishing, or spear phishing emails to a group of 117 university students. Overall, the results indicated that students tended to classify emails as genuine rather than fraudulent and were worse at detecting spear phishing attempts over generic phishing attempts. It was also found that, where spear phishing emails utilized an authority-style social engineering strategy (i.e., the apparent sender of the email held authority over the reader), students were less able detect spear-phishing.

Variables Associated with Phishing Risk

Individuals, once aware of their own potential victimization, are thought to become more cautious or defensive as they navigate risky environments. A national representative survey undertaken in 2017 by the Australian Institute of Criminology (AIC) offers some general insights into how individuals respond to identify theft, a key predicate offence. The survey found that identity theft was on the rise with 13.1% of the 9, 947 respondents reporting misuse of personal information in the past 12 months compared to 8.5% in 2016. The study suggested this was “. . . due to an increase in phishing attacks and information being obtained by telephone and in face-to-face meetings” (Goldsmid, Gannoni, & Smith, 2018, p. 59). About two-thirds of the sample perceived that the risk of victimization was likely to increase, with recent victims more likely to perceive an increased risk and more likely to adopt security measures such as password changes, signatures, and voice recognition (Goldsmid et al., 2018, p. 52; 50ff). Understanding the impact of participant ‘priming’ or awareness of potential risk can inform the development of programs aimed at preventing online victimization.

Participants who are informed that they are being tested on their ability to detect phishing emails fare better than those who are not informed (Pattinson et al., 2012), although the extent to which

priming or training assists in preventing victimization has been challenged (Alsharnouby, Alaca, & Chiasson, 2015; Caputo et al., 2014). In the present study, our subjects were all simply primed as ethical approval required off-line formal consent and agreement for attempts to deceive them with a scam email as required by the relevant Australian National University (ANU) Human Research Ethics Committee (per protocol HREC #2015/038). However, a subsample of our subjects was designated as 'Hunters,' and further primed to be more alert than other subjects.

Some studies have found females to be more susceptible to online scams (Iuga, Nurse, & Erola, 2016), and others have found no such connection (Butavicius et al., 2017; Oliviera et al., 2017). The AIC survey, for example found males and those aged 25-34 years (93.5% of our sample were 25 years of age or under) were significantly more likely to fall victim to identify theft than females or other age groups (Goldsmid et al., 2018). Despite these contradictory findings, recent studies have sought to re-frame the relationship between gender and phishing susceptibility. In Goel, Williams, and Dincelli's (2017) study, the act of falling for a phishing scam comprised two steps: first, the opening of a phishing email, and second, the clicking of the malicious link within. They found that whilst women were more likely than men to open risky email messages, they were also less likely to click on embedded links, although differences were not statistically significant.

It is also suggested that technical knowledge and experience improve an individual's online security safeguards (Sun et al., 2016). However, the extent that IT competence impacts phishing susceptibility is still not understood. In their scenario-based role-play experiment, Iuga and colleagues (2016) examined the relationship between personal computer 'usage' and phishing detection by asking participants to differentiate legitimate web pages from phishing pages. It was found that those who had been using computers for longer achieved better detection scores.

Pattinson and colleagues (2012) operationalized the notion of computer 'familiarity' by combining the concepts of usage and proficiency and asked their participants how frequently they engage in certain online activities. This variable was tested for both those that were informed of the experiment (i.e., primed to phishing attempts) and those that were not (i.e., the control group). For those that were informed, familiarity correlated significantly with detection rates, and it was determined that those highly familiar with computers were better at managing phishing emails. This was not the case for the control group however, suggesting that individuals need to be actively conscious of phishing in order for their computer familiarity to be relevant.

In both online and offline settings, perceptions of safety alter individual behavior and safety precautions. The literature has broadly examined the influence of perceptions of Internet safety and phishing vulnerability (e.g., Abassi, Zahedi, & Chen, 2016), however, the relationship between 'feeling safe on the Internet' and the actual risk of deception via a 'phish' has not yet been quantified.

Abassi and colleagues' (2016) study questioned the assumption that those who feel unsafe are more likely to be vigilant and employ protective countermeasures as this did not necessarily translate into decreased vulnerability. Their research drew on a sample of 509 university students, staff, and members of the general public from two cities in the United States. The study categorized individuals into clusters based on shared online experiences and analyzed their interactions with fake phishing pages. They found the best detectors of a 'phish' were those that were keenly aware of phishing, familiar with websites, positive about the effectiveness of anti-phishing tools, and had experienced financial losses due to phishing. However, some of these same traits also negatively affected an individual's ability to successfully detect phishing attempts. This was because past encounters and phishing awareness seem to have accentuated an individual's over-confidence in their ability to detect malicious

websites, and familiarity with frequented websites induced over-reliance and trust. The study did not explore how these traits form perceptions of safety, but the results indicate that more robust notions of Internet risk and vulnerability may assist in phishing avoidance.

More generally, Van Wilsem (2013), drawing from a large household panel study (Longitudinal Internet Studies for the Social Sciences [LISS]) of Dutch cybercrime victims, suggested that low self-control (manifested as impulsivity) contributed to online victimisation and was a general risk factor in victimization online or not. An earlier study, drawing also from the LISS panel, noted the overlap between digital and traditional crime showed "... the pervasive influence of online activities on victimization experiences, both on the Internet and in 'traditional' life" (Van Wilsem, 2011, p. 125).

The Present Study

Susceptibility is not homogenous amongst Internet users, as a myriad of factors impact individual vulnerability, judgment, and online behavior. Accordingly, the present study seeks to determine the extent the factors set out above influence the risks of cybercrime for students at the ANU. To accomplish these goals, participants were exposed to various fake email scams, and their interactions with these scams were observed. The observation was conducted over a period of nine months (i.e., February – November, 2017), during which email content was socially engineered to replicate three different types of phishing: generic, tailored, and spear phishing. These required emails to be broad and impersonal, tailored to participants' institution of study, or highly specific to participants' own personal circumstances.

Participants were also compared across two conditions: the 'Hunter' condition and the 'Passive' condition. In the 'Hunter' condition, participants were regularly instructed to be on the lookout for all forms of cybercrime and report any suspicious content to researchers. This condition primed participants to think about the dangers of phishing and was assumed to increase cybercrime awareness. In the 'Passive' condition, no such instructions were received. The number of successful scams (those that participants were deceived by and referred to as the 'scam count'), both overall and for each scam event, provided a measure of susceptibility. Falling for scams was defined as the act of clicking on the fake links provided in the emails. A small pilot study was conducted over several months in 2016 to understand the responses of the ANU Information Security system and involved a sample of 61 students who were recruited to help test various 'spoof' emails. Drawing from the results of the pilot and the literature several hypotheses were tested:

- H1: Scam susceptibility increases as emails became increasingly tailored to the individual. That is, participants were expected to be more likely to be deceived by spear phishing emails than tailored emails, and tailored emails more than generic emails.
- H2: Scam susceptibility varies as a function of cybercrime awareness. The scam count was expected to be lower for Hunter participants, who were primed to remain vigilant for cybercrime. (Note, however, in this study all participants were informed that they were going to receive scam emails, and so the "Hunter" role was a reinforcement rather than a 'primer' or awareness stimulus.)
- H3: An association between gender and scam susceptibility: females were expected to exhibit higher scam susceptibility than males.
- H4: An association between IT competence and scam susceptibility: participants with lower IT competence were expected to exhibit higher scam susceptibility.

H5: An association between perceived Internet safety and scam susceptibility: feeling safe may increase susceptibility.

Method

Participants

Table 1. Characteristics of Sample

Characteristic	Participant (N=138)
<i>Gender</i>	%
Male	50
Female	49.3
Other	0.7
<i>Age</i>	
Under21	64.5
21 - 25	29
26 - 30	3.6
>30	2.9
<i>Student Status:</i>	
Domestic	83.8
International	16.2
<i>Residential Status</i>	
Home	45.6
On Campus	38.2
Other	16.2
<i>Faculty / Study</i>	
Science	29.0
Arts/Social Sciences	25.4
Commerce/Economics	13.8
Science/Engineering	12.3
Law	11.6
Asia Pacific Studies	5.1
Other	1.4
Medicine	0.7
Administration	0.7
<i>Year of Study</i>	
1 Year	53.6
2 Years	17.4
3 Years	11.6
4 Years	11.6
>4 Years	5.8

One hundred and forty-four students from ANU (73 males, 70 females, 1 other) were recruited for this study, and most (53.6%) were commencing their first year of study. Recruitment occurred during orientation week. Students either signed up at a stall belonging to the ANU Criminology Society, or upon being approached by researchers on campus. All participants provided informed, written consent prior to their participation in the study as required by the relevant ANU ethics protocol. Those who

completed the post observational survey received a free hamburger voucher from a popular store as an incentive to complete the follow-up survey.

Data analysis was conducted on a final sample of 138 participants after excluding several due to indecipherable personal details and/or incomplete survey responses. General demographic data, and attitudes to the Internet were obtained from participants via a pre-test and follow-up survey. We asked about gender, age, student (domestic or international) and residential status (home, residential college, other), year of study, and study discipline (course or degree enrolled). An Internet safety component included questions about overall IT competence (54.3% thought they were above average or advanced), social media access (95.7% used social media daily), past experiences with cybercrime (nine respondents reported being a victim of cybercrime), self-reported ability to spot Internet scams (89.8% agreed or strongly agreed that they could detect scams), and feelings about Internet-related safety (87.6% reported being safe or somewhat safe). To reduce respondent burden during field recruitment, the survey format was limited to a single question for each potential variable. The face-to-face consent and pre-observation survey were designed to be completed in less than ten minutes.

Upon finishing the experimental phase, participants were asked to complete a second survey. This involved responding to the same questions as the Internet safety component of the pre-test survey. In addition, participants were asked if they had fallen for any fake scams, whether the study impacted their perceived risk and awareness of cybercrime, and how participating in the study influenced Internet-related behaviors. This information was collected for the purpose of comparing participants' responses at the beginning of the study (Time 1) with their responses at the end of the study (Time 2) and examining the impact of the study on participants' Internet-related attitudes and behaviors.

Software, materials and data recording

This experiment required the re-design of different elements of available software. We needed to manage the creation and distribution of the phishing emails, and design a method for recording data about the interaction participants had with the fake phishing emails. In addition a service that hosted a number of different web sites (copies of legitimate web services) that our participants could visit if deceived by the fake phishing emails. We developed our own software system to enable:

- Use of the university's mail-server to spoof originating email addresses.
- Record of when emails were sent, when they were opened (in some cases), when the participant clicked on one of our fake or 'dodgy' hyperlinks, and when they then entered their credentials into the fake website.
- Web based software to send and monitor these fake emails.

The research team set up a web server that ran an industry standard web server configuration (a LAMP stack) made up of an Apache web server, a MySQL database, and both Python and PHP scripting languages. A framework for sending emails and recording participant responses to those emails was developed using standard scripting languages. The emails were crafted to appear to have been sent by a (fake) person or organization. In order to distribute these 'spoofed' emails, access to an open SMTP server is required. To send these emails, the script would connect to the SMTP server and send the email data. This data includes the sender's email address. Email clients, such as Hotmail, Gmail, and Outlook, include the sender's email address in the emails that they send, however the SMTP standard does not require the correct originating email address. This weakness allows cybercriminals to send emails that appear to have originated from other people or organizations.

During the observation phase, emails were sent to our participants containing a link to a falsified 'Login Page.' The login page was a copy of the university website's student portal login page and, was hosted on the server used in this study. All data was transmitted and received between the server and the participant. Each participant was assigned a unique identifier and every phish email that was sent to a participant contained this unique identifier, allowing any actions taken by participants in response to the phish to be recorded. Three different types of responses were recorded:

- No response. The email never got past the spam filters into the participant's inbox (however, see below regarding web beacon de-activation and non-response ambiguity).
- Received but ignored. The participant opened the email, but chose not to take any action. This may or may not be because they identified the email as fraudulent.
- Received and responded. The participant takes action in response. This could be sending an email in reply, clicking on a link within the email, and/or completing a web form as a result of clicking a link. We did not include downloading and opening an attachment in this study due to the enhanced security associated with spoofing attachments.

The study was not able to use the university Internet Service Provider's service to create a 'white list' for tracking the students participating in the study and alternative means of monitoring had to be devised. We duplicated the login screens of a number of ANU web services, including the email system and the student online management services. They were designed to record the time, date and IP address of each access by our participants, as well as whether or not the participant then proceeded to log into the fake website. We also tracked when a recipient had opened an email by embedding a hidden web beacon into the content of the email. The beacon or website monitor and attached 'cookies' made contact with the study web server every time the email was rendered on a participant's computer screen, and a record was made of the date and time of the contact, and of the IP address of the participant's computer.

It was not straightforward to track when a participant opened an email to read. Techniques such as using a web beacon are well known amongst spammers. The presence of a web beacon, or similar tool, triggers the spam filters employed by many email providers. It is also possible that the web beacon failed to connect to the study web server due to the presence of beacon and/or cookie de-activation software. Moreover, many email clients disable the automatic loading of content when an email is opened. Without this automatic load, the web beacon is unable to signal to the server that the email has been read. Due to this, it was not possible to track emails which had been received but ignored. In this circumstance we are unable to identify if the email was actually read by the participant. Thus, the absence of a record is not conclusive evidence that a participant did not open one of our spoof emails. The data counts are thus conservative.

Operationalizing Scam Susceptibility

The number of fake scams that successfully deceived recipients operationalized scam susceptibility. Scam content was varied across three levels of specificity or individualization via nine fake emails that tested participants' susceptibility:

Generic: the content of fake scams was not personally relevant to participants and replicated real world mass scams. Three common emails were sent, with two of these displaying a 'Mailbox Full' notification and the other alerting the receiver to 'Unread Messages.'

Tailored: the content of fake scams at this level related to the ANU. While these emails were not specific to the individual, they were tailored to the institution and thus provided a mid-point of specificity between generic and spear phishing emails. Four mails purporting to be from ANU's Student Administration included: a notice about changes to the 'Exam Timetable,' an email about a refund from the Higher Education Contribution Scheme (HECS) with subject heading 'HECS Overcharge,' an email about 'Semester 1 Results,' and an email requesting an update of the students record on the Interactive Student Information System (ISIS) with subject heading 'Outdated ISIS Details' (see example below).

Spear-phishing: fake content was made to be personally relevant to the individual. Spear phishers take time and effort to understand their targets in order to maximize the perceived legitimacy of their emails. Such emails may relate not only to relevant institutions, but also to the individual's personal and social lives. Two individual crafted spear phishing mails were sent to a sub-set of participants for whom sufficient personal information was found online. Examples of the scam emails used for each level of specificity are provided below.

Generic Email

Mailbox Full: Upgrade Now

Hi,

Your mailbox is currently at capacity and you are eligible for a free upgrade. Click here to upgrade.

Thanks,

The Outlook Team

Tailored Email

Final Examination Timetable: Update

This is an automatically generated email from an unattended email account; please do not reply

Student ID:

Name:

Dear (insert name),

IMPORTANT: There have been changes to the final examination timetable. Please disregard previous email sent on Friday 28 April 2017.

To access your examination timetable login to ISIS and select 'My Timetable' from the menu on the left.

The examination timetable can be viewed at: <https://exams.anu.edu.au/timetable/>

Further general information about examinations is available at: <http://www.anu.edu.au/students/program-administrati> on/assessments-exams/examination-conduct and <http://www.anu.edu.au/students/program-administration/assessm> ents-exams/examination-timetable. For noting, all examinations are taking place on Acton campus [see campus map at <http://www.anu.edu.au/maps/#>] or at 7-11 Barry Drive, Turner, ACT, 2612 [see map at <http://quicklink.anu.edu.au/xni6>]

Spear Phish Email

From: ANU Sport <sport@anu.edu.au>

Subject: Sportsperson of the Year – Nominated

Good afternoon <participant name>,

Congratulations! We are delighted to let you know that someone has nominated you for the following award: ANU Sportsperson of the Year 2016-17.

We heard that you competed in the Pacific Athletics Championships mid last year. This is an amazing feat that should be celebrated.

If you are interested in officially entering as a nominee, please follow the link and enter your details:

<Link to fake ANU Sports person of the Year Nomination form that requires the following: First name, surname, and email address, mailing address, phone number, ANU student number>

Best,
Mike Brody
Chief Executive Officer - ANU Sport

Procedure

The observational phase occurred over a period of several months. Prior to participation, all participants read an information sheet detailing the study. The voluntary nature of their participation was emphasized, and a consent form was signed (per ANU Ethics Protocol #2015/038). Participants completed the general demographic and Internet safety questionnaire and provided their university identification and email address. Two months after signing up, participants were emailed a reminder that they were part of the study and an opportunity to opt out prior to commencement was provided. Participants were randomly assigned to one of two conditions (Hunter vs. Passive group). Hunters were asked every 4-6 weeks via email to remain constantly vigilant for both fake and real forms of cybercrime, and to forward all suspicious content to the researchers.

Personal information about each participant was extracted, if possible, from their Facebook and LinkedIn profiles, in order to create content for the spear phishing emails. These social media sites provided information about age, current and previous jobs, social relationships, religious and political preferences, hobbies and interests, club memberships and affiliations, and frequently visited locations. After extracting and documenting personal information, a tailored, personally relevant fake attack was created for each participant. For example, searching the Facebook profile of one participant revealed that they had competed in the 2016 Pacific Athletics Championships (a pseudonym). This information allowed for an email impersonating ANU Sport to be created (see above). All spear phishing emails were created in a similar manner and varied depending on the online personal information available. Personal information could not be collected for all participants due to an absence of a social media presence or restricted privacy settings. Specialized emails were created for only 25 participants with adequate online information.


During the observation phase, participants received between seven and nine fake email scams (depending on whether they could be spear phished). All emails attempted to elicit personal information from participants (e.g., university login and password), or attempted to entice participants to click on a fake compromised link. Participants who clicked links or attempted to login to fake pages were redirected to a landing page informing them that they had fallen for a fake attack and reminding them to be more vigilant in future (see Appendix 1).

Results

We first separately examined the effects of the different scam types used, namely Generic, Tailored and Spear Phishing. Altogether three generic and four tailored scams were randomly sent to 138 subjects and two 'spear' or individualized scams were sent to 25 subjects for whom sufficient personal data was obtained from open sources such as Facebook. The total numbers of scams are compiled for each category and we obtain the proportion by normalizing or adjusting the total count by both the number of subjects and number of scams in each category. Participants were most susceptible to a scam with the heading "Final Examination Timetable: Update," which was a scam tailored to the

participants’ university study. Participants were almost equally susceptible to a generic scam titled “Messages.” Figure 1 shows the number of participants who fell for each scam by gender and the number of successful scams by gender is shown in Figure 2.

Figure 1: Number of participants deceived by gender, ordered by scam counts.




Notes: The sample for spear phish attempts is 25 and 135 participants (after removing ‘other’ gender and missing values) for all other scams. Following the scam type, we indicate the level of specificity by G = generic, T = tailored and distinguished from a ‘spear phish’. We note the order of a scam delivery in the observation timeline by 1<2<3<4, where ‘1’ is earlier than ‘2’, which is earlier than ‘3’. For example, ‘Exam Timetable (T1)’ is a scam notifying changes to the exam table that was the first of the tailored scams received by participants.

Overall, there appeared to be an increasing trend in relation to the scam type and scam susceptibility in the normalized proportions in Table 2 with increasing ‘success’ for more individualized and tailored scams. However, a Wilcoxon signed rank test showed these proportions do not differ significantly but the comparison between generic and tailored approached significance ($p = 0.093$, $W = 2785$). Note that a chi-square test is not appropriate for this comparison because Table 2 is not a contingency table and the adjusted proportions do not sum to 1. A t-test is also not adequate for this pair-wise comparison due to the discrete nature of scam counts. Low numbers ($n=25$) for the spear phishing sample significantly reduces the power of the Wilcoxon signed rank test when paired with the corresponding generic and tailored group.

To test the effects of various variables of interest listed in Hypotheses 2 to 5, we fitted a Generalized Linear Model (GLM) with a Poisson error distribution and log link to the response variable total

Figure 2: Number of participants fallen for no scams, 1 scam or more by gender.



scam count as a measure of scam susceptibility. Allowance was made for the fact that only 25 subjects received individualized spear phishing emails. We defined an offset of $\log(7)$ or $\log(9)$ for each subject depending on the total number of scams they were exposed to. Explanatory variables included in this model (Model 1) were the initial hypothesis variables: Gender, IT Competence, Cybercrime Awareness, (the Hunter vs. Passive condition), and Perceived Internet Safety. The likelihood ratio test of Model 1 against the null model gave a non-significant p-value of 0.17.

Table 2. Scam type by number of successful deceptions

Scam type	Generic	Tailored	Spear phishing
Total count	113	165	17
Adjusted proportion	0.27	0.30	0.34

In a further analysis of other variables of interest, the best model is obtained from a stepwise variable selection procedure, which includes only Years of Study (first year university student and later years) and Student Status (international or domestic students) with no significant interaction effect. We call this Model 2. Adding in the hypothesis variables included in Model 1 to Model 2 produced no significant change ($p > .35$). In Model 2, both variables are significant with p-value .012 for Years of Study and p-value 0.017 for Student Status respectively (see Appendix 2 for details of the GLM analysis).

Table 3. Survey responses pre (T1) and post (T2)

Question Percent(%)	Response at T1 (N=138)	Response at T2 (N=85)
<i>Use Social Media Daily:</i>		
Use Facebook	95.7	96.5
Use Instagram	94.7	96.5
Use Snapchat	51.9	69.4
Use Google+	46.6	76.5
Use Twitter	7.6	18.8
Use Tumblr	6.9	30.6
Use LinkedIn	6.1	18.8
Use Other	4.6	23.5
<i>IT Competence(%):</i>		
Poor	3.1	15.5
Adequate	8.7	3.5
Above Average	37.0	28.2
Advanced	44.2	48.2
	10.1	20.0
<i>Cybercrime Victim</i>	6.5	4.7
<i>Can spot cybercrime</i>		
Strongly Disagree	3.6	0.0
Disagree	6.5	4.7
Agree	65.2	71.8
Strongly Agree	24.6	23.5
<i>Purchase Online Goods</i>		
Never	3.6	2.4
Rarely	16.7	21.2
Sometimes	51.4	51.8
Frequently	28.3	24.8
<i>Internet Safety</i>		
Very Unsafe	0.7	1.2
Somewhat Unsafe	11.6	10.6
Somewhat Safe	71.7	71.8
Very Safe	15.9	16.5

^a Average response between T1 and T2 were significantly different ($t_{84} = -2.689, p < .01$). On average, people had lower self-reported IT competence before the study.

The mean scam count was found to differ significantly between domestic and international students ($t = -3.2749, p < .003$). A greater number of international students fell for 3 or more scams than domestic students. Similarly, the mean scam count differed significantly between first and later year students ($t = 3.1724, p < .002$).

A Fisher's Exact Test was used as low cell counts were observed in cross-tabulations between some of the variables investigated. Self-reported IT competence was found to significantly differ by gender (Fisher's Exact Test $p < .001, \phi = .38$). More males rated their IT competence as above average or advanced, while more females reported having only poor or adequate IT competence. Males were also significantly more likely than females to self-report an ability to spot fake scams (Fisher's Exact Test p

$< .05$, $\phi p = .30$). Scam susceptibility, however, was not associated with IT competence nor was it found to significantly differ by perceptions of Internet safety.

Responses to the Internet survey at time 1 (before the observations) and time 2, are reported below in Table 3, however, only 62% of the respondents completed the follow-up survey, limiting the reliability of pre- and post-study differences. Participants rated their IT competence more highly post-study but perceptions of online safety remained largely unchanged although overall there was more diversity in the use of social media.

Discussion

The literature generally suggests that the specificity of a scam may influence cybercrime susceptibility. That is, individuals are more likely to be deceived by scams that are tailored to their personal circumstances compared to those with generic content. To determine whether participants were more susceptible to spear phishing attacks than generic attacks, we used three different scam types: generic, tailored, and spear phishing. Results revealed no significant relationship between scam type and scam susceptibility. However, the email content that deceived most participants provided insight into the types of scams that may succeed. The most successful attack related to an urgent email sent during the exam period about the participants' final exam timetable. This email likely succeeded because it was both relevant and salient, and instilled fear in participants as the email required urgent changes by participants.

The hypothesis that scam susceptibility would vary as a function of awareness about cybercrime was not supported. Despite participants in the Hunter condition being primed to remain vigilant for cybercrime, this did not reduce scam susceptibility. Over several months, Hunters received four emails reminding them about the dangers of cybercrime and prompting them to remain vigilant but only one 'hunter' reported a single suspicious email. This kind of general prompt may have been too weak to raise cybercrime awareness, and thus created minimal differences in awareness between Hunter and Passive conditions. The ineffectual prompting apparent in the present study suggests that increasing the public's level of cybercrime awareness would require constant effort and specific rather than general prompts or warnings about cybercrime.

The gender, IT competence, and perceived Internet safety hypotheses were also not supported. In line with more recent studies (e.g. Butavicius et al. 2017) results from the present study revealed no significant differences in scam susceptibility between male and female participants, low IT competence and high IT competence participants, or participants who rated the Internet as a safe versus unsafe place. This sample was perhaps too small and/or atypical to detect differences even if significant relationships between gender, IT competence, feelings of Internet safety, and cybercrime susceptibility have been identified in other studies (e.g. Iuga et al., 2016; Halevi, Memon, & Oded, 2015).

While none of the initial hypotheses were supported, post-hoc analyses revealed that international students were significantly more susceptible to email scams than domestic students. Although the nature of this relationship is unclear, it is theorized that international students were possibly disadvantaged by language barriers, and/or had different experiences with cybercrime in their countries of origin. Similarly, first year students were significantly more susceptible to email scams than later year students. This may be due to a multitude of factors including age, cybercrime experience, or overall confidence. Perhaps later year students had experienced more real-world scams or may have been more confident in navigating the university email systems compared to first year students. Like international students, first year students were more at risk of cybercrime, suggesting that awareness measures targeted to new and international students would be beneficial. More broadly, exploring the

influence of age and experience on scam susceptibility would allow the nature of this relationship to be better understood. A lack of age variability in the present study prevented this from being examined.

Conclusion

It is important to acknowledge the limitations of this small exploratory study. Firstly, the experimental manipulation (Passive versus Hunter) may not have adequately distinguished levels of cybercrime awareness. It remains important that future research explores how phishing and cybercrime awareness impacts susceptibility to scams.

Secondly, the ability to observe whether emails were actually opened and read was not always feasible, often due to the action of web beacons. Consequently, it was unknown during the initial phase of the study whether participants were actively identifying the emails as attacks or simply ignoring them. Opening an email and identifying it as a scam is different from ignoring the content entirely; however, the inability to distinguish between these actions meant that participants who ignored the emails altogether were also treated as less susceptible to fake scams. Therefore, our interpretation of non-response was conditional because it was not always possible to distinguish between an unread or unopened email. Our observation is thus limited to what action our respondents took, if any, in respect to the response demand of the phish.

Finally, the present study did not account for practice effects. The second of our generic 'Mailbox Full' scams deceived (n=16) half as many participants than the first round (n=34) suggesting that a practice effect may be in play. Each time a participant was deceived by our fake phishing mail they were directed to a web-landing page (see Appendix 1). This informed them that they had been deceived and offered cyber-safety advice. Thus repeated practice with phishing emails may have overshadowed the influence of different scam types. While results did not reveal an overall decrease in susceptibility over time, it would have helped to distinguish between the effects of scam types and the role of practice. This would have allowed results to be attributed confidently to the experimental manipulation of scam type and could have shed light on whether repeated exposure to fake scams increased cybercrime awareness and decreased cybercrime susceptibility. Observing the presence of practice effects could provide information about how to teach and increase cybercrime awareness (see Canfield, Fischhoff, & Davis, 2016).

Future research could apply a more robust quasi-experimental design to determine the variables that influence scam susceptibility. Understanding the factors that influence susceptibility will help to protect against phishing and other forms of cybercrime. While the present study was exploratory, our attempt to observe cybercrime victimization in a real-world setting may be scaled up with larger samples and a greater variety of social engineering methods.

Acknowledgement

We are grateful to the Australian Criminology Research Council (Grant # CRG 51 16-17) for funding this research. We thank Charlotte Ho Chung, Ross Maller, Donald Maxim, Bianca Sabol, Khoi-Nyguen Tran, Hannah Woodford-Smith, and ANU IT Security for their assistance. We also thank the anonymous reviewers for their helpful suggestions.

References

Abassi, A., Zahedi, F. M., & Chen, Y. (2016). Phishing susceptibility: The good, the bad, and the ugly. *2016 IEEE Conference on Intelligence and Security Informatics*, 169-174, Tucson: IEEE.

- Alazab, M., & Broadhurst, R. (2016). Spam and Criminal Activity. *Trends & Issues in Crime and Criminal Justice*, no. 526, Canberra: Australian Institute of Criminology.
- Alsharnouby, M., Alaca, F., & Chiasson, S. (2015). Why phishing still works: Use strategies for combating phishing attacks. *International Journal of Human Computer Studies*, 82,69-82.
- Benenson, Z., Gassmann, F., & Landwirth, R. (2016). *Exploiting curiosity and context: How to make people click on a dangerous link despite their security awareness*, viewed 15 January 2018, retrieved from: <https://paper.seebug.org/papers/Security\%20Conf/Blackhat/2016/us-16-Benenson-Exploiting-Curiosity-And-Context-How-To-Make-People-Click-On-A-Dangerous-Link-Despite-Their-Security-Awareness-wp.pdf>
- Butavicius, M., Parsons, K., Pattinson, M., McCormac, A., Calic, D., & Lillie, M. (2017). Understanding susceptibility to phishing emails: Assessing the impact of individual differences and culture. *Proceedings of the Eleventh International Symposium on Human Aspects of Information Security & Assurance*, 12-23.
- Butavicius, M., Parsons, K., Pattinson, M., & McCormac, A. (2015). Breaching the human firewall: Social engineering in phishing and spear phishing emails. *Australasian Conference on Information Systems*, 12-23, Adelaide: ACIS.
- Canfield, C. I., Fischhoff, B., & Davis, A. (2016). Quantifying Phishing Susceptibility for Detection and Behaviour Decisions. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 58 (8), 1158-1172.
- 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.
- Chaudhry, J. A., Chaudhry, S. A., & Rittenhouse, R. G. (2016). Phishing attacks and defenses. *International Journal of Security and its Applications*, 10(1), 247-256.
- De Kimpe, L., M. Walrave, W., Hardyns, L. Pauwels & K. Ponnet (2018). 'you've got mail! Explaining individual differences in becoming a phishing target. *Telematics and Informatics*, <https://doi.org/10.1016/j.tele.2918.02.009>
- Gavett, B. E., Zhao, R., John, S. E., Bussell, C. A., Roberts, J. R., & Yue, C. (2017). Phishing suspiciousness in older and younger adults: The role of executive functioning. *PLOS ONE*, 12(2), 1-16.
- Goel, S., Williams, K., & Dincelli, E. (2017). Got phished? Internet security and human vulnerability. *Journal of the Association for Information Systems*, 18(1), 22-44.
- Goldsmid, S., Gannoni, A., & Smith, R.G. (2018). *Identity crime and misuse in Australia: Results of the 2017 online survey*. Statistical Report 11, Australian Institute of Criminology, Canberra.
- Gratian, M., Bandi, S., Cukier, M., Dykstra, J., & Ginther, A. (2018). Correlating human behaviour and cyber security behaviour intentions. *Computers & Security*, 73, 345-358.
- Halevi, T., Memon, N. & Oded, N. (2015). Spear-Phishing in the wild: A real-world study of personality, phishing self-efficacy and vulnerability to spear-phishing attacks, viewed 25 January 2015, retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2544742.
- Iuga, C., Nurse, J. R. C., & Erola, A. (2016). Baiting the hook: Factors impacting susceptibility to phishing attacks. *Human-Centric Computing and Information Sciences*, 6(1:8), 1-20.

- Mayhorn, C. B., Welk, A. K., Zielinska, O. A., & Murphy-Hill, E. (2015). Assessing individual differences in a phishing detection task. *Proceedings of the 19th Triennial Congress of the IEA*, Melbourne: IEA.
- Oliviera, D., Rocha, H., Yang, H., Ellis, D., Dommaraju, S., Muradoglu, M., . . . & 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. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 6412-6424, Denver: ACM.
- Pattinson, M., Jerram, C., Parsons, K., McCormac, A., & Butavicius, M. (2012). Why do some people manage phishing e-mails better than others? *Information Management & Computer Security*, 20(1), 18-28.
- Parsons, K., McCormac, A., Pattinson, M., Butavicius, M., & Jerram, C. (2015). The design of phishing: Challenges for researchers. *Computers & Security*, 52, 194-206.
- Sun, J. C. Y., Yu, S. J., Lin, S. S. J., Tseng, S. S. (2016). The mediating effect of anti phishing self-efficacy between college students' internet self-efficacy and anti phishing behaviour and gender difference. *Computers in Human Behaviour*, 59, 249-257.
- Symantec (2014), *Internet Security Threat Report 2014*, viewed 25 January 2018, retrieved from http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018_en-us.pdf.
- Talos (2018), *Email & Spam Data*, viewed 25 January 2016, retrieved from https://www.talosintelligence.com/reputation.center/email_rep#global-volume.
- Van Wilsem, J. V. (2013). Hacking and harassment—Do they have something in common? Comparing risk factors for online victimization, *Journal of Contemporary Criminal Justice*, 29(4), 437-453.
- Van Wilsem, J. V. (2011). Worlds tied together? Online and non-domestic routine activities and their impact on digital and traditional threat victimization. *European Journal of Criminology*, 8(2), 115-127.

Appendix 1: Landing Page



Oops! This was a test from the ANU Cybercrime Observatory Experiment “Wi-Fi Usage and Cybercrime Risks in University Student Communities”.

We are testing participants’ susceptibility to spear phishing and scam/spam attempts.

This notification indicates that you have fallen for a FAKE spam/phishing attempt. Don’t worry, as this was not a legitimate spam/phishing attempt any identifying information you may have provided is not compromised. Your participation in this study is of great value for our research as we attempt to identify where our on-line vulnerabilities lie and how we can protect ourselves.

PLEASE NOTE: during this research phase you may in fact be exposed to ‘real’ scams/spam/phishing attempts. Remain vigilant and be careful about your online activity.

For useful cyber safety information visit http://dmm.anu.edu.au/7JPtR_cybersafety/package.php

For queries regarding the study you are welcome to contact roderic.broadhurst@anu.edu.au or cyberobs.anu@gmail.com with the subject line as “Queries about Wi-Fi usage study (protocol number: 2015/038)”

Appendix 2: GLM Results

The analysis is performed with R version 3.5.1 (2018-07-02). Three participants with missing values are omitted from the GLM analysis, one of them had indicated Gender as ‘Others’ and two others have missing Student Status. All variables included in Model 1 are listed below:

- Y_i represents the number of scams the i^{th} individual has fallen for: 0-7;
- X_{1i} is the indicator of the i^{th} individual’s gender: male or female;
- X_{2i} represents the i^{th} individual’s IT competence: Poor < Ave. < Above Ave. < Advance;
- X_{3i} is the indicator of the i^{th} individual’s Cybercrime Awareness: Hunter or Passive;
- X_{4i} represents the i^{th} individual’s perceived internet safety: Very Unsafe < Somewhat Unsafe < Somewhat safe < Very safe;
- $Offset_i$ adjusts for the total number of scams the i^{th} individual is exposed to: log(7) or log(9).

Model 1 regression equation

$$\log(E(Y_i)) = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \beta_3 X_{3i} + \beta_4 X_{4i} + offset_i$$

Table 2.1: Analysis of deviance table of Model 1*

	Residual df	Residual deviance	Additional df	Change in deviance	p-value
Null model	134	154.48			
Model 1	126	142.78	8	11.696	0.1653

*Note: Analysis of deviance table for Model 1 compared to the null model where only the intercept is fitted with the offset; note df = degree of freedom.

The likelihood ratio test of Model 1 against the null model where only the intercept is fitted gives a p-value of 0.1653. Thus, there are no significant effects overall in Model 1. Post-hoc analysis identified two additional variables that are highly correlated to scam count, they are

- W_{1i} is the i^{th} individual’s year of study: Year 1, 2, 3, 4, 5 (treated as numeric);
- W_{2i} is the indicator of the i^{th} individual’s residential status: domestic or international.

Both forward and backward stepwise variable selection procedures are performed on all variables mentioned above based on the Akaike Information Criterion (AIC). The final best model we call Model 2 and is described next (see below).

Model 2 regression equation

$$\log(E(Y_i)) = \beta_0 + \beta_1 W_{1i} + \beta_2 W_{2i} + offset_i$$

Table 2.2: Analysis of deviance table of Model 2*

	Residual df	Residual Deviance	Additional df	Change in Deviance	p-value
Null model	134	154.48			
Model 2	132	136.68	2	17.7918	0.0001369
Full model	124	127.86	8	8.8271	0.3570867

Note*: Analysis of deviance table for Model 2 compared to the null model and the full model where all hypothesis variables are included.

Table 2.3: Summary of coefficients in Model 2*

	Estimate	Standard Error	Z-value	p-value
intercept(β_0)	-1.0351	0.1250	-8.282	0
years of study(β_1)	-0.1381	0.0549	-2.516	0.0119
status: international(β_2)	0.3453	0.1447	2.387	0.0170

*Both variables in Model 2 are significant at 0.05 level.