Australian Government
Australian Institute of Criminology

Trends & issues in crime and criminal justice

ISSN 1836-2206 (Online) | ISBN 978 1 925304 25 1 (Online)

No. 587 February 2020

Abstract | In an exploratory quasi-experimental study, 138 students recruited during a university orientation week were exposed to social engineering directives in the form of fake emails, or phishing, over several months in 2017. The study assessed the risks of cybercrime for students by observing their responses. Three types of scam emails were distributed that varied in the degree of individualisation: generic, tailored, and targeted or 'spear'. The study explored the influence of scam type, cybercrime awareness, gender, IT competence and perceived internet safety on susceptibility to email scams.

Although tailored and individually crafted email scams were more likely to induce engagement than generic scams, differences were not significant. Analysis of the variables showed that international students and first year students were deceived by significantly more scams than domestic students and later year students.

Phishing risks in a university student community

Roderic Broadhurst, Katie Skinner,
Nicholas Sifniotis, Bryan Matamoros-Macias
and Yuguang Ipsen

As individuals become increasingly connected to the virtual world, the avenues for exploitation by cybercriminals also increase. Although developments in technology have attempted to mitigate these risks, human error continues to be the weakest link in cybersecurity (Mayhorn et al. 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 judgement rather than their virtual security measures (Alazab & Broadhurst 2016; Gratian et al. 2018). 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, the worldwide average daily volume of email spam was approximately 422 billion in January 2018, constituting about 85 percent of all daily global email traffic (Talos 2018).

CRG | CRIMINOLOGY RESEARCH GRANT

The ubiquitous threat of malware-laden spam has significant economic and social consequences, but victimisation and susceptibility vary. Previous studies have suggested that gender (Sun et al. 2016), age (Gavett et al. 2017), and technical experience (Pattinson et al. 2012) influence an individual's susceptibility to spam and phishing attempts. This study was designed to explore how these factors—and the type of scam, cybercrime awareness and IT competence—influenced susceptibility among a sample of university students. To accomplish this, participants were exposed to various fake emails sent from the research team's web server, and their interactions with these scams were observed.

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 a catch. The lure often involves an email message appearing to be from a legitimate person or organisation, 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; and
- empathy—such as emails impersonating a friend or relative who needs financial assistance or personal information.

This list is not exhaustive and can be augmented by appeals to other emotions such as greed (eg a winning lottery ticket), lust, or vanity (eg an adoring admirer or a prestigious job opportunity). De Kimpe et al. (2018) list characteristics that can either facilitate or hinder the success of phishing emails (eg the presence of spelling, design or formatting errors, or an offer of prize money). Once recipients are convinced the mail is authentic, the next stage is to convince them to divulge sensitive information. Various social manipulators—such as liking or trusting the email source, implicating reciprocity (eg returning favours) or 'social proof' (ie others are participating), creating a sense of scarcity, or evoking an authoritative source—help the deception to succeed.

When phishing emails make use of personalised data in their lures, they become examples of 'spear-phishing.' Spear-phishing emails often contain information that would be familiar or important to specific recipients (De Kimpe et al. 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 et al. 2014). These emails tend to impersonate well-known companies, trusted relationships or contexts that have personal relevance to the individual (De Kimpe et al. 2018).

The success of any phishing email is linked to how well it is able to deceive its recipient. While the research literature has focused on phishing email structure (eg use of visual cues and the presence of misspellings or attachments: Parsons et. al. 2015), this study explores email contextualisation and personalisation. Phishing emails containing personalised information have been shown to be effective in deceiving their targets (Benenson, Gassmann & Landwirth 2016).

Butavicius et al. (2015) tested the effects of different social engineering strategies by sending a series of genuine, phishing, or spear-phishing emails to a group of 117 university students. The results found that students tended to classify emails as genuine rather than fraudulent and were worse at detecting spear-phishing attempts than generic phishing attempts. Where spear-phishing emails used an authority-style social engineering strategy (ie the apparent sender of the email held authority over the reader), students were less able to detect spear-phishing.

Variables associated with phishing risk

Individuals, once aware of their own potential victimisation, are thought to become more cautious or defensive as they navigate risky environments. Understanding the impact of participant 'priming' or awareness of potential risk can inform the development of programs aimed at preventing online victimisation.

Participants who know 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 victimisation has been questioned (Alsharnouby, Alaca & Chiasson 2015; Caputo et al. 2014). In the present study our subjects were all primed, as ethical approval required offline formal consent for attempts to deceive them with a scam email. However, a sub-sample of our subjects (designated as 'hunters') were 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), while others have found no such connection (Butavicius et al. 2017; Oliveira et al. 2017). Despite these contradictory findings, recent studies have sought to reframe the relationship between gender and phishing susceptibility. In Goel, Williams and Dincelli (2017), the act of falling for a phishing scam was framed as two steps: first, the opening of a phishing email and, second, the clicking of the malicious link within. They found that while 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.

Technical knowledge and experience should improve an individual's online security safeguards (Sun et al. 2016); however, the extent to which IT competence affects phishing susceptibility is unclear. In their scenario-based role-play experiment, Iuga, Nurse and Erola (2016) examined the relationship between personal computer use and phishing detection by asking participants to differentiate between legitimate web pages and phishing pages. It was found that those who had been using computers for longer achieved better detection scores.

Pattinson and colleagues (2012) operationalised the notion of computer familiarity by combining usage and proficiency, and asked their participants how frequently they engaged in certain online activities. This variable was tested both for those who were informed of the experiment (ie primed to phishing attempts) and for those who were not (the control group). For those who 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 behaviour and safety precautions. The literature has broadly examined the influence of perceptions of internet safety and phishing vulnerability (eg Abbasi, 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.

The present study

Susceptibility is not homogeneous among internet users, as myriad factors impact individual vulnerability, judgement, and online behaviour. Accordingly, the present study seeks to determine the extent to which the factors set out above influence the risks of cybercrime for students at the Australian National University (ANU). To accomplish this goal, 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 from February to November 2017, during which email content was socially engineered to replicate three different types of phishing: generic, tailored, and ‘spear’. These required emails to be either broad and impersonal; tailored to participants’ institution of study; or highly specific to a participant’s 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 to 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, referred to as the ‘scam count’), both overall and for each scam event, provided a measure of susceptibility. Falling for a scam was defined as the act of clicking on a fake link provided in an email.

Several hypotheses were tested:

- H1—scam susceptibility increases as emails become increasingly tailored to the individual.
- 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. (However, since all participants were informed that they would receive scam emails, the hunter role was a reinforcement rather than a primer of awareness.)
- H3—there is an association between gender and scam susceptibility: females were expected to exhibit higher scam susceptibility than males.
- H4—there is an association between IT competence and scam susceptibility: participants with lower IT competence were expected to exhibit higher scam susceptibility.
- H5—there is an association between perceived internet safety and scam susceptibility: feeling safe may increase susceptibility.

Method

Participants

One hundred and forty-four students from ANU (73 males, 70 females, 1 other) were recruited for this study, and most (54%) were commencing their first year of study. Recruitment occurred during orientation week. Students signed up either 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 ethics protocol. Those who completed the post-observational survey received a free hamburger voucher, offered as an incentive to complete the survey.

Data analysis was conducted on a final sample of 138 participants after several were excluded due to indecipherable personal details and/or incomplete survey responses. Demographic data and attitudes to the internet were obtained from participants via a pre-test survey and a follow-up survey. We asked about gender, age, student status (domestic or international) and residential status (home, residential college or other), year of study, and study discipline (course or degree enrolled in). An internet safety component included questions about overall IT competence (54% thought they were above average or advanced), social media access (96% used social media daily), past experiences with cybercrime (nine respondents reported being a victim of cybercrime), self-reported ability to spot internet scams (90% agreed or strongly agreed that they could detect scams), and feelings about internet-related safety (88% reported being safe or somewhat safe). The face-to-face consent and pre-observation survey were designed to be completed in less than ten minutes in total.

Participants were asked to re-take the same survey at the completion of the phishing phase. In addition, participants were asked if they had fallen for any fake scams, whether the study had impacted on their perceived risk and awareness of cybercrime, and how participating in the study had influenced internet-related behaviours. This information was collected for the purposes 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 behaviours.

Table 1: Sample characteristics (%)

Gender		**Faculty/study**	
Male	50.0	Science	29.0
Female	49.3	Arts/social sciences	25.4
Other	0.7	Commerce/economics	13.8
Age		Science/engineering	12.3
Under 21	64.5	Law	11.6
21–25	29.0	Asia–Pacific studies	5.1
26–30	3.6	Medicine	0.7
>30	2.9	Administration	0.7
Student status		Other	1.4
Domestic	83.8	**Years of study**	
International	16.2	1 year	53.6
Residential status		2 years	17.4
Home	45.6	3 years	11.6
On campus	38.2	4 years	11.6
Other	16.2	>4 years	5.8

Software, materials and data recording

This experiment required the redesign of available software. We needed to manage the creation and distribution of the phishing emails, and design a method for recording data about the interactions participants had with the fake phishing emails. We created a service that hosted fake websites (copies of legitimate web services) that our participants would visit if deceived by the fake phishing emails. We developed a system to enable:

- use of the university's mail server to spoof originating email addresses;
- records of when emails were sent, if and when they were opened, if and when a participant clicked on one of our fake or 'dodgy' hyperlinks, and if and when they then entered their credentials into the fake website; and
- web-based software to send and monitor these fake emails.

The emails were crafted to appear to have been sent by a (fake) person or organisation. In order to distribute these spoofed emails, access to an open SMTP server was required. To send these emails, the script would connect to the SMTP server and send the email data. These data included 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 originating email address to be correct. This weakness allows cybercriminals to send emails that appear to have originated from other people or organisations.

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 were 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 and linked to that individual. 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);
- 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; and
- received and responded—the participant took 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 opening an attachment in this study due to the enhanced security associated with spoofing attachments.

We duplicated the login screens of a number of ANU web services, including the email system and the online student 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 in to 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 contacted 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 and read it. Techniques such as using a web beacon are well known among spammers. The presence of a web beacon may have triggered 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 would have been unable to signal to the server that the email had been read. This meant it was not possible to track emails which had been received but ignored. We were therefore unable to identify whether the email was actually read by the participant. The data counts are thus conservative.

Operationalising scam susceptibility

The number of fake scams that successfully deceived recipients operationalised scam susceptibility. Scam content was varied across three levels of individualisation 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. The four emails purporting to be from ANU's Student Administration were:
  - 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 student's record on the Interactive Student Information System (ISIS) with the subject heading 'Outdated ISIS Details'.
- 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 maximise the perceived legitimacy of their emails. Such emails may relate not only to relevant institutions but also to an individual's personal and social life. Two individually crafted spear-phishing mails were sent to a subset of participants for whom sufficient personal information was found online.

Procedure

The observational phase occurred over a period of several months. Before participating, all participants read an information sheet detailing the study. The voluntary nature of their participation was emphasised, and a consent form was signed. 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 was provided to opt out prior to commencement. Participants were randomly assigned to one of two conditions (hunter or passive). Hunters were asked every four to six weeks via email to remain 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, memberships and affiliations, and frequently-visited locations. After documenting the 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. 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. Specialised emails were only created for the 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 (eg university login and password) or to entice participants to click on a fake link. Participants who clicked links or attempted to log in 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.

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 individualised scams were sent to the 25 subjects for whom sufficient personal data was obtained from open sources such as Facebook. The total numbers of scams were compiled for each category and we obtained the proportion by normalising 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 especially 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.

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

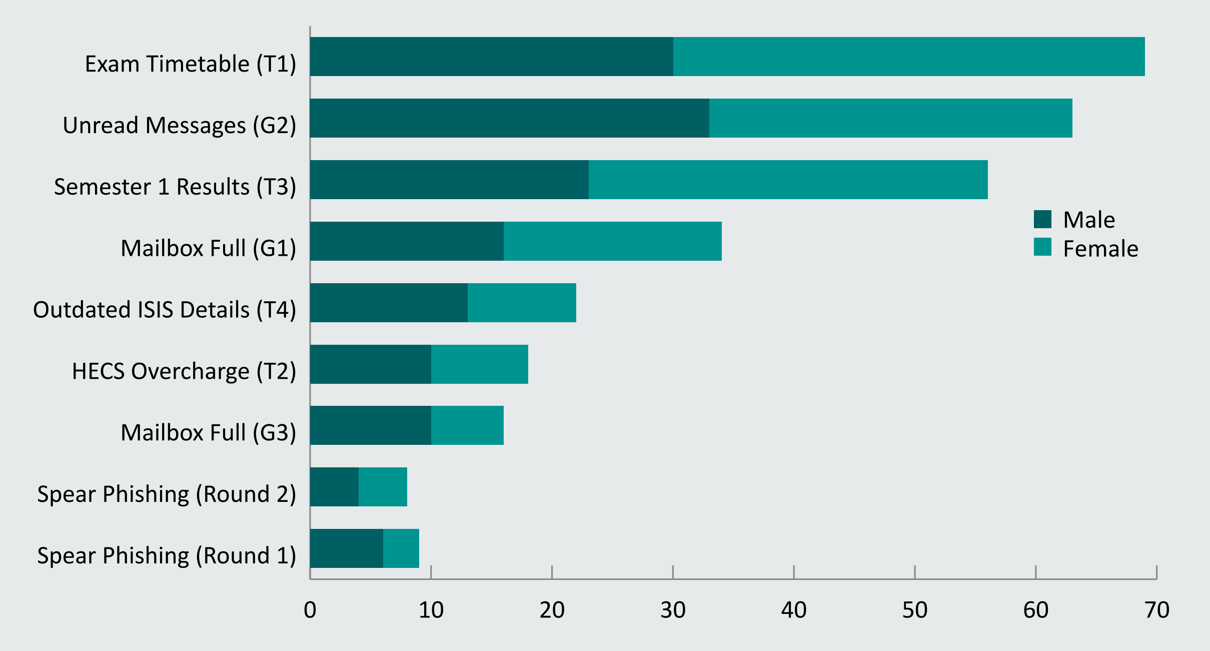

Note: The sample size for spear-phishing attempts is 25 participants and for all other scams is 135 participants (after removing 'other' gender and missing values). Following the scam type, we indicate the level of specificity by G=generic, T=tailored as distinct from spear-phishing. We note the order 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 of changes to the exam table that was the first of the tailored scams received by participants

Overall, there appeared to be a trend in relation to the scam type and susceptibility, with increasing success for more individualised and tailored scams. However, a Wilcoxon signed-rank test showed these differences were not significant, although a comparison between generic and tailored emails approached significance (p=0.093, W=2785). Low numbers (n=25) for the spear-phishing sample significantly reduced the power of the Wilcoxon signed-rank test when paired with the corresponding generic or tailored group.

To test the effects of variables of interest listed in hypotheses 2, 3, 4 and 5, we fitted a generalised linear model with a Poisson error distribution and log link to the response variable 'total scam count' as a measure of scam susceptibility. Allowance was made for the fact that only 25 subjects received individualised 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 (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.

In a further analysis of other variables of interest, the best model was obtained from a stepwise variable selection procedure, which included only years of study (first year or later year university student) and student status (international or domestic) with no significant interaction effect. We called this Model 2. Adding in the hypothesis variables included in Model 1 to Model 2 produced no significant change. In Model 2, both variables were significant, with a p value of 0.012 for years of study and a p value of 0.017 for student status.

The mean scam count was found to differ significantly between domestic and international students (t=-3.2749, p<0.003). A greater number of international students fell for three or more scams than domestic students. Similarly, the mean scam count differed significantly between first year and later year students (t=3.1724, p<0.002).

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 differ significantly by gender (Fisher's exact test p<0.001, ϕp=0.38). 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<0.05, ϕp=0.30). 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 in Table 2; however, only 62 percent 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 there was more diversity in the types of social media used by respondents at time 1 compared with time 2.

Table 2: Survey responses pre-study (T1) and post-study (T2)

Question	Response at T1 (%) (*n*=138)	Response at T2 (%) (*n*=85)
Use social media daily	95.7	96.5
Use Facebook	94.7	96.5
Use Instagram	51.9	69.4
Use Snapchat	46.6	76.5
Use Google+	7.6	18.8
Use Twitter	6.9	30.6
Use Tumblr	6.1	18.8
Use LinkedIn	4.6	23.5
Use other	3.1	15.5
IT competence[a]		
Poor	8.7	3.5
Adequate	37.0	28.2
Above average	44.2	48.2
Advanced	10.1	20.0
Cybercrime victim	6.5	4.7
Can spot cybercrime		
Strongly disagree	3.6	0.0
Disagree	6.5	4.7
Agree	65.2	71.8
Strongly agree	24.6	23.5

Table 2: Survey responses pre-study (T1) and post-study (T2)

Question	Response at T1 (%) (*n*=138)	Response at T2 (%) (*n*=85)
Purchase online goods		
Never	3.6	2.4
Rarely	16.7	21.2
Sometimes	51.4	51.8
Frequently	28.3	24.8
Internet safety		
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 responses between T1 and T2 were significantly different (*t*84=-2.689, *p*<0.01). On average, people had lower self-reported IT competence before the study

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 than by scams 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 scam 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 attention.

The hypothesis that scam susceptibility would vary as a function of awareness of cybercrime was not supported. Despite participants in the hunter condition being primed to remain vigilant for cybercrime, this did not reduce their 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 general prompt may have been too weak to raise cybercrime awareness, and thus created minimal differences in awareness between the hunter and passive conditions of this study. This 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 (eg Butavicius et al. 2017), results from the present study revealed no significant differences in scam susceptibility between male and female participants; between participants with low IT competence and high IT competence; or between participants who rated the internet as a safe versus an 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 (eg Halevi, Memon & Nov 2015; Iuga, Nurse & Erola 2016).

While none of the initial hypotheses was 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 theorised 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 factors including age, cybercrime experience and overall confidence. Perhaps later year students had experienced more real-world scams, or they 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 at new and international students would be beneficial.

Conclusion

It is important to acknowledge the limitations of this small exploratory study. Firstly, the experimental manipulation (passive versus hunter conditions) may not have adequately distinguished levels of cybercrime awareness. Future research should explore how phishing and cybercrime awareness impacts susceptibility to scams.

Secondly, observing whether emails were actually opened and read was not always feasible, often due to the action of web beacons. It was therefore unknown during the initial phase of the study whether participants were actively identifying the emails as attacks or simply ignoring them. This meant our interpretation of non-response was conditional, because it was not always possible to distinguish between an unread and an unopened email. Our observation is thus limited to what action our respondents took, if any, in response to the demand of the phish.

Finally, the present study did not account for practice effects. The second of our generic 'Mailbox Full' scams deceived half as many participants (*n*=16) as 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 landing page. This informed them that they had been deceived and offered cybersafety advice. 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 be able 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 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).

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 victimisation in a real-world setting may be scaled up with larger samples and a greater variety of social engineering methods.

References

URLs correct as at November 2019

Abbasi A, Zahedi FM & Chen Y 2016. *Phishing susceptibility: The good, the bad, and the ugly.* 2016 IEEE Conference on Intelligence and Security Informatics. Tucson: IEEE: 169–74. https://doi.org/10.1109/ISI.2016.7745462

Alazab M & Broadhurst R 2016. Spam and criminal activity. *Trends & issues in crime and criminal justice* no. 526. Canberra: Australian Institute of Criminology. https://aic.gov.au/publications/tandi/tandi526

Alsharnouby M, Alaca F & Chiasson S 2015. Why phishing still works: User 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. Paper to Black Hat USA 2016 conference, Las Vegas, 30 July–4 August. 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 2015. *Breaching the human firewall: Social engineering in phishing and spear phishing emails*. Australasian Conference on Information Systems 2015 Proceedings. Adelaide: ACIS: 12–23

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. University of Plymouth: 2017: 12–23

Canfield CI, 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): 1158–72

Caputo DD, Pfleeger SL, Freeman JD & Johnson ME 2014. Going spear phishing: Exploring embedded training and awareness. *IEEE Security & Privacy* 12(1): 28–38

Chaudhry JA, Chaudhry SA & Rittenhouse RG 2016. Phishing attacks and defenses. *International Journal of Security and its Applications* 10(1): 247–56

De Kimpe L, Walrave M, Hardyns W, Pauwels L & Ponnet K 2018. You've got mail! Explaining individual differences in becoming a phishing target. *Telematics and Informatics* 35(5): 1277–87. http://hdl.handle.net/1854/LU-8554543

Gavett BE, Zhao R, John SE, Bussell CA, Roberts JR & 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

Gratian M, Bandi S, Cukier M, Dykstra J & Ginther A 2018. Correlating human behaviour and cyber security behaviour intentions. *Computers & Security* 73: 345–58

Halevi T, Memon N & Nov O 2015. *Spear-phishing in the wild: A real-world study of personality, phishing self-efficacy and vulnerability to spear-phishing attacks*. https://dx.doi.org/10.2139/ssrn.2544742

Iuga C, Nurse JRC & Erola A 2016. Baiting the hook: Factors impacting susceptibility to phishing attacks. *Human-Centric Computing and Information Sciences* 6(1:8): 1–20

Mayhorn CB, Welk AK, Zielinska OA, Murphy-Hill E 2015. *Assessing individual differences in a phishing detection task.* Proceedings of the 19th Triennial Congress of the IEA. Melbourne: IEA: np

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.* Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. Denver: ACM: 6412–6424

Parsons K, McCormac A, Pattinson M, Butavicius M & Jerram C 2015. The design of phishing: Challenges for researchers. *Computers & Security* 52: 194–206

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

Sun JCY, Yu SJ, Lin SSJ & Tseng SS 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: 249–57

Talos 2018. Email & spam data. https://www.talosintelligence.com/reputation_center/email_rep#global-volume

Roderic Broadhurst, Professor of Criminology, Australian National University School of Regulation and Global Governance

Katie Skinner, Research Assistant, Australian National University Cybercrime Observatory

Nicholas Sifniotis, Research Assistant, Australian National University Cybercrime Observatory

Bryan Matamoros-Macias, Research Assistant, Australian National University Cybercrime Observatory

Yuguang Ipsen, Lecturer, Research School of Finance, Actuarial Studies and Statistics, Australian National University

General editor, *Trends & issues in crime and criminal justice* series: Dr Rick Brown, Deputy Director, Australian Institute of Criminology. Note: *Trends & issues in crime and criminal justice* papers are peer reviewed. For a complete list and the full text of the papers in the *Trends & issues in crime and criminal justice* series, visit the AIC website at: aic.gov.au

ISSN 1836-2206 (Online)

ISBN 978 1 925304 25 1 (Online)

©Australian Institute of Criminology 2020

GPO Box 1936
Canberra ACT 2601, Australia

Tel: 02 6268 7166

Disclaimer: This research paper does not necessarily reflect the policy position of the Australian Government

aic.gov.au