CRIMINOLOGY
RESEARCH GRANT

Phishing and cybercrime risks in a university student community

Roderic Broadhurst
Katie Skinner
Nick Sifniotis
Bryan Matamoros-Macias
Yuguang Ipsen

Report to the Criminology
Research Advisory Council
Grant: CRG 51/16–17

February 2020

This project was supported by a Criminology Research Grant.
The views expressed are the responsibility of the author and are not necessarily those of the Council.

Published by the Australian Institute of Criminology
GPO Box 1936 Canberra ACT 2601
Tel: (02) 6268 7166
Email: front.desk@aic.gov.au
Website: crg.aic.gov.au

ISBN 978 1 925304 26 8 (Online)

This research was supported by a Criminology Research Grant. The views expressed are those of the author and do not necessarily reflect the position of the Criminology Research Advisory Council or the Australian Government.

This report was reviewed through a double-blind peer review process.

Edited and typeset by the Australian Institute of Criminology.

Contents

Figures

Tables

Acknowledgements

We gratefully acknowledge the support of the Criminology Research Advisory Council (CRG grant no. 51/16–17) in funding this research project. We also thank Bianca Sabol, Charlotte Ho Chung, Donald Maxim, Nguyen Tran and Hannah Woodford-Smith for their assistance.

Abstract

In a quasi-experimental study, 138 students recruited during a university orientation week were exposed to social engineering directives in the form of fake phishing emails 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.

Executive summary

In an exploratory quasi-experimental observational study, 138 student participants recruited during a university orientation week were exposed to social engineering in the form of fake email phishing spam from February to November 2017. Unsolicited emails or spam can involve harmless advertising, text messages or social network messages, but spam may also contain viruses or malware designed to exploit personal or sensitive information from its recipients. These spam-like phishing email attacks attempted to elicit personal information from participants, or to entice them to click on links which in a real-world setting may have been compromised.

The success of any phishing email depends on 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 context and personalisation. Phishing emails containing personalised information have been shown to be effective in deceiving their targets (Benenson, Gassmann & Landwirth 2016).

The study aimed to determine the risks of cybercrime for Australian university students by observing their responses to phishing emails and by exploring attitudes to cybercrime risks before and after the phishing phase. 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;
- 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; and
- H5—there is an association between perceived internet safety and scam susceptibility: feeling safe may increase susceptibility.

The experiment was conducted in 2017 over a period of nine months from February to November. During this time, email content was socially engineered to replicate three different types of phishing: generic, tailored, and spear-phishing. This meant the respective emails had to be either broad and impersonal, tailored to participants' institution of study, or highly specific to a participant's personal circumstances. To differentiate participants on the basis of cybercrime awareness, some participants were primed throughout the study to remain vigilant to all scams (the 'hunter' condition) while others received no such instruction (the 'passive' condition).

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 explored 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.

Contrary to the hypotheses, none of these factors were associated with scam susceptibility. 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, these differences were not significant, although a comparison between generic and tailored emails approached significance. Low numbers (*n*=25) for the spear-phishing group reduced the power of statistical tests comparing the generic and tailored groups. However, tailored and 'spear' type scams were more likely to induce engagement than generic emails.

Post-hoc analysis of all the variables showed that international students and first year students fell for significantly more scams than domestic students and later year students. A generalised linear model analysis was undertaken to further explore the role of all the variables of interest. The results were consistent with the descriptive findings showing that student status (domestic compared to international) and year of study (first year students compared to students in second, third and later years of study) had a stronger association with the risk of scam deception. 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 students in second, third or further years of study. 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.

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.

Introduction

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 today (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). As a result, these methods provide cybercriminals with an attractive initial point of contact with individual victims, acting as major vectors for the propagation of cyberattacks.

Although spam may merely involve harmless advertising through unsolicited emails, SMS texts, or social network messages, spam may also contain viruses or malware that are designed to exploit personal or sensitive information from its recipients. While many forms of cybercrime target 'low-volume, high-value' victims and require advanced hacking expertise, spam differs in the preference for 'high-volume, low-value' victims and the relative ease of distribution (Alazab & Broadhurst 2016). Though spam may seem insignificant at the individual level, recent estimates indicated that the average daily spam email volume was approximately 422 billion in January 2018, constituting about 85 percent of all daily global email traffic (Talos 2018). Such mass distribution is possible because spammers are capable of sending tens of thousands of messages in a matter of seconds via botnets, with the recipients potentially vulnerable due to ineffective antivirus software and other countermeasures (Alazab & Broadhurst 2016).

However, while the ubiquitous threat of spam has significant economic and social consequences (Hong et al. 2013; Pattinson et al. 2012), experiences of victimisation and susceptibility are not universal. Previous studies have suggested that factors such as gender (Sun et al. 2016; Sheng et al. 2010), age (Gavett et al. 2017; Oliveira et al. 2017) and technical experience (Pattinson et al. 2012; Sheng et al. 2010) may influence an individual's susceptibility to spam and phishing attempts. Accordingly, the present study was designed to determine how these factors and others—namely the type of scam, cybercrime awareness, IT competence, and perceived internet safety—influenced cybercrime susceptibility among a sample of Australian National University (ANU) students. To accomplish this, research participants were exposed to various fake emails sent from an experimental web server, and their interactions with these scams were examined. Notably, engaging university students offered the advantage of using a single institutional internet service which could be monitored to track real spam events, and reduced key ethical concerns associated with an open or public sample. Our sample of university students offered relatively young and well-educated subjects, who have been noted in previous studies as being more likely to receive spam than others (see De Kimpe et al. 2018), making it suitable to test factors that may increase or decrease the risk of being deceived by well-designed email scams, especially in light of developing cyberattack techniques (Alazab & Broadhurst 2016; Kumaraguru et al. 2010; Sheng et al. 2010). Initial findings of this project did not support a link between phishing susceptibility and the five key variables set out in this paper. Post-hoc analyses, however, found that both international students and first year students were more likely to be deceived by phishing scams than other students.

Literature review

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 (Chaudhry, Chaudhry & Rittenhouse 2016). 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.

Malware may be distributed through two kinds of spam: attachments containing viruses or Trojans that install themselves when recipients click 'download'; and hyperlinks to malicious webpages that appear legitimate (Tran, Alazab & Broadhurst 2013). Common forms of malware include key loggers, screen grabbers, and spyware, which enable phishers to use and exploit sensitive information (Chaudhry, Chaudhry & Rittenhouse 2016). Through these, phishers also gain the capacity to use compromised computers for other phishing attacks, which may target the victim's acquaintances in the form of spam or distributed denial-of-service (DDoS) attacks.

Phishing may also involve the creation of compromised websites, the acquisition of email lists by using botnets, and the spoofing of emails in order to change the apparent sender of messages. Social engineering aspects here serve to deceive the victim into downloading attachments or visiting compromised sites by convincing them that the malicious message comes from a legitimate and trustworthy source.

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 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 (eg a winning lottery ticket), lust, or vanity (eg an adoring admirer or a prestigious job opportunity). De Kimpe et al. (2018) further listed characteristics that can either facilitate or hinder the success of phishing emails (eg the presence of spelling, design or formatting errors, or offers of prize money). 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 (eg returning favours) or 'social proof' (ie others are participating); creating a sense of scarcity; or evoking an authoritative source will help the deception to succeed.

The hook is the compromised link or attachment included in the email, while the catch involves the attacker obtaining and using collected information. Although this may appear simple, techniques and procedures involved in phishing are constantly evolving, and can reflect new social trends occurring in the world (eg Gudkova et al. 2017) or employ new methods of bypassing security protocols and evading detection (Alazab & Broadhurst 2016). The continued rapid spread of the internet in particular has allowed attacks to increase in frequency and diversity, which enhances the likelihood of their success (Chaudhry, Chaudhry & Rittenhouse 2016).

When phishing emails make use of personalised 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 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). Although spear-phishing attackers use far fewer emails than mass generic phishing attackers and limit their attacks to entire companies or institutions, they nonetheless manage to cause sizable economic and social costs (Sun et al. 2016; Wang et al. 2012). It has been suggested that, because they appear more reliable to their targets, spear-phishing emails are likely to increase in prevalence due to the financial gain they offer (Symantec 2014).

Variables associated with phishing risk

Scam type

The success of any phishing or spear-phishing email is inextricably linked to how well it is able to deceive its recipient. While there exists a body of literature that focuses on phishing email structure such as the use of visual cues and the presence of misspellings et cetera (see, for example, Parsons et al. 2015; Wang et al. 2012), this study is primarily concerned with email contextualisation and personalisation. Thus far, research has generally indicated that phishing emails that contain personalised information relevant to recipients are effective in deceiving their targets (Benenson, Gassman & Landwirth 2016; Jagatic et al. 2007; Jakobsson & Rathkiewicz 2006).

One study tested the effects of different social engineering strategies by sending a series of genuine, phishing and spear-phishing emails to a group of 117 university students (Butavicius et al. 2015). Overall, the results indicated that students tended to classify emails as genuine rather than fraudulent, and were worse at detecting spear-phishing attempts than generic phishing attempts. It was also found that, 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 not able detect spear-phishing reliably at all. These findings are supported by Goel, Williams and Dincelli (2017), who sought to directly examine the effect email contextualisation had on victim susceptibility. After sending emails that varied in motive and contextualisation to a large sample of university students, the study found that contextualisation did in fact influence the success of phishing attempts. In particular, emails that addressed a specific student, appeared to be sent from within the university, and threatened the loss of course registrations were found to be the most effective among participants.

Cybercrime awareness

It is thought that individuals who have been made aware of their own potential victimisation become more cautious or defensive as they navigate risky environments. Though this notion poses a practical challenge to researchers in the form of cognitive bias (Pattinson et al. 2012), studying the impact of participant 'priming' has the potential to inform the development of training programs aimed at preventing instances of online victimisation. Certainly, 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), and are more likely to adopt appropriate behavioural countermeasures to phishing attempts (Parsons et al. 2015). In the present study, our subjects were primed because ethical approval (ANU Human Research Ethics Committee protocol no. 2015/038) required offline formal consent and agreement for attempts to deceive them with a scam email.

However, while primed participants may appear to be less susceptible to phishing than non-primed participants, Alsharnouby, Alaca and Chiasson's (2015) study challenges the extent to which priming actually assists in preventing victimisation. When asked to differentiate between legitimate and phishing websites, participants in that study were only able to correctly detect fraudulent websites about 50 percent of the time, reflecting a relatively poor detection ability. This ability was similarly called into question by an experiment conducted by Caputo et al. (2014) that evaluated the effectiveness of an embedded training program designed to assist in phishing identification. Following an initial attack to determine baseline susceptibility rates, users who fell for the phishing email were either given training materials ('training' condition) or were simply notified that they had been spear phished ('awareness' condition). A second email was then sent and, in comparing the click rates of both groups, it was found that the training cohort performed only slightly better than the awareness group (34% to 36%). While it is important to note that the click rates of both groups improved from the baseline rates by almost 50 percent, the results also indicate that the exact impacts of priming and non-priming on susceptibility are not yet fully understood.

Gender

Gender is often included within phishing vulnerability studies as a demographic variable; however, results regarding its impact on phishing detection are mixed. While some studies have concluded that a statistically significant relationship between the two does exist (Halevi, Memon & Nov 2015; Iuga, Nurse & Erola 2016; Jagatic et al. 2007; Sheng et al. 2010), others have found no such connection (Butavicius et al. 2017; Flores et al. 2015; Mohebzada et al. 2012; Oliveira et al. 2017). Where a correlation has been discovered, females are shown to be more susceptible to phishing attempts than males, and it has been suggested that this is due, in part, to less technical training and computer knowledge among females (Sheng et al. 2010).

Despite these contradictory findings, more recent studies have begun to reconceptualise the relationship between gender and phishing susceptibility in order to better understand it. In Goel, Williams and Dincelli's study (2017) for example, the act of falling for a phishing scam is framed as two steps: first, the opening of a phishing email and, second, the clicking of the malicious link within. The study found that, while women were more likely than men to open risky email messages, they were also less likely to click on embedded links in an email. The differences in click rates between males and females was not statistically significant. Susceptibility is thus nuanced by this two-stage process, which has implications for the development of targeted intervention programs aimed at specific populations and behaviours.

IT competence

It is commonly suggested that technical knowledge and experience develop and improve an individual's online security safeguards (Sun et al. 2016). However, the extent to which an individual's IT competence affects their phishing susceptibility is still not well understood. In their scenario-based role-play experiment, Iuga, Nurse and Erola (2016) examined the relationship between PC usage 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 possessed better detection scores, suggesting that extended interactions with computers increased an individual's phishing awareness and improved their detection strategies. Observing the related concept of technical proficiency, Alsharnouby, Alaca and Chiasson (2015) employed a similar methodology in their experiment involving a group of 21 participants. Their measure of technical proficiency related to the range of online activities that users regularly engaged in (eg designing a website, installing an operating system etc). This was used as a measure of general proficiency rather than as a measure of technical specialisation. Though no significant correlation between proficiency and improved detection scores was found, it was acknowledged that using a different measure (ie assessment of technical proficiency) may have led to different results.

In operationalising computer familiarity, Pattinson et al. (2012) essentially combined the concepts of usage and proficiency by asking 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 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 may need to be regularly informed about phishing risks and actively conscious of phishing in order for their familiarity with computers to reduce their vulnerability to more subtle forms of phishing. These studies perhaps point towards a difficulty in conceptualising IT competence, and towards a need for standardised and uniform definitions.

Perceived internet safety

In both online and offline settings, perceptions of safety alter individual behaviour and determine safety precautions, which aim to mitigate risks prospectively. These perceptions are formed at the intersection of various factors, which can range from past experiences of victimisation to an individual's unique personality traits. The literature has broadly examined the influence of perceptions of internet safety on phishing vulnerability (eg Abbasi, Zahedi & Chen 2016; Parrish, Bailey & Courtney 2009); however, the relationship between 'feeling safe on the internet' and the actual risk of deception via a 'phish' has not yet been quantified.

In drawing from existing research, it is difficult to determine if those who view the internet as unsafe are actually less susceptible to phishing attempts. For instance, while it may be thought that those who feel unsafe are more likely to be vigilant and to employ protective countermeasures, Abbasi, Zahedi and Chen's (2016) study reported that this may not necessarily translate into decreased vulnerability. Their research categorised individuals into clusters or profiles based on shared online experiences, and analysed their interactions with fake phishing pages. It was found that the best detectors (clusters 1 & 2) were those who were keenly aware of phishing, were highly familiar with websites, had positive perceptions of anti-phishing tools, and had experienced past losses to phishing. At the same time, however, some of these traits also negatively affected an individual's ability to successfully detect phishing attempts, as shown with clusters 3 and 6. With the former group, it was suggested that past encounters and phishing awareness caused individuals to be overconfident in their ability to detect malicious websites; with the latter group, that familiarity with frequented websites induced over-reliance and trust. Certainly, although the study makes no mention of how these traits form perceptions of safety, results such as these indicate that more robust notions of internet risk and vulnerability may assist in phishing avoidance.

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 to what extent the factors set out above influence the risks of cybercrime for students at the ANU. To accomplish this goal, participants were exposed to various fake email scams, and their interactions with these scams were observed. The experiment was conducted over a period of nine months from February to November 2017. During this time, email content was socially engineered to replicate three different types of phishing: generic, tailored, and spear-phishing. This meant the respective emails had to be either: broad and impersonal, tailored to participants' institution of study, or highly specific to a participant's 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 (ie 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 scams was defined as the act of clicking on the fake links embedded 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 become increasingly tailored to the individual. Participants were expected to be more likely to be deceived by spear-phishing emails then tailored emails, and more likely to be deceived by tailored emails 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 condition was a reinforcement rather than a primer or awareness stimulus.)
- 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 ANU ethics protocol. Those who completed the post-observational survey received a free hamburger voucher from a popular store, offered 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 information about attitudes to the internet were obtained from participants via a pre-test and the 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 feeling 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 in total.

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

Upon finishing the experimental phase, participants were asked to complete a second survey. This involved responding to the same questions that were asked in 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 on their perceived risk and awareness of cybercrime, and how participating in the study had influenced their internet-related behaviours. 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 behaviours.

Software, materials and data recording

This experiment required a redesign 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 interactions participants had with the fake phishing emails. We also needed a service to host a number of different websites (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:

- the use of the university's mail server to spoof originating email addresses; and
- a record of when emails were sent and if and when they were opened, if and when a participant clicked on malicious hyperlinks, and if and when they then entered their credentials.

The research team set up a web server that ran an industry-standard web server configuration 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 email responses was also developed. The emails were crafted to appear to have been sent by a (fake) person or organisation. Access to an open SMTP server was required for distributing the spoof emails. The SMTP standard does not require the originating email address to be correct, and can therefore be exploited to send emails that appear to be from another person or organisation.

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, who was assigned a unique identifier. Sent emails were also attributed unique identifiers, which allowed 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 ambiguity);
- received but ignored—the participant opened the email but chose not to take any action. This may or may not have been because they identified the email as fraudulent; and
- received and responded—the participant took action in response. This could have been 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's internet service provider to create a whitelist to track the students participating in the study, and so alternative means of monitoring had to be devised. A number of ANU web services login screens were duplicated, including those for the email system and online student management services. They were designed to record the time, date and IP address of each access by one of our participants, as well as whether or not the participant then proceeded to log into the fake website. We also embedded a hidden web beacon into the content of the email to track when a recipient had opened an email.

The presence of the web beacon would have triggered the spam filters used by many email providers. This made it difficult to track the rate at which emails were opened. It is also possible that the web beacon sometimes 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 not have been able to signal to the server that an email had been read. It was therefore not possible to track all emails which had been received but ignored, as we could not determine whether the email was actually read by the participant. Thus, the absence of an access record is not conclusive evidence that a participant did not open one of the spoof emails. The data counts are therefore conservative.

Operationalising scam susceptibility

The number of fake scams that successfully deceived recipients operationalised scam susceptibility. There were nine fake emails that tested participants' susceptibility, with scam content varying across three levels of specificity or individualisation (see Figure 1):

- 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 mails 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 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 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 (as per ANU ethics protocol no. 2015/038). Participants completed a 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 before the study began. Participants were randomly assigned to one of two conditions (hunter or passive). Hunters were reminded every four to six weeks via email to remain constantly vigilant for both fake and real forms of cybercrime, and to forward all suspicious content to the researchers (see *Appendix E*). A reminder email detailing these instructions was sent to hunters every two weeks following the initial email (see *Appendix F*).

To create content for the spear-phishing emails, personal information about each participant was extracted, if possible, from their Facebook and LinkedIn profiles. 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 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 (see Figure 1). 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 created only for the 25 participants with adequate online information.

During the observation phase, participants received between seven and nine fake emails, 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 compromised 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 (see *Appendix G*).

All fake emails (see *Appendix H*) were delivered from the experiment's server hosted at the ANU Cybercrime Observatory. Any websites that participants viewed or interacted with were hosted on this server. The server recorded all participant interactions with these emails, including whether emails were opened, whether links were clicked, and whether participants attempted to log in to fake pages. Upon completing the experimental phase, participants' scam counts were recorded.

At the end of the experimental period, participants were asked to complete the internet survey (see *Appendix I*). This survey was completed in the participants' own time via the Qualtrics online survey service. Upon completion of the survey, participants were provided with a hamburger voucher (offered as an incentive to do the survey) and an opportunity to write feedback. For clarity, a complete timeline of the experiment can be found in *Appendix K*.

Figure 1: Examples of scam emails at each level of specificity

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-administration/assessments-exams/examination-conduct and http://www.anu.edu.au/students/program-administration/assessments-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]

Figure 1: Examples of scam emails at each level of specificity

Spear-phishing 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 Sportsperson 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

Results

Data cleaning and screening

Screening, cleaning, and subsequent analyses were performed using Statistical Package for the Social Sciences (SPSS v24) software. The dataset was examined for out-of-range and missing values, and adherence to assumptions of chi-square and analysis of variance (ANOVA). All data were within range. Some missing values were found among the questions about student and residential status but this did not impact on results.

Visual inspection of histograms revealed approximately normal distributions for all relevant variables, and skewness and kurtosis statistics revealed no significant violations of normality. A non-significant Levene's test indicated that the homogeneity of variance (HOV) assumption had not been violated for any relevant demographic variables. Low cell counts were discovered in cross-tabulations between many variables, indicating a violation of the sample-size assumption of the chi-square test. This was resolved by using Fisher's exact test in subsequent analyses.

Figure 2: Participants deceived, by gender, ordered by scam counts (*n*)

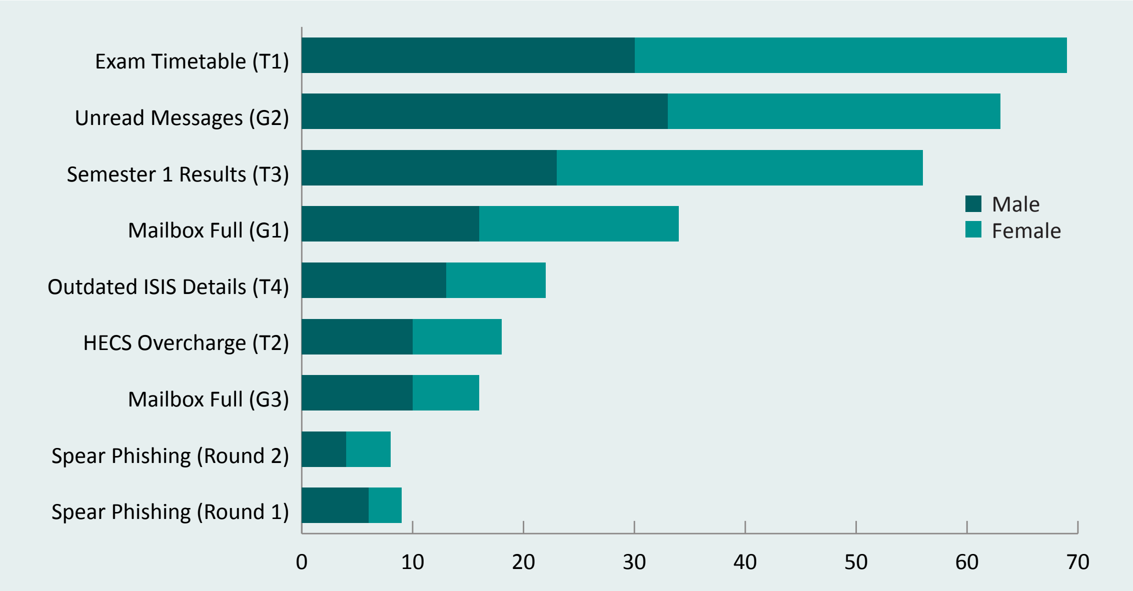

Note: The sample size for spear-phishing attempts is 25 participants and for all other scams is 135 participants (after removing 'other' gender and missing values). Following the scam type, we indicate the level of specificity by G=generic, T=tailored as distinct from spear-phishing. We note the order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

Overall 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 were obtained from open sources such as Facebook. The total numbers of scams are 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 tailored to the participants' university study. Participants were almost equally susceptible to a generic scam titled 'Messages'. Participants did not, however, differ significantly in terms of gender according to Fisher's exact test. Figure 2 shows the number of participants who fell for each scam by gender. The number of successful scams by gender is shown in Figure 3.

Figure 3: Participants deceived by no scam, one scam, or more, by gender (*n*)

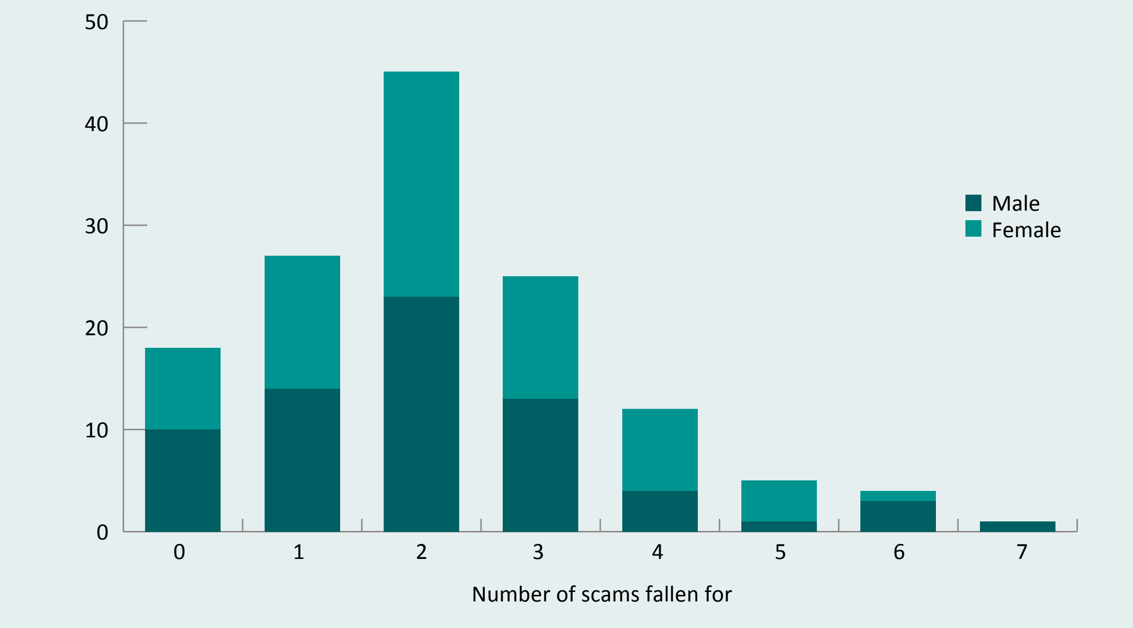

Overall, there appeared to be an increasing trend in relation to the scam type and scam susceptibility in the normalised proportions as shown in Table 2, with increasing success for more individualised and tailored scams. However, while a Wilcoxon signed-rank test showed that these proportions do not differ significantly, the comparison between generic and tailored scams 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. The low numbers (n=25) for the spear-phishing sample significantly reduce the power of the Wilcoxon signed-rank test when paired with the corresponding generic and tailored samples.

Table 2: Successful deceptions by scam type

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

Figure 4: Mean scam counts, passive and hunter groups

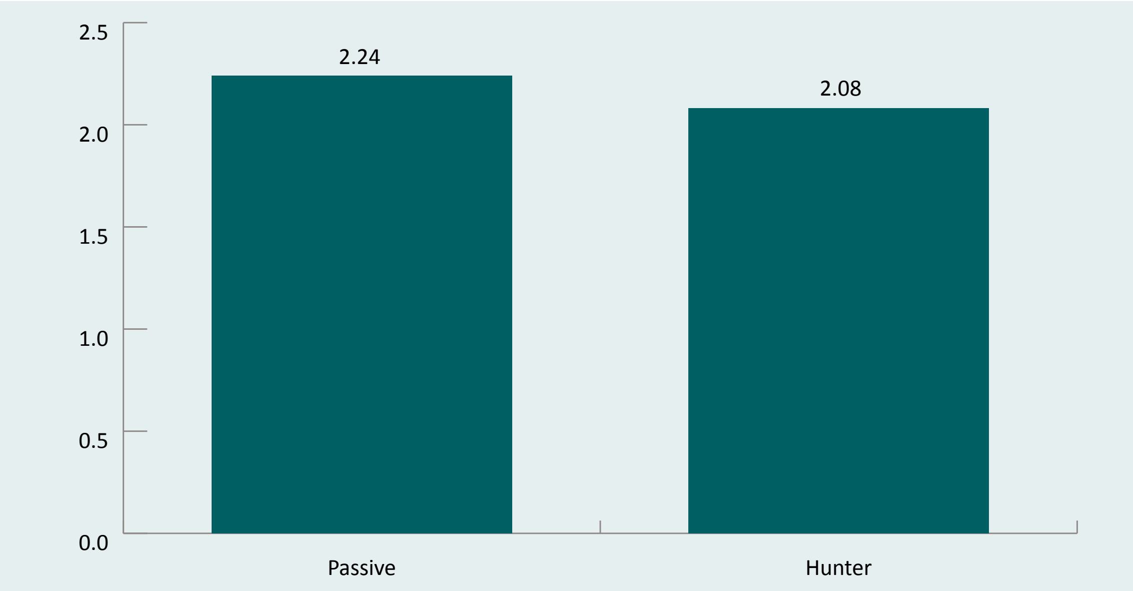

Bivariate tests: Cybercrime awareness, IT competence, and perceived internet safety

In order to assess whether scam susceptibility varied as a function of an individual's cybercrime awareness, a one-way ANOVA was used, with condition (passive or hunter) as the independent variable (IV) and scam count as the dependent variable (DV). The hypothesis was not supported, although on average those assigned to the 'hunter' condition were slightly less likely to fall for a scam than those assigned to the 'passive' condition. The ANOVA addressing this question revealed a non-significant difference in scam susceptibility between the passive group and hunter group, $F(1, 136)=0.543$, $p=0.462$. This indicates that there was no reliable difference in scam susceptibility due to participants' primed awareness of cybercrime throughout the study. Figure 4 presents scam susceptibility in relation to passive and hunter groups. Results indicate that scam susceptibility did not vary as a function of cybercrime awareness, demonstrated by the non-significant difference between groups.

Multivariate modelling

To test the combined effects of the variables of interest listed in hypotheses 2, 3, 4 and 5, we also 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; see *Appendix J*) were the initial hypothesis variables: gender, IT competence, cybercrime awareness (hunter or 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 (see Table 3).

Table 3: Analysis of deviance table of Model 1

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

Post-hoc analyses: Student status and years of study

Initial bivariate correlational analyses revealed significant relationships between scam susceptibility and student status (*R*=-0.287, *p*<0.01), and scam susceptibility and years of study (*R*=-0.297, *p*<0.01). Accordingly, it was explored (post-hoc) whether scam susceptibility might reliably differ as a function of student status (domestic vs international students) and/or years of study (first year vs later year students). Two additional ANOVAs were performed to examine this possibility. The mean scam count was found to differ significantly between domestic and international students (*F*1, 134=12.01, *p*<0.01). As can be seen in Figure 5, a greater number of international students fell for three or more scams, indicating significantly higher scam susceptibility among international compared to domestic students. Similarly, mean scam count differed significantly between first year and later year students (*F*1, 136=26.1, *p*<0.001). Figure 6 demonstrates a higher number of first year students falling for three or more scams, indicating a significant difference in scam susceptibility among first year compared to later year students.

Figure 5: Participants deceived by no scam, one scam, or more, by student status (%)

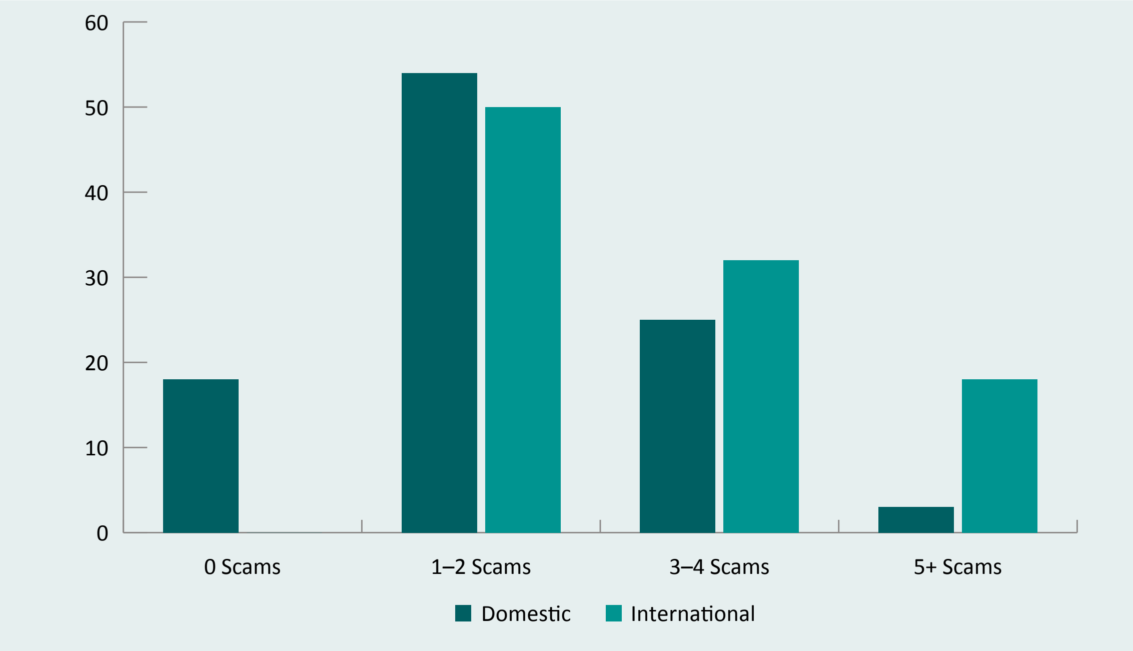

Figure 6: Participants deceived by no scam, one scam, or more, by years of study (%)

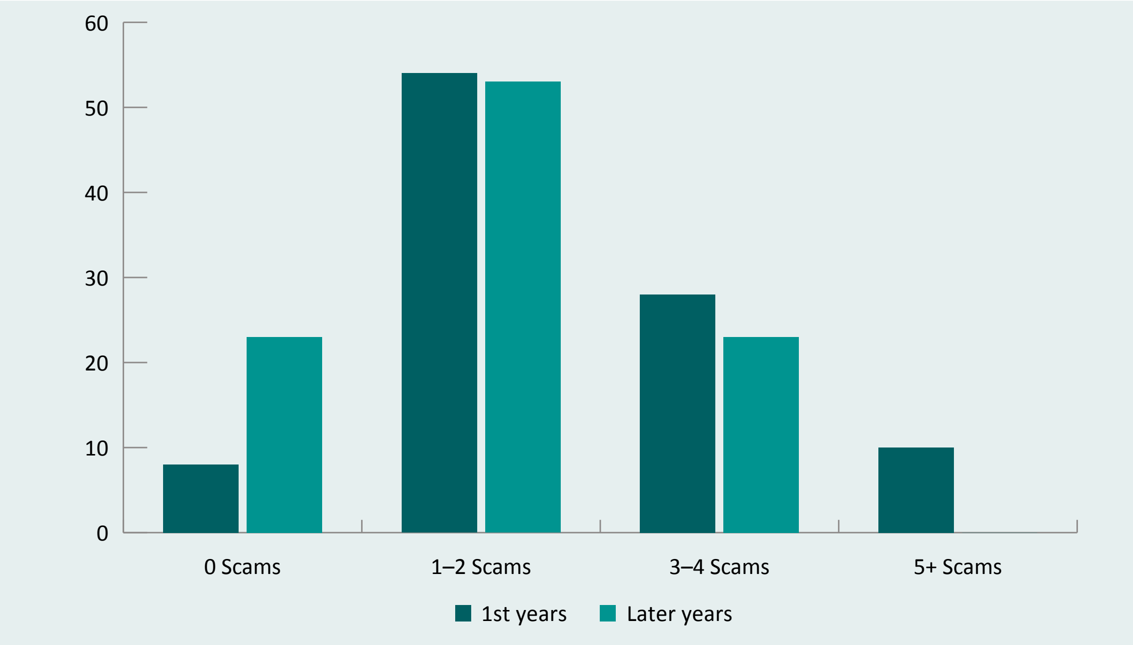

In a further analysis of these variables, it was revealed that the best model for susceptibility is obtained from a stepwise variable selection procedure, which includes only ‘years of study’ (first year or later year university student) and ‘student status’ (international or domestic student) with no significant interaction effect. We call this Model 2 (*Appendix J*). As seen in Table 4, adding in the hypothesis variables included in Model 1 to Model 2 produced no significant change (p>0.35). In Table 5, however, it is shown that both variables are significant with p value 0.012 for ‘years of study’ and p value 0.017 for ‘student status’, respectively.

Table 4: Analysis of deviance table of Model 2 and the Full Model (Models 1 and 2)

Model	Residual *df*	Residual deviance	Additional *df*	Change in defiance	*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

Table 5: 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

Note: Both variables in Model 2 are significant at 0.05 level

Analyses of gender against other variables

Self-reported IT competence was found to significantly differ by gender (Fisher's exact test $p<0.001$, ϕp=0.38). As shown in Figure 7, 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), as seen in Figure 8.

Figure 7: Self-reported IT competence by gender (%)

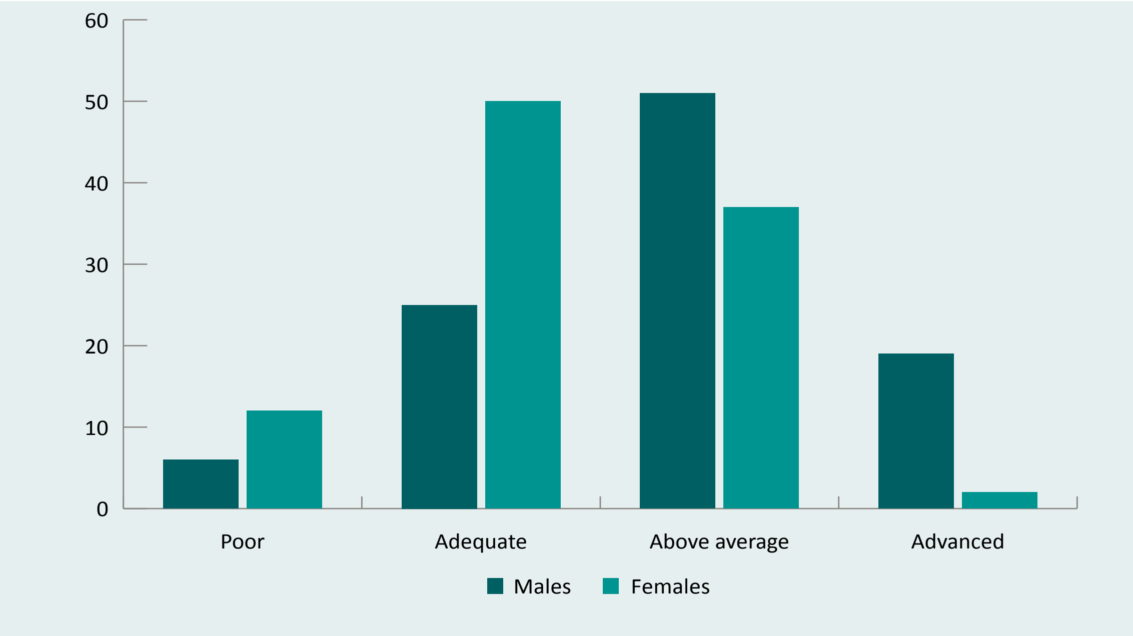

Figure 8: Self-reported ability to spot scams by gender (%)

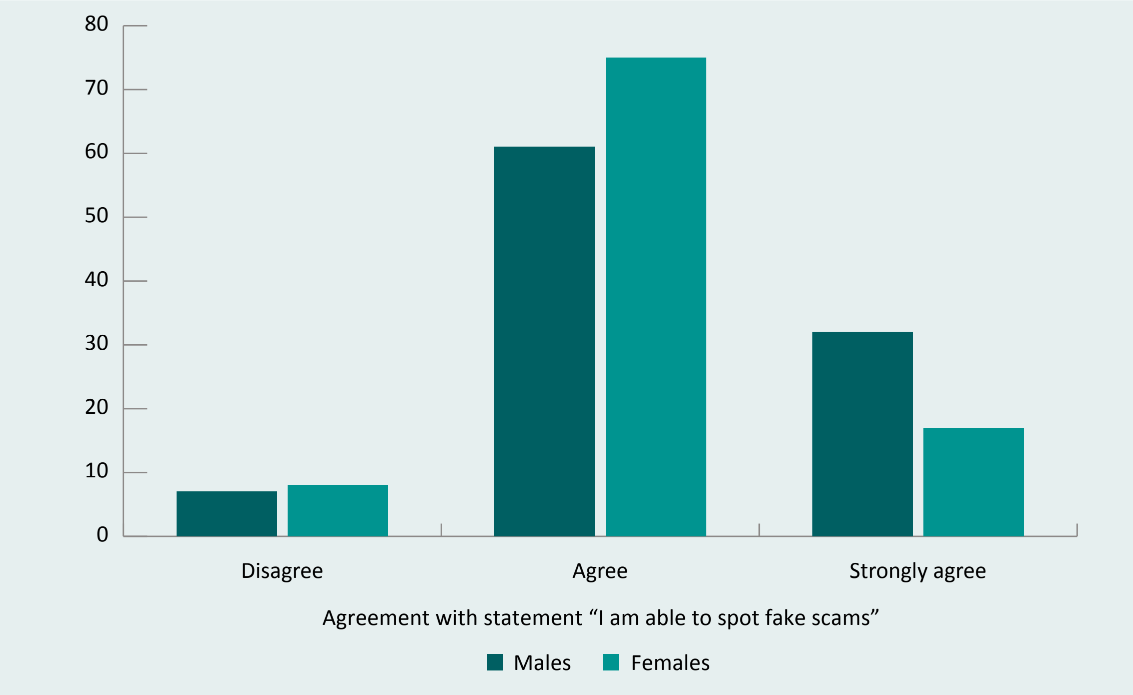

Pre-study and post-study differences

Responses to the internet survey at time 1 (before the observations) and time 2 are reported in Table 6. However, only 62 percent of the respondents completed the follow-up survey, limiting the reliability of pre- and post-study differences. Importantly, participants rated their IT competence as higher post-study, but their perceptions of online safety remained largely unchanged. Additionally, there was more diversity in the types of social media used by respondents at time 1 compared with time 2.

Table 6: 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
Facebook	94.7	96.5
Instagram	51.9	69.4
Snapchat	46.6	76.5
Google+	7.6	18.8
Twitter	6.9	30.6
Tumblr	6.1	18.8
LinkedIn	4.6	23.5
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		
Disagree	6.5	4.7
Agree	65.2	71.8
Strongly agree	24.6	23.5
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 participating in the study (95% CI [-0.33, -0.05])

Discussion

This study was designed to determine the risks of cybercrime for students at the Australian National University. At the heart of this study was an interest in how scam type, level of cybercrime competence and awareness, and campus demographics influenced susceptibility to cybercrime. A review of relevant literature suggests that cybercrime susceptibility may be influenced by the level of specificity in a scam. That is, individuals may be more likely to be deceived by scams that are tailored to their personal circumstances (spear-phishing) compared to scams with generic content (phishing). Certain other variables have also been flagged as potential contributors to scam susceptibility. In the literature these variables include level of cybercrime awareness, IT competence, and gender. Accordingly, this study proposed that scam susceptibility in ANU students may be associated with these variables. To explore this possibility, participants were exposed to social engineering directives in the form of fake email attacks. These fake emails attempted to either elicit personal information from participants or compel them to click links that in the real world could contain malware.

To determine whether participants were more susceptible to spear-phishing attacks than generic attacks, email content was socially engineered to replicate three different scam types: generic, tailored, and spear-phishing. These scam types differed in their level of personal relevance (specificity) to the individual. Results revealed no significant relationship between scam type and scam susceptibility. That is, participants in general were not found to be more susceptible to spear-phishing attacks compared to generic and tailored attacks. However, the email content that deceived most participants also provides insight into the types of scams that may succeed. The most successful attack related to participants' final exam timetable. This tailored email was of an urgent nature and was sent during the ANU exam period. It likely succeeded because it was both relevant and salient, and potentially instilled a sense of fear in participants. Unnoticed changes to the exam timetables could have fostered concerns in participants, compelling them to click on and respond to the malicious links and prompts. Thus it may be said that the success of this fake scam can be attributed to a combination of personal relevance and fear, indicating that individuals in the real world may be more susceptible to scams which tap into salient and urgent life circumstances.

The hypothesis that scam susceptibility would vary as a function of cybercrime awareness was not supported. Despite participants in the hunter condition being primed regularly to remain vigilant for cybercrime, this did not correspond to a reduction in scam susceptibility. The most likely explanation for this is that the hunter condition did not adequately increase cybercrime awareness in participants. Over nine months, hunters received a total of four emails reminding them about the dangers of cybercrime and prompting them to remain vigilant. This kind of general prompt may have been too weak to raise cybercrime awareness, and thus only minimal differences may have been generated between the hunter and passive conditions. For this reason, the relationship between cybercrime awareness and scam susceptibility remains unexplored despite its theoretical relevance as outlined in previous research (Alsharnouby, Alaca & Chiasson 2015; Parsons et al. 2015; Pattinson et al. 2012). The ineffectual prompting apparent in the present study suggests that increasing the public's level of cybercrime awareness requires constant effort and specific rather than general warnings about cybercrime.

The gender, IT competence, and perceived internet safety hypotheses were also not supported. In line with more recent studies (Alsharnouby, Alaca & Chiasson 2015; Butavicius et al. 2017; Goel, Williams & Dincelli 2017; Oliveira 2017; Pattinson et al. 2012), 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 safe versus unsafe. While it is possible that these relationships may not exist, it is more likely that the present study was perhaps too small and/or atypical to detect the significant relationships that have been identified in other studies (eg Halevi, Memon & Nov 2015; Iuga, Nurse & Erola 2016; Sheng et al. 2010).

Student status and years of study

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 exact nature of this relationship is unknown, it can be theorised that international students were possibly disadvantaged by a language barrier, or had different experiences with cybercrime in their countries of origin. This latter point is in line with a few studies (eg Butavicius et al. 2017; Flores et al. 2015) that have proposed that national individualism plays a role in phishing vulnerability due to the interplay between individual protection and interpersonal harmony. These studies suggest that individuals from countries low in individualism may be more inclined to comply with scam requests, as these nations promote the needs of the wider group and encourage the development of large social networks (Butavicius et al. 2017; Hofstede 1984). The current project, however, is unable to determine the impact of this notion within the specific ANU context, as the relevant demographic data were not collected to begin with (see *Appendix A*). Nevertheless, the present study does highlight that international students are easier targets for scammers, indicating that the university would benefit from providing these students with targeted educational resources to improve overall campus cybersecurity. As a research note, further exploration of the factors which contribute to this significant relationship would assist in explaining how phishing varies in an inter-state context.

Similarly, a post-hoc analysis revealed that 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 and overall confidence. Perhaps later year students had been exposed to more real-world scams and had thus developed more of an ability to detect scams through those experiences. Alternatively, later year students may have been more confident in navigating the ANU email systems compared to first year students. As with international students, the present study identifies first year students as being more at risk of cybercrime, and indicates that awareness measures should be targeted to new rather than continuing ANU students. On a broader scale, exploring the influence of age and experience on scam susceptibility would allow the exact nature of this relationship to be understood in greater depth. A lack of age variability in the present study prevented this from being explored.

Conclusion

It is important to acknowledge the limitations of this small exploratory study. Firstly, as mentioned previously, the experimental manipulation (hunter vs passive) may not have adequately distinguished participants with different levels of cybercrime awareness. This prevented sufficient observation of the relationship between cybercrime awareness and scam susceptibility. This relationship makes theoretical sense, and so it is important that future research explore how cybercrime awareness affects susceptibility to scams. This could be done in a manner similar to the approach taken in the present study, but with stronger and more consistent prompting.

Secondly, the ability to observe whether emails were opened was not technically possible until halfway through the study. It was therefore 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 substantially different from ignoring the content entirely; however, the study's temporary inability to distinguish between these two actions meant that participants who ignored the emails altogether were being treated in the same way as those who identified the scams—that is, as less susceptible to fake scams. Given this, our interpretation of non-response is conditional, because it was not always possible to distinguish whether an email was unread or unopened. Our observation is thus limited to what action our respondents took, if any, with respect to the response demand of the phish. For this reason, it is important that future studies of similar design are able to determine whether an email was either opened and interacted with or completely ignored.

Finally, the present study did not account for practice effects. It is possible that the influence of a scam type was overshadowed by the practice participants received through repeated exposure to fake scams. Of note, for example, is that our generic 'Mailbox Full' scams deceived half as many participants (*n*=16) in the second round as in the first round (*n*=34), suggesting that a practice effect may be in play. While results did not reveal an overall decrease in susceptibility over time, it would have been beneficial to distinguish between the effects of scam types and the effects of practice. This would have allowed results to be attributed solely and 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 have provided useful 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 victimisation in a real-world setting may be scaled up with larger samples and a greater variety of social engineering methods. Such study is necessary, ultimately, against the ever-increasing magnitude and impact of phishing worldwide.

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: 169–174. 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): 1158–1172

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): 247–256

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): 1277–1287. http://hdl.handle.net/1854/LU-8554543

Flores WR, Holm H, Nohlberg M & Ekstedt M 2015. Investigating personal determinants of phishing and the effect of national culture. *Information & Computer Security* 23(2): 178–199

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: 345–358

Gudkova D, Vergelis M, Shcherbakova T & Demidova N 2017. *Spam and Phishing in Q3 2017*. https://securelist.com/spam-and-phishing-in-q3-2017/82901/

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

Hofstede G 1984. *Culture's consequences: International differences in work-related values*. Newbury Park, California: SAGE Publications

Hong KW, Kelley CM, Tembe R & Mayhorn C 2013. *Keeping up with the Joneses: Assessing phishing susceptibility in an email task*. Proceedings of the Human Factors and Ergonomics Society Annual Meeting. sl: SAGE Publications: 57(1): 1012–1016. https://doi.org/10.1177%2F1541931213571226

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

Jagatic T, Johnson N, Jakobsson M & Menczer F 2007. Social phishing. *Communications of the ACM* 50(10): 94–100

Jakobsson M & Ratkiewicz J 2006. *Designing ethical phishing experiments: A study of (ROT13) rOnl query features*. Proceedings of the 15th International Conference on the World Wide Web. New York: ACM: 513–522

Kumaraguru P, Sheng S, Acquisti A, Cranor LF & Hong J 2010. Teaching Johnny not to fall for phish. *ACM Transactions on Internet Technology* 10(2): 7:1–7:31

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

Mohebzada JG, El Zarka A, Bhojani AH & Darwish A 2012. *Phishing in a university community: Two large scale phishing experiments*. 2012 International Conference on Innovations in Information Technology. Abu Dhabi: IIT: 249–254

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

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

Parrish JL, Bailey JL & Courtney JF 2009. *A personality model for determining susceptibility to phishing attacks*. Oklahoma City: Southwest Decision Sciences Institute

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

Sheng S, Holbrook M, Kumaraguru P, Cranor LF & Downs J 2010. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In Mynatt E, Fitzpatrick G, Hudson S, Edwards K & Rodden T, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York: ACM Press: 373–382

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: 249–257

Symantec 2014. *Internet security threat report 2014*. Mountain View, California: Symantec Corporation. http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf

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

Tran, K. N., Alazab, M., & Broadhurst, R. (2013). Towards a feature rich model for predicting spam emails containing malicious attachments and urls. Paper to the 11th Australasian Data Mining Conference (AusDM), Canberra, 13-15 November

Wang J, Herath T, Chen R, Vishwanath A & Rao HR 2012. Phishing susceptibility: An investigation into the processing of a targeted spear-phishing email. *IEEE Transactions on Professional Communication* 55(4): 345–362

Appendix A: General demographic and internet safety questionnaire

Q.1 Gender

□ Male □ Female □ Other

Q.2 Age

□ Under 21 □ 21-25 □ 26-30 □ Over 30

Q.3 Student status

□ International Student □ Domestic Student

Q.4 Residential Status

□ University accommodation □ At home □ Other

Q.5 Course/Faculty

□ Science □ Engineering/Computing □ Commerce/Economics □ Law

□ Arts/Social Science □ Asia Pacific Studies □ Medicine □ Other □ Staff

Q.6 Year of University Course/Study

□ First □ Second □ Third □ Fourth □ More than four years

Q.7 How do you rate your overall IT (Information Technology) competence?

□ Poor □ Adequate □ Above average □ Advanced

Q.8 Do you access social media daily?

□ Yes □ No

If yes, which of the following do you access the most?

□ Facebook □ Instagram □ Twitter □ LinkedIn □ Google+ □ Tumblr
□ Snap Chat □ No social media account
□ Others, please specify ______________________________

Q.9 Have you been a victim of a cybercrime?

□ Yes □ No

If yes, please briefly describe ______________________________

Q.10 I am usually able to spot a fake email, spam or internet scam.

□ Strongly disagree □ Disagree □ Agree □ Strongly agree

Q.11 I use the internet to purchase goods and services.

□ Never □ Rarely □ Sometimes □ Frequently

Q.12 How safe do you feel when you access the internet?

□ Very safe □ Somewhat safe □ Somewhat unsafe □ Very unsafe

Appendix B: Internet survey

Q.1 Please provide your university ID: ______________________________

Q.2 As a participant in the ANU Cybercrime Observatory spear-phishing study, do you recall falling for any fake email scams?

□ Yes □ No

Q.3 Has this study had an impact on your cybercrime awareness?

□ Yes, I have become more aware of cybercrime □ No change

Q.4 Do you think you are more at risk of cybercrime now compared to the beginning of the study?

□ I am more at risk now □ No change □ I am less at risk now

Q.5 Do you think your online behaviours are less risky, about the same, or more risky now as they were at the beginning of the study?

□ Less risky □ About the same □ More risky

Q.6 In the past 6 months have you changed your privacy settings on any social media accounts that you use?

□ Yes □ No

Q.7 How have your privacy settings changed in the past 6 months?

□ They have become less privat □ They have become more private

Q.8 How do you rate your overall IT (Information Technology) competence?

□ Poor □ Adequate □ Above average □ Advanced

Q.9 Which of the following social media sites do you access (please mark all that apply)?

□ Facebook □ Instagram □ Twitter □ LinkedIn □ Google+ □ Tumblr

□ Snap Chat □ No social media account
□ Others, please specify ________________

Q.10 Do you access social media daily?

□ Yes □ No

Q.11 Which social media site/s do you access the most? ________________________________

Q.12 Have you been a victim of real cybercrime (i.e. not our fake attacks) since signing up to our study?

□ Yes □ No □ Unsure

If yes, please briefly describe your experience as a victim of cybercrime: ________________

____

Below are a few statements and questions relating to the internet. Please respond with the answer that is most applicable to you.

Q.10 I am usually able to spot a fake email, spam or internet scam.

□ Strongly disagree □ Disagree □ Agree □ Strongly agree

Q.11 I use the internet to purchase goods and services.

□ Never □ Rarely □ Sometimes □ Frequently

Q.12 How safe do you feel when you access the internet?

□ Very safe □ Somewhat safe □ Somewhat unsafe □ Very unsafe

Q.13 Would you like to provide us with any comments, concerns, or feedback about the study? (Please write N/A if no)

____

____

____

____

Appendix C: Information sheet

Researcher: My name is Roderic Broadhurst, and I am a professor of criminology in the Research School of Social Sciences (RSSS) in the College of Arts and Social Sciences (CASS) at the Australian National University (ANU). I am an experienced criminological researcher who leads the ANU Cybercrime Observatory and has published widely on cybercrime, crime victims, violence and recidivism. I have held several Australian Research Council, Australian Criminology Research Grants and overseas grants on a variety of topics including cybercrime and I have been a consultant to the United Nations Office of Drugs and Crime and Australian and overseas police agencies. I will be assisted by research assistants and interns working for the ANU Cybercrime Observatory: Don Maxim, Katie Skinner, Nick Sifiniotis, Hannah Woodford Smith, Bianca Sabol, Charlotte Chung and Virginia Chow.

Project Title: Cybercrime Risks in University Student Communities

General Outline of the Project

Description and Methodology: This research study has two main aims: (1) to determine the risks of cybercrime for staff and students at the ANU, and; 2) to determine the coverage and usage of Wi-Fi across the ANU campus. Security and Wi-Fi usage across the ANU campus have been ongoing concerns at ANU and not studied from a social science or victim's perspective. In particular, this study is interested in how risks and awareness of cybercrime varies across the campus. As a participant, we will ask you for permission to collect the URL links that you visit. These URL links allow us to estimate the distribution of web usage and determine current cybercrime risks (e.g. access to known blacklisted scam websites). We will also conduct social engineering experiments (in the context of cybercrime) to determine if some students are more vulnerable than others to phishing attacks and other on-line deceptions (e.g. fake-scams designed to elicit personal information – that is steal details of your identity). Participants will get fake email scams sent to them over coming weeks which will attempt to entice them to reveal private information.

Participants: The target participant group are ANU students/staff, and we particularly want to recruit new and/or first year students to participate. We look to gauge the internet access of students and their susceptibility to cybercrime and seek volunteers to help us understand on-line risks. To participate, please sign up in "0" week or class or attend one our research briefings and/or visit our website (http://sociology.cass.anu.edu.au/centres/anu-cybercrime) to indicate your interest in participating in this study. You can indicate your interest on-line and we will follow-up with you to explain the research and possible risks involved.

Use of Data and Feedback: Your web browsing data (we collect only your URL data not the page visited or the content) and information collected by our social engineering spam techniques will be used to improve cybersecurity at the ANU. The study will help determine the most vulnerable groups, provide recommendations to improve cybersecurity practices, develop better cyber safety awareness, and improve the quality of cybersecurity research. We will provide updates about our research through our website (*http://sociology.cass.anu.edu.au/centres/anucybercrime*), the ANU Cybercrime Observatory website, and through your ANU email address (required for sign up). We also notify you if we detect any risks of cybercrime associated with this study.

Participant Involvement

Voluntary Participation & Withdrawal: Your participation in this study is voluntary and you may withdraw from this research study at any time without any penalty. When you withdraw, we will ask you explicitly if we can continue to store and use your data. We will remove your data if we do not have your permission or in the absence of explicit permission from you to retain your data.

What does participation in the research request of you? Your participation in this research requires access to your URL browsing data on the ANU network. You will also be exposed to social engineering directives (in the form of spam) that attempt to get you to reveal your personal information. These directives or scam spams will be revealed if you are deceived by them. You will then be notified that you have fallen for a (fake) scam and we will provide you with cyber safety materials to avoid becoming a victim to (future) true scams. You will also be asked to complete short questionnaires or surveys throughout your participation. Note you will remain at risk of 'real' scams that can sometimes avoid the spam traps and filters designed to prevent such emails reaching your mail service. We will alert you to these once known but you should be aware that not all scam mails are identified and when identified may not be identified swiftly and a time gap between receiving these deceptions and warnings about them are common.

Location and Duration: The research study, its investigators, and data collection are all located and conducted on the ANU campus. The research study period is planned to commence from 15 FEBRUARY 2017 and end 18 DECEMBER 2017. Your data will be collected between 20 FEBRUARY 2017 and 31 JULY 2017. You will be notified of additional data collection periods. Data collection of URLs will be undertaken in two 24/7 periods and we will remind you when this takes place by an on-line banner that data collection will occur (to the effect "your URL visits will be collected for research").

Incentives: No incentives to participate are offered. However, we may offer a deceptive advantage (a fake) that is common in cases of a real malware/cybercrime events. For example the fake mail might ask you to login into a website to get a 'free gift'.

Risks: The main risk factor is the disclosure of your personal information to true cybercriminals. This risk is not increased by your participation in this study. When you sign up, we will ask you to remain vigilant and be careful about cybercrimes and email scams in particular. We will also notify you if we find that you accessed known malicious URLs that you may not be aware of doing. Another risk is that you may visit web pages that you prefer not to disclose to anyone. Although we will not be "reading" URLs we will alert you to this by reminding you of your participation every time you access a browser during the data collection periods. Participants also remain subject to the terms and conditions associated with the ANU Secure and re-sign these as condition of participation – these terms and condition require responsible use of ANU Secure.

Implications of Participation: Your participation will allow us to estimate the state of Wi-Fi usage, the landscape of websites accessed across the ANU campus, and the level of cybercrime risks/awareness at the ANU. We will use the results in our study to provide recommendations to improve internet access on the ANU campus, and to improve cybersecurity and cyber-safety at the ANU. Direct benefits to you may be an increased awareness and vigilance about social engineering methods (and other cybercrimes) that are intended to defraud or steal your personal information or other malicious intentions.

Confidentiality

Your data is stored on a RAID-1 array of disks that is encrypted using AES-256 key encryption (one of the current most secure encryption methods). This study will only store information you have provided to us and the websites that you have visited (no information about the content, interactions, or cookies are tracked or included). No identifiable data (that is who you are) will be published or kept beyond the study period. We will NOT share your collected information with anyone except the investigators listed in the Ethics Protocol. Our website (http://sociology.cass.anu.edu.au/centres/anu-cybercrime) details the specific data we are collecting, and information relating to this research study.

We will make every effort to safeguard and keep confidential the data we collect to the extent permitted by law.

Data Storage

Where: Your web browsing data (stored on our web server) is physically located in the Beryl Rawson building. The server is physically locked to a wall mount, behind a locked door with only the investigators having physical key access, in a secured area that requires ANU ID card access (separate from building access). The log and web-server will be kept separately and secured by ANU IT security with special attention to possible risks of attack or hacking from any source.

How long: Your data will be collected between 20 FEBRUARY 2017 and 31 JULY 2017. Your information may be kept for up to 5 years after any publications arising from the research study. However details such as your name and email address will not be kept and the data will not link any individual with the URL data collected.

Destruction of Data: At the end of the storage procedure, your data stored on disk will be erased securely, then the physical disk will be physically destroyed and disposed in a secure manner.

Queries and Concerns

Contact Details for More Information: Contact me by email: roderic.broadhurst@anu.edu.au to raise queries or concerns about the study. Please include “Queries about cybercrime study (protocol number: 2015/038)” in your subject line so I can prioritise a response. I can be contacted by phone on 6125 4665 and welcome any queries about the research and our work at ANU Cybercrime Observatory.

Ethics Committee Clearance:

The ethical aspects of this research have been approved by the ANU Human Research Ethics Committee in protocol 2015/038. If you have any concerns or complaints about how this research has been conducted, please contact:

Ethics Manager

The ANU Human Research Ethics Committee

The Australian National University

Telephone: +61 2 6125 3427

Email: Human.Ethics.Officer@anu.edu.au

Appendix D: Study reminder email

Good evening,

I am writing about the research titled **Cybercrime Risk in University Communities** that you kindly agreed to participate in. We signed most participants up during O-week.

You may recall the study involves the sending of fake email scams, which will attempt to entice you into revealing private information.

As we are due to commence the research in the coming weeks we wanted to alert you to the start of the experiments and also give an opportunity to opt out of the research.

Please do not hesitate to contact us if you have any questions or concerns.

If you wish to withdraw from the study please contact us by email, either myself (Roderic.Broadhurst@anu.edu.au) or Nick Sifniotis nick.sifniotis@anu.edu.au.

Thank you once again for agreeing to be a research volunteer.

Professor Rod Broadhurst and the Cybercrime Observatory Team

Appendix E: Initial hunter email

Hi,

Thank you once again for choosing to participate in our ANU Cybercrime experiment.

We have selected you to actively hunt for cybercrime throughout the duration of this study. Email scams are one of the most commonly encountered forms of cybercrime and it important that you remain constantly vigilant.

Your task: report any suspicious email content that you come across during this study. This is a simple task, we simply request that you forward suspicious content and suspected scams to us at cyberobs.anu@gmail.com.

There will be a prize for the person who reports the most suspicious content.

Your participation contributes greatly to the identification of online vulnerabilities and the protection of the ANU.

For the duration of this study, you are a **cybercrime hunter**. Good luck!

Cybercrime Observatory Team

Appendix F: Hunter reminder email

Hi,

Just a reminder that you are a **cybercrime hunter** for our study. Email scams are one of the most commonly encountered forms of cybercrime and it is important that you remain *constantly vigilant*.

Your task: report any suspicious email content that you come across during this study. Remember to forward any suspicious content and suspected scams to us at cyberobs.anu@gmail.com.

Good luck!

Cybercrime Observatory Team

Appendix G: Landing page

Australian
National
University

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 online 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 H: Fake scam content (generic and tailored emails)

(1) Generic email #1: 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

(2) Tailored email #1: Final Examination Timetable: Update

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

Student ID:

Name:

Dear,

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-administration/assessments-exams/examination-conduct and http://www.anu.edu.au/students/program-administration/assessments-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]

(3) Generic email #2: Messages

You have (3) unread messages from Australian National University, Click on review to read them.

(4) Tailored email #2: Potential HECS Overcharge

Dear,

Due to a rounding error on the 2017 HECS student contributions, you have been overcharged. The amount could be as high as $430.00 across your Summer/Semester 1/Autumn session enrolments. Please login to ISIS to verify your new student contributions.

2017 enrolments and associated HECS debt have not yet been reported to the Australian Government (Department of Education and the Australian Taxation Office). Summer session and Semester 1 enrolments will be reported at the end of May and we will ensure that the correct amount is reported. New Commonwealth Assistance Notices (CANs) will be issued to those students whose fees have changed.

Our apologies for any inconvenience this has caused.

Kind Regards - Student Central- Division of Student Administration - Building X-005,
121 Marcus Clarke, ANU, 2601

ANU Student Central T: 135 ANU (135 268) E: student@anu.edu.au W: www.anu.edu.au
Commitment to Service - Feedback and Grievances

The Australian National University, Canberra | CRICOS Provider : 00120C | ABN : 52 234 063 906

Never miss a thing: follow us on social media today!

This email message and any accompanying attachments may contain confidential information. If you are not the intended recipient, do not read, use, disseminate, distribute or copy this message or attachments. If you have received this message in error, please delete it immediately. Any views expressed in this message are those of the individual sender, except where the sender expressly, and with authority, states them to be the view of The Australian National University.

(5) Tailored email #3: Results for First Semester, 2017

Dear,

Your results for first semester, 2017 are now available on ISIS. Login here to access them.

Grades

For information on grading scale for your courses, please see the ANU Grading Scale. If you are missing any grades/marks, please contact the respective College for information on its availability.

Supplementary Examinations

If you have been offered a supplementary examination, you have seven working days from the date of notification of the result to notify the relevant College of your acceptance of the offer of a supplementary examination. For more information, please see the Supplementary Exams page.

Assessment Appeals

For any appeals regarding your grades or marks, please contact the Course Convenor or relevant College. For more information, please see the Assessment Appeals page.

Printed Results

You can print a copy of your Statement of Results from ISIS which will be available later today. If you require a certified copy of your results, you may request an ANU Academic Transcript from the ANU Student Exchange.

More information

Please see http://www.anu.edu.au/students/program-administration/assessments-exams for more information on Assessment and Exams.

For further information, please contact Student Central: student@anu.edu.au

Yours sincerely,

Mark Erickson

Registrar, Student Administration

(6) Generic email #3: 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

(7) Tailored email #4: [Students.all] URGENT: ISIS details out of date (sent on 18/8/17)

Good afternoon,

Further to our last email, we have identified that some of your data is out of date. Please log in to ISIS right away and ensure that your details are up to date.

Failure to do so may affect your graduation schedule.

Kind regards,

ANU Student Central

Appendix I: Email distributing the internet survey

Dear,

We would like to thank you for participating in our ANU cyber experiment. Your contribution to our research has been ongoing since the beginning of 2017 and we greatly appreciate your time, effort, and willingness to endure 9 months of fake cybercrime.

The phishing component of our study has officially come to an end. That means no more sneaky emails from us!

For the final part of this study, we would like to ask for 3 minutes of your time. Please help us with our research and complete a short survey about your experiences with the study and with the internet:

https://anu.co1.qualtrics.com/jfe/form/SV_1HLIPCD7bQtPCVT

As an added incentive, you will receive a BrodBurger voucher upon completing the survey, organised by us. This voucher entitles you to one free burger of your choosing, from the BrodBurger in the ANU Popup Village.

In case you were wondering, the survey link is legitimate and the BrodDog voucher is not a scam (we promise)!

Thank you once again for your participation in the cyber experiment. We are expecting to have a summary of our research findings out by the end of November. If you would like to be kept informed of the results of this study, please let us know.

In the meantime please contact us if you have any questions or concerns.

Kind Regards,

Cybercrime Observatory Team (cyberobs.anu@gmail.com)

Professor Rod Broadhurst (roderic.broadhurst@anu.edu.au)

Nick Sifniotis (nick.sifniotis@anu.edu.au)

Katie Skinner (u5184092@anu.edu.au)

Appendix J: Generalised linear model analysis and results

The analysis is performed with R version 3.5.1 (2018-07-02). Three participants with missing values are omitted from the generalised linear model analysis. One of them had indicated gender as ‘Other’ and two others had not specified their student status. All variables included in Model 1 are listed below:

- Y_i represents the number of scams the ith individual has fallen for: 0-7;
- X_{1i} is the indicator of the ith individual’s gender: male or female;
- X_{2i} represents the ith individual’s IT competence: poor < average < above average < advanced;
- X_{3i} is the indicator of the ith individual’s cybercrime awareness: Hunter or Passive;
- X_{4i} represents the ith individual’s perceived internet safety: very unsafe < somewhat unsafe < somewhat safe < very safe; and
- $Offset_i$ adjusts for the total number of scams the ith individual is exposed to: log(7) or log(9).

Model 1 regression equation

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

Table J1: Analysis of deviance table of Model 1

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

The likelihood ratio test of Model 1 against the null model where only the intercept is fitted gives a p value of 0.1653 (Table J1). Thus, there are no significant effects overall in Model 1. Post-hoc analysis identified two additional variables that are highly correlated to scam count:

- W_{1i} is the ith individual’s year of study: Year 1, 2, 3, 4 or 5 (treated as numeric); and
- W_{2i} is the indicator of the ith 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.

Model 2 regression equation

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

Table J2: Analysis of deviance table of Model 2 and the Full Model (Model 1 and Model 2)

Model	Residual *df*	Residual deviance	Additional *df*	Change in defiance	*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

Table J3: Summary of coefficients in Model 2

Model	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

Note: Both variables in Model 2 are significant at 0.05 level

Appendix K: Timeline of phishing experiment

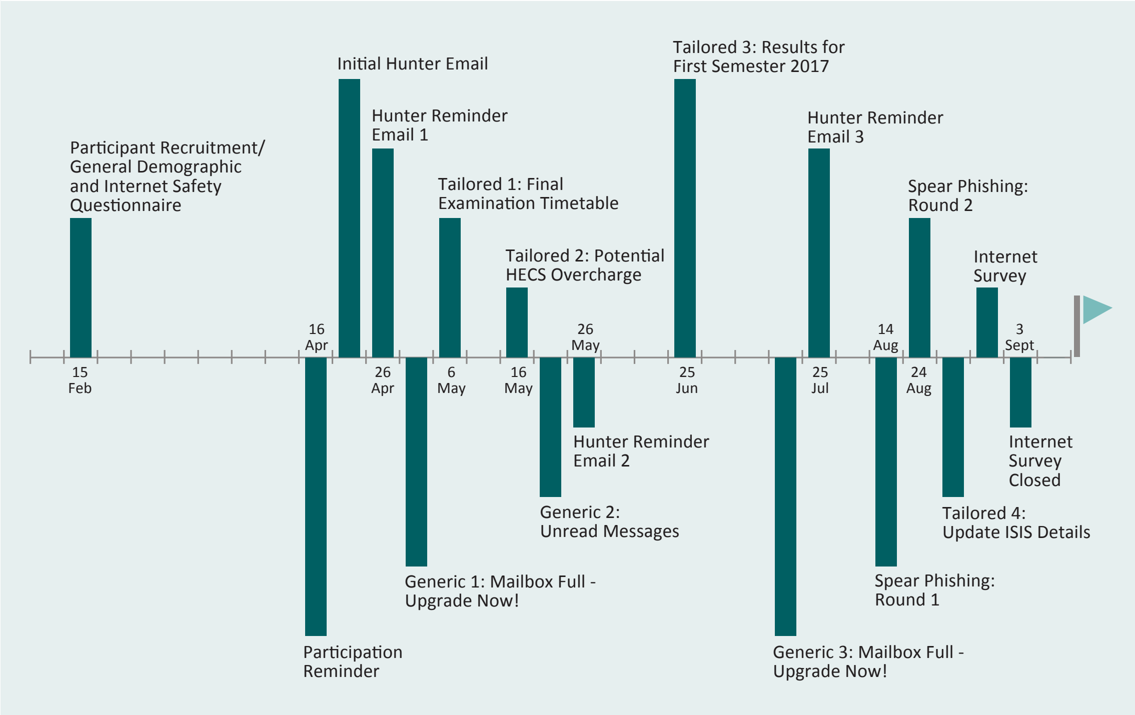

CRG reports

CRG 51/16–17

Roderic Broadhurst is Professor of Criminology at the Australian National University.

Katie Skinner is a Research Assistant at the Australian National University Cybercrime Observatory.

Nick Sifniotis is a Research Assistant at the Australian National University Cybercrime Observatory.

Bryan Matamoros-Macias is a Research Assistant at the Australian National University Cybercrime Observatory.

Yuguang Ipsen is a Lecturer at the Australian National University Research School of Finance, Actuarial Studies and Statistics.

crg.aic.gov.au