



A Qualitative Investigation of Bank Employee Experiences of Information Security and Phishing

Dan Conway, Ronnie Taib, Mitch Harris, Kun Yu, Shlomo Berkovsky,
and Fang Chen, *Data61 - CSIRO*

<https://www.usenix.org/conference/soups2017/technical-sessions/presentation/conway>

**This paper is included in the Proceedings of the
Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017).**

July 12–14, 2017 • Santa Clara, CA, USA

ISBN 978-1-931971-39-3

**Open access to the Proceedings of the
Thirteenth Symposium
on Usable Privacy and Security
is sponsored by USENIX.**

A Qualitative Investigation of Bank Employee Experiences of Information Security and Phishing

Dan Conway, Ronnie Taib, Mitch Harris, Shlomo Berkovsky, Kun Yu, Fang Chen

Data61, CSIRO, Australia
firstname.lastname@data61.csiro.au

ABSTRACT

Staff behaviour is increasingly understood to be an important determinant of an organisations' vulnerability to information security breaches. In parallel to the HCI and CSCW literature, models drawn from cognitive and health psychology have suggested a number of mental variables that predict staff response to security threats. This study began with these models, but engaged in a broader, discovery-orientated, qualitative investigation of how these variables were experienced, interacted subjectively, and what further variables might be of relevance. We conducted in-depth, semi-structured interviews consisting of open and closed questions with staff from a financial services institution under conditions of strict anonymity. Results include a number of findings such as a possible association between highly visible security procedures and low perceptions of vulnerability leading to poor security practices. We also found self-efficacy was a strong determinant of staff sharing stories of negative experiences and variances in the number of non-relevant emails that they process. These findings lead to a richer, deeper understanding of staff experiences in relation to information security and phishing.

1. INTRODUCTION

The roles that staff play in information security (IS) breaches have, of late, become increasingly recognised as important determinants of an organisation's IS defence posture. While the exact classification of breach types remains controversial, reports such as IBM's 2014 Cyber Security Intelligence Index claim that 'human error' was implicated in over 95% of significant data breaches of their systems [1]. As such, it is becoming apparent that purely technical solutions to information security will not be sufficient to address the growing threat to our networks and data posed by cyber criminals and hostile entities.

There are a number of much discussed user failures to comply with IS policies that have shown to be largely explicable using investigations based around user-education and the usability-security trade-off. Examples include; the difficulties in complying with password policies [2, 3], giving away too much personal information when not required [4], and ignoring warning messages when engaging in unsafe behaviour [5]. As these examples

suggest, this body of work is typically based on 'user studies' where the dependent variables are either behavioural, or subjective observations of the behaviours in question. This body of literature also typically focusses on raising user awareness of cyber threats, with the assumption that knowledge will allow people to recognise and deal with attacks. However, a further class of problems requires a different investigative lens. Cyber attackers are now recognised as understanding and leveraging the inherent cognitive biases and weaknesses of the human information processing system [6, 7], enabling them to bypass effortful, deep information processing by the user [8]. This is particularly evident in phishing attacks, which consist of generic, non-targeted emails, distributed widely, that attempt to entice the user to click on a link or open an attachment leading to a malware infection or security credentials being revealed to the attacker. These types of exploits are crafted with increasing sophistication aimed at bypassing conscious processing of the victim and eliciting more automatic behaviours characterised by shallow information processing and as such these methods require new approaches to mitigate [9]. In the face of these kinds of attacks, analyses based on more behavioural methods are likely to fall short, explication requiring a deeper engagement with the cognitive processes that staff experience when facing threats. In this paper we discuss cognitive models that include constructs such as threat Self-Efficacy (SE) and perceived Vulnerability (V). These variables, in particular, have been shown to predict users deploying protective behaviours to a greater extent than knowledge alone [10, 11]. Knowledge is now seen as necessary, but not sufficient to arm users against attackers.

This paper aims to extend the understanding of the human end-user within the IS landscape, specifically seeking to understand the underlying, presumably causative, cognitive variables that drive these behaviours. This work draws on the literature of cognitive psychology and aims to extend the approaches adopted by the HCI and CSCW community. Our study involved the staff of a major financial services institution in Australia and New Zealand. The study was aimed principally at understanding factors implicated in victimisation via phishing attacks, but had as a secondary objective to understand the challenges that staff faced in relation to IS more generally. We were interested in the following research questions:

- What cognitive variables may be implicated in staff's behaviour in relation to phishing emails?
- How do staff experience information security within the organisation and how does this differ from their perceptions at home?
- What environmental and organisational factors affect staff behaviour in relation to phishing attacks and information security more generally?

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2017, July 12 -- 14, 2017, Santa Clara, California.

After carrying out our study and analysing the results, we established a number of emergent themes from the data. These themes are detailed individually in the results section of this work and then the implications are considered and additional context is provided in the discussion section. Many of the themes, such as staff's low feelings of vulnerability, variable proportions of non-relevant emails and willingness to share victimisation experiences only if they have high self-efficacy, immediately suggest further working hypotheses, the primary of which are discussed in the future work section.

2. RELATED WORK

Differing approaches to research in cyber security have resulted in subtly, but fundamentally different bodies of literature around the subject. Each has its own characteristics such as assumptions, methods, and investigative lenses. One body of work, emerging largely from the HCI and CSCW domain, has provided us with a rich picture of the behavioural characteristics of users in response to IS challenges. Acquisiti et al. [4] provided an excellent overview of the way users make poor decisions about privacy. Dhamaja et al. [12] demonstrated the inability of people to detect well-crafted phishing emails, even in ideal conditions, and noted the poor response to security indicators such as status bar warnings. And finally users have also been shown to frequently disclose more information on-line than they need to [13], and are often willing to sacrifice privacy for remarkably small rewards [14]. Overall the picture built up by this research is concerning since it indicates that people are extremely vulnerable to cyber attacks.

Much, but not all, of this work is based on an underlying assumption that educating the user will fix the problem. The core issue is often understood to be 'how do we help users learn more about security so that they can make better decisions'. Influential papers such as that by Kumaraguru et al. [15] are focussed almost entirely on the education issue and the dependent variables of interest in the study are all based around the acquisition, retention, and transfer of knowledge.

Arising from this standpoint, an entire commercial ecosystem has emerged offering to address the 'human problem' of cyber security purely through training and education campaigns. However these approaches are rapidly approaching the point of diminishing returns, where security professionals are frustrated by the persistence of poor user IS behaviour leading some authors to suggest that human-based solutions are not feasible and that technical solutions are the only way to effectively safeguard systems from attackers [16].

However, another body of cyber security research has concerned itself more with understanding the underlying cognitive variables, or mental constructs, underpinning the behaviours of interest. This vein of research has its roots in both health psychology and cognitive psychology and promises to extend the efficacy of mitigation methods beyond that offered by simple education. For example, Samaya et al. [10] recently showed, in an excellently designed study of 3,500 participants across seven countries, that user self-confidence in being able to respond to security threats was a more than four times larger predictor of their measure of good cyber security behaviour, than was knowledge of cyber security threats. Findings such as this, that identify the cognitive constructs that drive behavioural models, promise to be able to extend the effectiveness of mitigation strategies beyond the limitations of current 'education and training' approaches.

Models originally based in health psychology [17] are remarkably suited for deployment in the IS domain since the environments are in many ways analogous. Both IS and health involve individual behaviour, in situations of uncertainty and in response to threats which are often poorly understood, where costs can often be temporally far removed and not deemed likely, and where compliance with desired protocols (often referred to as response costs) is either arduous, or not immediately desirable. According to these models, the challenge of eating well and exercising regularly is almost perfectly analogous to deploying strong and different passwords on every system you use.

Many of these current psychological models of behaviour in response to threats are derived to some degree from the hugely influential Theory of Reasoned Action (TRA) [18]. This model proposes that beliefs about a behaviour and evaluations of the outcome of a behaviour result in attitudes towards the behaviour, and that social influences and motivation result in subjective norms. These two constructs; attitude towards the behaviour and subjective norms, then interact to result in behavioural intention, which in turn predicts the behaviour itself. Ajzen's later reformulation of the model into the Theory of Planned Behaviour (TPB) [19] involved the addition of variables that accounted for a person's own beliefs about their ability to carry out the behaviours in question. Referred to as perceived behavioural control and later disambiguated further to variables such as locus of control [20], self-efficacy [17, 21] and response-efficacy [22], these variables have a long history of being shown to be significant predictors of behavioural intentions. These variables are also deployed in contemporary models of protective behaviour such as the Protection Motivation Theory (PMT) [23]. In short, the constructs encompassed in perceived behavioural control suggest that people are unlikely to attempt to engage in a behaviour if they think that they will not be able to carry out the behaviour in question. Constructs such as these are central to our investigations and are discussed in more detail in the discussion section of this paper.

Another prominent model emerging from health psychology literature is Rogers' Protection Motivation Theory (PMT) [11] which was derived from ideas about people's response to fear, and suggested that encountering a communication that induced fear would induce a threat appraisal process which, mediated by variables such as response efficacy, self-efficacy and response costs, would result in protection motivation leading to either an adaptive or maladaptive response to the threat.

More recently still, in the cognitive domain, dual route models of information processing such as the Heuristic Systematic Model (HSM) have begun to be applied specifically to the problem of phishing victimisation with notable successes in predicting user behaviour [9, 24]. Dual Route models suggest that users often engage in little elaborative, deep ('system 2') information processing when scanning emails, and rely instead of more shallow (system 1) information processing based on simple heuristics such as calls to authority, urgency cues and social proof [25] to make fast decisions about whether to respond or not [8]. These models suggest that regardless of a users' knowledge of threats, when scanning an inbox for messages to respond to, users often engage in very shallow cognitive processing of email cues such as sender and subject line, meaning that they are not deploying the knowledge that they have about these cues. This results in important signals such as malformed email addresses (type-jacking) escaping attention. According to these theories, in this low level of cognitive

engagement with the task, people are more likely to respond to simple heuristic rules of thumb such as ‘this email is marked as urgent’ and ‘Oh this is a reputable brand – it must be ok’ and are therefore enticed to click on emails that would, if given more thought, appear suspicious.

3. THEORETICAL FOUNDATIONS

In this section, we discuss the theoretical standpoint from which this research was carried out.

We sought to set out from the previously established findings but engage in a more discovery-orientated investigation. Thus we aimed at uncovering the ‘unknown unknowns’ of this particular space and thereby be better equipped to later select specific models to apply – or develop new models altogether. As such, we deployed qualitative methods, with the aim of gaining insight into the cognitions, work contexts, motivations, normative influences, and everyday practices of staff as they experience phishing attempts. Although our approach was not a full implementation of grounded theory, we deployed many of the techniques prescribed by this method, seeking knowledge from the ground-up and asking questions with varying degrees of specificity in order to probe specific areas of interest.

Qualitative methods are useful for identifying new and undiscovered phenomena, providing deeper insights into user experiences than quantitative measures can provide, may be transferable to populations equivalent to the sample group and can uncover themes that may be later tested with more quantitative approaches [26]. Furthermore, the depth of detail and nuanced, semantic-based responses provide a richer, deeper understanding of the problem-space than offered by higher-*n* quantitative studies with less time devoted to each subject [27].

In this, our research was successful in that it uncovered evidence of both a number of widely reported dynamics and phenomena in the field, as well as promising results that were novel or even contradicted prevailing knowledge in the literature.

4. INSTRUMENT DEVELOPMENT

Starting out from the variables deployed by the theories above, we developed a 38-item questionnaire with questions grouped by topics: Knowledge, Attitudes, History, Practices, Contexts and Identity (see appendix A for the complete instrument). Since we sought to extend our investigation beyond the known, specific constructs of the models in question, we formulated many open questions designed to elicit non-structured, wide-ranging responses. An overview of the literature consulted in the process of developing the questions for each topic is included below.

4.1 Knowledge

Individuals’ knowledge of cyber security threats as well as computer literacy and expertise have been proposed as important determinants of protective behaviours. Furthermore, specific variables such as threat awareness and countermeasure awareness have been posited as predictors of IS policy compliance [17]. However, as Stephanou showed [28], while education/training campaigns have measurable impact on staff knowledge of the desired behaviours, they are not necessarily correlated with actual subsequent behaviour suggesting that education is necessary but not sufficient to mitigate victimisation. As such, the questions included in the Knowledge topic were designed to gauge the depths

of people’s understanding of the domain generally, as well as elicit more emotional and relational perceptions of their role and interactions with others in this context. We wanted to understand how participants thought of and spoke about IS and how it affected them in their everyday work lives. Therefore we developed five questions (Q1, Q1A-D, see Appendix A) and grouped these under the Knowledge topic. These were deliberately broad, open questions, designed to elicit conceptions about security in the most general terms possible and in ways that were most cognitively available and important to participants. Q1-B was designed to elicit conceptions around who were the actors in the IS space – both within the organisation and outside – to try to understand whom participants interacted with and had relationships with in relation to the subject.

4.2 Attitudes

Ever since La Pierre showed that when questions about attitudes are posed broadly they are poor predictors of specific behaviours [29] attitudes have long been known to have a complicated relationship with behaviour. As such we set out to understand how the most commonly implicated variables in IS behavioural models of attitudes were experienced by staff and what kind of situational factors fed into these variables. We were also interested in people’s value systems and how ideas about the importance of IS impacted on their intention to comply with mandated security protocols.

Perceptions of vulnerability have been found to be important predictors of people engaging in protective behaviours in a number of models such as the PMT [30]. Perceptions of fear are also central to the threat appraisal process described in this model [11, 31]. Thus we were specifically interested in ideas around vulnerability and fear and formulated questions Q2 to Q4B in a manner that illuminated the contexts in which they are evoked, and categorised these as belonging to the ‘Attitudes’ Topic.

4.3 History

We were interested in the effects of previous phishing or fraud victimisation on subsequent behaviour and attitudes. Research into related constructs such as ‘threat awareness’ and ‘domain knowledge’ has yielded mixed results. Wang et al. [32] showed that prior ‘scam knowledge’ decreases attention to ‘visceral triggers’ and increases attention to deceptive elements in fraudulent emails. However Vishwanath et al. [33], deploying a dual process model of information processing, found that domain specific knowledge (a construct including experience and exposure) significantly predicted elaboration likelihood (deeper processing of information according to dual route theories) in only one half of a split-half test population and the relationship was therefore only partially supported. However, in both these cases experience was not directly related to previous victimisation and instead consisted of education and exposure to information without the negative outcomes associated with actual victimisation. On other hand, Böhme et al. [34] found that experience with e-commerce fraud was likely to reduce subsequent on-line purchasing behaviours, but that as a predictor, the effect size was less than ‘general concerns’ and ‘personal concerns’. In a result that may explain some of the variance in findings above, Yu [35] found that victimisation experience significantly affected subsequent fear of cybercrime for cyber bullying, digital piracy and computer viruses – but not for on-line scams, suggesting that post-incident fear perceptions are dependent on the nature of the crime itself.

In light of these seemingly variable findings, we developed a number of questions aimed at understanding participants' real world experience with phishing victimisation and cyber-fraud and how these experiences affected their subsequent and ongoing practices and cognitions around IS (Q5-Q12).

4.4 Practices

We were also interested in staff's experiential relationship with specific known challenges for IS. The questions in this topic were designed to elicit discussion on behaviours around passwords, use of free Wi-Fi, and then, in more detail; email practices (Q13-Q17). The relationship between system usability and restrictive security procedures has been much discussed. For example Post and Kagan [36] showed that increased requirements around the complexity and diversity in user-generated passwords resulted in ever increasing cognitive demands often leading to more risky behaviours (such as writing passwords down). This particular security-usability trade-off is also highly explicable to users and as such, we wanted to understand both their practices and attitudes towards password use as a proxy for more general behaviour around IS.

Since the financial institution involved in our study had a long-standing and significant investment in IS education including an extensive training program, an information portal, phishing drills, and awareness events we were also interested in staff awareness of and response to these engagements (Q18-22).

There has also been discussion about what educational formats are most effective at engaging staff [37] – so we asked about both staff appetite for more learning – and their ideal format for educational materials (Q23).

4.5 Contexts

Behaviour does not occur in a vacuum, and as such we were interested in gaining insight into the environmental and organisational factors that impact on the work practices associated with phishing victimisation. Much of this section was specifically designed to elicit information about staff practices in relation to email as the primary vector for phishing attacks.

Dual route models of information processing suggest that users typically engage in little elaborative information processing when scanning emails, and rely instead of more shallow evaluations based on heuristics such as calls to authority, urgency cues and social proof to make fast decisions about responding [3, 16]. Importantly these attentional-based theories suggest that education is unlikely to be sufficient to curb risky behaviour if the user never engages their implicit knowledge of the subject matter in order to evaluate threats. These attentional models also suggest that a new range of variables – such as workload, attentional resources and task demands (as well as individual differences such as need for cognition [8]) are important determinants of phishing victimisation. Mark et al. [38] showed that some email usage patterns result in users feeling cognitively overloaded and stressed. This, in conjunction with the finding by Vishwanath et al. [33] that the number of emails that users engage with daily (email load) significantly increased the likelihood of falling for phishing attacks, suggest that the sheer volume of emails people respond to provides significant challenges to people's available cognitive resources to evaluate threats. Mark et al. also noted different patterns of behaviour around responding to emails, such as users who process in 'batches' at pre-determined times, to those who check their inbox

constantly throughout the day, and those who respond to notifications in real-time. The ramifications of these different patterns of email interaction on phishing responses has yet to be investigated so we formulated Questions 24-29 in order to better understand staff behaviours in this area.

Additionally, two much discussed variables in IS behaviour are those relating to the punitive measures that organisations deploy in response to poor staff security behaviour, specifically perceived certainty of sanctions and perceived sanction severity. These variables are controversial since while prevailing thought within the criminologist domain suggests that increasing these variables leads to more desired behaviours [39] – other studies have shown that they are only weakly predictive [29], particularly when there are avenues to neutralise the effects of their non-compliance [41].

Furthermore, we were also interested in the normative environment in which participants existed and informational and cultural influences on staff attitudes. Dodge et al. [42] amongst many other has shown that staff who have leaders that espouse IS protocols and lead by example are more likely to adopt good practices themselves. Flores et al. [43] showed that both the variables of transformation leadership (where leaders involve subordinates in decision making and driving change) and IS culture were both significant predictors of IS awareness which in turn had significant effects on intrinsic beliefs and then intentions. In terms of social norms and peer influence, Ifinedo [21] deployed Social Bond Theory to show that the four constructs of attachment (to an organisations values), commitment (to an organisations goals), involvement (in an organisations goals) and personal norms were all significant predictors of subjective norms and that subjective norms had a positive effect on attitudes towards IS policy compliance. As such, we asked a number of questions (Q31A-C, Q32A-C) about where people learnt about IS from, and then also about whether they talked about, learnt from or thought of as valued by both peers and managers.

4.6 Identity

In models derived from Ajzen's Theory of Planned Behaviour, as well as later variants such as the PMT, various elements of a person's ideas about their selves, such as locus of control, self-efficacy and response efficacy have been shown to be important predictors of behaviour [44]. Furthermore, motivation has been shown to have a causal relationship with elaborative processing, as expressed in more attention-based models such as the HSM. However, in a finding that poses challenges for motivational-based models, Floyd et al. [23] showed evidence that that self-efficacy was not correlated with elaboration likelihood in evaluating phishing emails whereas level of involvement was.

Therefore, in addition to the ideas about subjective norms which we included in the contexts section above, in this subset of questions (Q33-Q38) we wanted to understand how staff engaged with their work practices in ways that may be affected by such variables such as self-efficacy, response-efficacy, responsibility and locus of control.

5. METHOD

5.1 Participants

18 staff situated in Australia (4) and New Zealand (14) from a major financial services institution, took part in the study (8 female,

10 male). Staff were recruited via emails sent to a cross section of staff. Some organic recruitment also took place as staff began to forward the invitation to colleagues. Staff were offered the opportunity to win double movies passes as recompense for participation, were instructed as to the anonymity protocols involved and informed of the voluntary nature of the experiment via the automated, sign-up web service provided by the bank. Participation consent, and consent for the experimenters making a recording was provided at the beginning of the interview session. The functional roles of the participants are listed in Table 1 below.

Table 1: Functional roles of participants of the study.

Position Category	Number of participants
Customer services, support and sales	4
Technical + Operations	8
Managers	2
Finance and Risk	4

5.2 Apparatus

The 38-item questionnaire had questions grouped by topics: Knowledge, Attitudes, History, Practices, Contexts and Identity (see appendix A for the complete instrument). Many questions were open and designed to elicit extensive, wide-ranging responses. Since we had limited time (30 minutes per interview), insufficient to administer all questions, a randomly selected subset of topics was differentially applied to each participant – with the exception of Knowledge questions – which were applied to all participants. Coverage of topics across participants is detailed in table 2. As many topics were applied to each participant as time permitted. Some participants offered much more detailed, and therefore time-consuming responses than others, leading to an uneven number of topics covered by different participants.

5.3 Iteration

After the first two days of interviews, consisting of 12 participants, an initial analysis of responses was made to determine emergent themes. Based on this analysis, seven additional questions (Q201-207) were developed and then were administered to the remaining participants during the second interview session, referred to as ‘round two’. These questions were interleaved with the existing questions according to their topic. The responses arising from these questions are discussed in the results section according to the category that gave rise to each question.

5.4 Procedure

All interviews were carried out remotely with participants ‘dialling in’ to a conference call from their premises. The interviewers remained at *Data61* premises and were visible via webcam for the first four interviews – but then, after finding that this was of limited utility, for all subsequent interviews only audio was used. Participants only provided audio and were never visible to the interviewers. Throughout the recruitment process, participant anonymity was stressed, and owing to the protocols we deployed, participants were only known to the interviewers by their ‘Made-up’ ID provided by the bank. This approach seemed to reassure interviewees, and they appeared to speak freely and without evidence of much social desirability bias present in their responses.

5.5 Coding

Interviews were transcribed in full by various authors, with one interview being lost owing to a failure of audio recording equipment. For this participant detailed interviewer notes were used for analysis. All coding was then carried out by the main author, with frequent access to the original recordings for clarification. Coding began with categories suggested by the cognitive variables in related work as detailed in sections 3 and 4. Additional categories were then developed from the data itself as analysis progressed and ones where known variables did not yield fruitful new information were abandoned. For each category identified, the entire body of transcripts was then re-analysed for additional data pertaining to the category uncovered. Further distinctions were made within categories as the data suggested leading in some cases to new questions being developed and deployed in round two of interviews. Eventually a two level taxonomy of findings was established consisting of general categories of responses with sub-themes. Through this process, we achieved saturation, i.e.: a state where little fresh information emerges from subsequent interviews because all the main themes have been uncovered, within our 18 participants [27].

Table 2. Basic demographics and coverage of topics by participant.

	Age bracket	Sex	Years with organisation	Knowledge	Attitudes	History	Practices	Contexts	Identity	Round two questions
P1	45-54	M	6	Yes	P			Yes		
P2	45-54	M	10	Yes	Yes	Yes	P	Yes	Yes	
P3	35-44	M	6	Yes	Yes	P	Yes		Yes	
P4	35-44	F	16	Yes	P		P		P	
P5	35-44	F	8	Yes	Yes		P	Yes		
P6	35-44	M	10	Yes			Yes		Yes	
P7	35-44	F	20	Yes	Yes			P	Yes	
P8	35-44	M	2	Yes			Yes	Yes		
P9	35-44	F	16	Yes	Yes	P	Yes	Yes	Yes	
P10	20-24	M	2	Yes	Yes	P	Yes		P	
P11	25-34	F	7	Yes		P	P			
P12	25-34	M	1	Yes	Yes			Yes		
P13	35-44	F	5	Yes	Yes	Yes	P	Yes	Yes	Yes
P14	45-54	F	1	Yes	Yes	Yes	P	Yes	Yes	Yes
P15	35-44	F	14	Yes	Yes	Yes	Yes	Yes	Yes	Yes
P16	45-54	M	3	Yes	Yes	Yes		Yes		Yes
P17	25-34	M	10	Yes	Yes			Yes	Yes	Yes
P18	>64	M	3	Yes	Yes	Yes		Yes	Yes	Yes

Note: P = Partial coverage of questions for this topic.

6. RESULTS

Results are presented in three categories, organised by the themes emerging from the interviews themselves: information, education and knowledge sharing; experience of email practices; perceptions

of threat, vulnerability and responsibility. Note that these categories and themes arise from the data itself and are therefore not directly related to the topics in which the questions were originally grouped.

6.1 Information, education and knowledge sharing

6.1.1 *Passive learning is taking place, but active learning needs to be facilitated*

We asked several questions designed to establish staff's sources of information about IS and their engagement with and opinions of those sources of information (Q18, 19, 20, 21, 22, 23 and 31A, 31B). We found a wide variety of practices around learning about IS including:

- Learning at specific training events
- As part of their job requirements (noted for staff in more technical and IT related positions).
- The bank's intranet.
- Weekly email bulletins.
- Monthly email bulletins.
- Outside sources of information such as IS websites and third party company security warnings.

Staff generally stated that they received information about IS and viewed this information in a positive light and as a necessary part of their responsibilities within the bank (see more detail on this in section 6.3.5).

But importantly, it should be noted that most of the modes of education staff referenced are passive – and when asked about where staff would go if they had questions about IS issues – there was a high degree of uncertainty.

"I don't even know if we've got that kind of stuff on our website." P6

"I said before I don't think we have specific training on cyber crime. There's no specific modules around it..." P7

"I would probably go to <name of internal corporate intranet - redacted>. And I would search for security and probably email them or call them and let them know something that had happened, and if they're not the right person then ideally they help you find who the right person is." P14

6.1.2 *Happy to scan an information email for new knowledge*

While asking participants about their sources of information on IS we uncovered a recurring pattern of usage of information received via email bulletins. Participants, at all levels of security awareness, expressed positivity about receiving periodic information about IS. When prompted to elaborate on this engagement many responses were characterised by the idea of there not necessarily being much new in the content, but being willing to scan over the information to search for any new pieces of information.

"...yeah it's definitely good reminders... It's timely, I don't think it's overwhelming..." P5

"...there's nothing I would read word for word, but I would definitely scan my eyes over it." P13

"I would say it's mostly a repeat. I can see what they're trying to do, and that, the intent of the bank as an organisation needs to insure that all of the staff understands the whole deal. So, the information that comes to where I am is fairly regular." P16

This finding is an encouraging indicator that staff value periodic information provided by their employer and furthermore have developed nuanced and agentic levels of engagement with these channels to extract information that they see as pertinent to them.

6.1.3 *How would people like their information presented? Short, snappy and based on real-life scenarios*

Participants experience with information delivery was of particular interest to us so Q23 was specifically crafted to uncover ideas about preferred format of IS information and training. While short videos were mentioned occasionally, most respondents expressed a clear preference for text-based communication and brevity was seen as an important requirement.

Furthermore, a number of respondents all converged on a single underlying theme – the desire for education based on user experiences, outcomes and specific mitigation techniques.

"When you're building something around info security training if it's a real life thing that actually happened." P3

"I think something that is a brief short and sharp one or two reduced snippet sort of a thing which says look: 'here is what happens if you did this and here is how you could avoid that'." P8

"But I would really, really ask for something very brief. I feel as if people just LOVE filling up a page with words. And I think bullet points. Can't go past a bullet point." P13

Participant responses indicated a set of highly specific criteria for information consumption, characterised by ideas of brevity and the desire for information that is tied to their own personal experiences and practices.

6.1.4 *Communicating after a bad event – The implications of prior experience*

We asked participants about their previous experience with both phishing email victimisation and on-line crime more generally. We were interested in how staff experienced these incidents and what meanings they ascribed to the events and then further how it shaped their ongoing behaviours. (Q6-Q10). In analysing the resultant responses, we noted that responses to Q10 ('Did you tell anyone about it?') varied in what appeared to be a systemic manner that fell into two distinct groups of responses.

Group 1:

Participants reported telling friends and colleagues about a negative IS incident following the fact. In all cases, the stated motivation was to assist others in avoiding the same problem that they experienced. Interestingly in all cases where participants reported broadcasting their negative experience, they also demonstrated high levels of technical awareness and rated themselves as highly competent with computers.

"I certainly did. I spoke to my colleagues, my friends, sort of tried to make sure that people are not getting into that." P8

"Uhh yeah I did tell my colleagues about it, yes." P17

Group 2:

Participants suggested that they did not want to tell anyone about their experience and specifically thought that it would reflect badly on them. In these cases, participants saw themselves as being less technically adept.

"I may not umm more so if people think 'how dumb she is' <laughs>" P15

"Oh definitely - I was definitely embarrassed. A sense of 'how did I not see that?'" P6

6.2 Experience of email practices

6.2.1 Scanning your inbox – where mistakes happen.

When asked about the quantity of emails employees received in a given day, participants volunteered a wide variety of responses ranging from '10' to 'thousands'. However, despite this variety a common theme emerged of the process of quickly 'scanning' their mailbox for important items in order to identify items that were important or time sensitive.

"You tend to - where you might have glanced at an email before and read a few sentences from the subject heading - to know a bit more before you make that judgement - when it's busy and you're stressed - you look at the subject header and you look at the 'to' box and if you're not in there and if you're not called out in the subject as action - you don't look at it." P6

"So if it's someone I'm currently working with I'll read it straight away. If it's like - a general email to a lot of people - then I'll be like 'Well do I have time? Nope - I'll look at it later'." P2

"I would quickly look at who sent it and the content - oh not the content - the subject line and determine whether it's worth me looking at it straight away then I'll flag emails myself to what priority." P5

This finding on its own may not be significant, but when coupled with both the findings about the amount of non-relevant emails (section 6.2.2) and staff's periodic variance in workload (section 6.2.3) – this may be an indication of circumstances when people's cognitive processing of emails is more shallow during busy work periods and are therefore more likely to click harmful links and attachments.

6.2.2 Some people receive high volumes of non-relevant content

Email practices are obviously a primary consideration when investigating staff behaviour in response to phishing attacks. We asked a number of questions designed to elicit staff experiences around practices and contexts when processing incoming emails – both at work and at home (Q15-Q17, Q24-Q29, Q35, Q38). During the initial analysis of session one interviews, we noted a consistent theme emerging where participants would nominate a number of emails that they would receive each day, but then would later modify that amount in respect to how many they thought were actually relevant to them. As a result we added Q207 ("Do you get

a lot of emails that aren't really relevant to you? Or are trivial?') to the interviews for participants in session two.

We uncovered that some, but not all, participants talked about having to deal with large numbers of emails that were not particularly relevant to them, or were trivial. These included:

- FYI emails where people were generating a paper trail in order to share responsibility or visibility for a decision or process, but again no action was required of the recipient.
- Spam (non-phishing) emails.
- Magazines and informational emails (presumably via subscription).
- 'Marketing' and promotional emails (presumably un-invited and as a result of submitting user details to an external party).

"Umm yeah a lot of the emails are sort of marketing emails." P17

"Definitely. ... you'd get in any given day where you would skim read it and say 'ok, great, fine, filed'." P6

6.2.3 Periodic variations in workload

Participants were asked several questions focussed around email practices and time pressure at work (Q24 - Q29). After round one of interviews, initial analysis showed that many participants found it difficult to respond to these questions systematically because of the variation in their workload over time. We therefore developed two additional questions (Q201, Q203) that were asked of participants in session two interviews, specifically aimed at exploring this phenomenon.

We found that while a few described their positions as being reasonably stable in terms of workload and time pressure, others indicated a large amount of 'seasonal' variation of these attributes.

"It can get crazy busy, it can get insanely busy and at other times - it can be quite relaxed." P2

"Our days are very umm... no single day is the same." P5

"There are phases when it's very, very busy and you definitely do feel the pressure. But that phase comes once every few weeks. And then it <unintelligible> goes back to normal where it's not so much of a time pressure. ... I think my behaviour changes significantly at that point in time, or during those phases." P8

Furthermore, when we probed deeper and asked participants to expand on their email practices during these different periods – many staff stated that they thought these would vary considerably according to the workload at the time.

"When it's busy and you're stressed - you look at the subject header and you look at the 'to' box and if you're not in there and if you're not called out in the subject as action - you don't look at it." P6

"If it's like - a general email to a lot of people - then I'll be like 'Well do I have time? Nope - I'll look at it later.'" P2

6.3 Perceptions of threat vulnerability and responsibility

6.3.1 'At work I feel safe' – Lack of vulnerability

In order to explore staff feelings around feelings of vulnerability we asked the question: 'Do you feel vulnerable to IS threats?' (Q4A). The majority of responses indicated surprisingly low feelings of vulnerability in response to this question.

"Umm at work I feel confident. Umm that our technology team work very hard." P5

"No. Not in the slightest." (Q4A) P13

"...probably not so much at work... because I'm pretty sure I feel like we've got good processes in place at work..." (Q4A) P4

The few responses that did indicate some degree of vulnerability were only offered by staff with high levels of cyber security awareness, and were couched in terms of nothing being fully secure.

"Well everything is vulnerable - You never know." P12

"So I feel like because I'm aware and I know to speak up about it and double check things, that I am quite safe, myself... umm however I guess it always plays in the back of your mind." P6

6.3.2 Information security in home contexts – Far more vulnerable, but less to lose

While we noted a high degree of confidence in the bank's security apparatus generally to protect them from the worst of information attacks, (see section 6.3.1), there seemed to be an acute consciousness amongst staff interviewed that these defence mechanisms were not available at home or on their personal devices. Thus for interview session two we added Q202: 'So what is the difference between thinking about, or IS practices at home compared to at work?' Responses to this question reinforced the finding that feelings of vulnerability were higher at home than at work.

"Whereas at home - you're that person that is susceptible to all those things - and those safety measures that the organisations put in place so therefore you think that much more about it. Or you SHOULD think that much more about it." P6

"And personally - umm - outside of work umm <laughs> - not so protected!" P5

"I might be even more conscious because I know that if anything goes wrong I'm going to have to sort it out - whereas at work if it goes wrong at least we have support networks to help us." P9

Additionally many participants reported differences in the sense of ownership/responsibility of the problem-space compared with at work. This was particularly true for employees who take on a lead role in managing IT systems for their household.

"At home you are tech security - well I am. <laughs> Whereas at work I'm not." P3

"I'm forever telling my wife of the latest scam that's happened." P6

But counter-intuitively, participants often reported more permissive and less stringent IS behaviour in the home and specifically talked about this in the context of the consequences being less important. This finding is mysterious and requires further investigation – specifically operationalising constructs around locus of control, vulnerability and threat severity.

"...but at home there's more risk because I probably don't have as strong a firewall." P7

"When I'm at home, I'm a bit more loose with my emails but I don't click on links." P10

6.3.3 Got scammed? Money was returned so no real loss! Lack of vulnerability

After investigating feelings of vulnerability, we extended our research into the area of perceived consequences of IS breaches. After identifying those participants who had experienced an episode of cyber fraud, we probed extensively into their experiences and reactions (Q5-Q9). We found that in all cases, respondents reported that the fraudulent transactions affecting their accounts were reversed by the bank. In some cases this happened quite quickly, while in one case only a partial reimbursement took place. All staff had banking accounts with the financial institution in question. When discussing these events, respondents were highly positive about their bank's response and offered high estimations of the bank's processes in these cases.

"But <bank name redacted> were brilliant. Seriously. Within 24 hours I think they had replaced the limit and <unintelligible> take care of it. So for me -whatever happens after that, doesn't really matter." P13

Interviewer: "In your case there was no consequence because it was stopped pretty much immediately, is that right?"

Participant: "Yep." P18

This can be seen to equating to the much discussed variable of threat severity which has also been discussed extensively as contributing (when high) to protective behaviours.

6.3.4 Security failures equated with loss of trust in the bank

In order to try to understand ideas of loss and the perceptions of consequences of poor IS behaviours we asked Q1d ('Why is information security important?') and then also Q12 ('What is the worst thing that could happen as a result of a phishing attack on you?'). As well as a host of responses detailing specific worst case scenarios centred around data loss and fraud, we noticed a repeated theme amongst many of the participants who volunteered that one of the worst consequences of successful attacks would be the reputation of the bank itself. It would appear that staff are highly cognisant of the wider implications data breaches and place a consistently high value on the reputation of the bank.

"...because at the end of the day - it is going to affect the um - what is the word I'm looking for - the name of my employer. So <bank name redacted> at the end of the day will be affected and we don't want it to be named and shamed in any way. So the reputation is at risk." P15

"We're a bank, banks are built on trust, if we don't have the trust of our customers, we're out of business." P1

“And the potential is, if we do it wrong, really badly wrong - and we lose that money - it's not a good thing. And I think primarily trust. Customer trust in us.” P3

“Because as a bank, we have a very high trust mandate...” P16

6.3.5 Responsibility for security – and identifying with the bank

Since a higher internal locus of control has been shown to be a necessary but not sufficient antecedent of engaging protective behaviours we operationalised this concept in several questions designed to elicit staff attitudes towards who was responsible for preventing IS attacks (Q1b, Q1d, Q3). During the initial analysis and iteration process after interview session one, we noted that many people talked about this issue quite specifically in relation to their perceived identity – as an employee of a bank, so we added Q204 (‘So you work for a bank – does that bring with it any special responsibilities and roles?’) for interview session two.

We noted that staff universally offered two primary observations about their perceived responsibility for IS. Firstly – and almost always offered as a response to Q1b – people volunteered that they, as an individual, were the primary actor and determinant in this space.

“Me as an individual I am primarily responsible for my own security...” P14

“Well I think I'm the primary.” P13

“It's yourself and anybody who's responsible for public and private networks, and the ownership of those.” P18

“I think it starts with you as a person. But I think everyone is involved...” P11

This assertion of individual responsibility was then almost always followed by a secondary consideration – that of a shared responsibility with others – primarily the bank – but often institutions generally and the collective ideas of the staff at large – all seen as powerful outside forces in the equation.

“So, cybersecurity as a holistic level really comes back to the education of everyone.” P16

“...everybody's! <laughs> Everybody who is involved in the network and anybody who's responsible for their own approaches <unintelligible> and use of the system.” P18

“Everyone should be aware of what's happening...” P2

“Protect my details as much as I can possible, I absolutely would take full responsibility for that. But at the same time I'm happy to lean on the bank when things do go pear shaped.” P13

7. DISCUSSION

Analysis of participant responses revealed a number of novel observations as well as confirmed some findings well established in the literature. Here we discuss the wider implications of these findings on future work in this domain as well as possible real-world applications to combat cyber-crime.

7.1 Information, education and knowledge sharing.

Since knowledge of cyber threats has been shown as a necessary but not sufficient pre-requisite for users carrying out appropriate protective measures [40], the insights established when asking about how users receive and participate in information sharing are encouraging. In section 6.1.1 we learnt about how participants learnt about IS and engaged with a wide variety of sources of information. However, much of this information seemed to not comply with basic instructional design principles such as those discussed by Kumaraguru et al. [15]. Furthermore when asked about where users would go if they had questions, there was much uncertainty, suggesting that, in this particular context, more active modes of information acquisition should be further facilitated. In particular, this reflects a lack of the basic instructional design principles iterated by Kumaraguru et al. where training was most likely to be effective when offered at the right time – i.e.: when participants were interested in learning or those moments when the information is particularly salient. This distinction between passive and active learning raises interesting questions about what additional protocols can be developed to meet staff needs and how active learning behaviours varies across the population in question.

In section 6.1.2, we showed that the participants we spoke to were not fatigued by regularly scheduled information broadcasts from the bank and that these established communications conduits were perceived as useful. More specifically, a behaviour pattern that repeatedly presented itself whereby users would scan over periodic IS related bulletins, assess what was relevant to them and then engage with new material that they deemed relevant to them, indicating a nuanced and agentic engagement with these sources of information. This process is also interesting in relation to our finding in 6.1.3 where the desired modes of information presentation (brief and based on personalised and highly situated stories) were both immediately cognitively available to users as desirable and quite specific, again reflecting Kumaraguru's et al. [15] instructional design principles of personalisation, contiguity and story based agent environment as well as the findings by Harbach et al. [45] on the effectiveness of information personalisation. This reinforces the need for increasing the bandwidth of existing channels of information to staff by adopting principles of brevity and presenting narrative based information based on specific user personas. Furthermore, this supports an emerging trend by industry where education efforts are personalised for different categories of user in order to increase relevance and engagement.

When asking about information sharing with peers generally (section 6.1.4), no consistent pattern of responses occurred, however when probing into experiences of cyber-crime victimisation, an interesting theme emerged. Staff with a self-image of being highly computer-literate and technically ‘savvy’ seemed more likely to communicate to others about their negative experiences, than those with less certainty about their technical skills. For those with low technical self-efficacy, cognitions around victimisation were more centred around the risk of appearing foolish or careless if they talked about their experience. Since an oft stated goal of IS culture is to increase normative information sharing [46] - from an organisational standpoint this would imply that reaching out to those with high levels of confidence in their technical abilities may facilitate the development of localised ‘champions of change’ throughout the organisation. This approach

would extend the work by Sauvik Das et al. in the area of social proof, where normative influence was shown to significantly increase uptake of additional FB security features. Furthermore this approach specifically suggests avenues to overcome the significant challenge for deploying on social proof dynamics where engaging in protective technologies is not easily visible to others.

7.2 Experience of email practices

While the HCI literature has grappled extensively with the phenomena of email processing and problems such as overload [38] and interaction patterns such as task switching, interruption lag and resumption lag [47], leveraging these findings in order to mitigate phishing threats has remained elusive. However, more recently, models emphasizing the attentional aspects of phishing victimization have appeared promising. The Heuristic Systematic Model (HSM) with its focus on competing (and interacting) shallow and deep information processing mechanisms has been shown to predict some degree of phishing victimization. As such our discovery of a number of real-world phenomena present in staff email usage that are likely to impact on the application of models such as the HSM to the problem may prove useful in future efforts.

We found evidence of shallow processing of incoming emails when people talked about ‘scanning’ their inbox for emails that needed an immediate response (section 6.2.1). This reflects Neustadter’s behavioural findings of the email triage process [48], such as the common tactic of attempting to remove trivial emails in order to make it easier to find more important ones. However, this behavioural analysis needs to be extended by a deeper understanding of the mental processes involved in order to effectively mitigate phishing victimisation. Models based on attentional theories, such as the HSM offer good utility here. An example can be found in Xu’s [8] exposition of the likelihood of elaborative processing in email processing according to personality traits such as need for cognition and contextual variables such as recipient expertise and recipient involvement (a motivational factor). Based on a similar dual process model, Vishwanath et al. [33] showed that most email is processed peripherally and that SE was an important factor in users engaging in elaborative processing. Furthermore, in findings that extend the exposition by Neustaeder et al. on the email triage process, Floyd [23] showed that calls to scarcity or urgency cues in phishing emails resulted in higher levels of user response owing to the dominance of peripheral/shallow processing strategies.

This understanding of engagement with incoming emails is likely to be complicated by our finding of substantial variations in workload over time (see section 6.2.3). Participants asserted that the way they process their inbox is qualitatively different depending on how heavy their workload is at the time – possibly explaining some of the variation in the effectiveness of other known predictors of systemic information processing such as desired confidence and motivation. This may also go some way to explaining a much discussed phenomena in the security services industry around phishing emulations where specific emails that have been ‘benchmarked’ according to their effectiveness, for use as calibration tools across organisations, nonetheless exhibit a wide degree of variance in victimisation rates.

In addition, we discovered that participants varied substantially in their reported numbers of non-relevant emails they received each day (see section 6.2.2). This reinforces the relevance of the assertion by Parson et al. [49] that the categories of emails that user

needs to process may have a profound effect on the mental processes involved and leads to a great deal of uncertainty in terms of experimental design. It further complicates Neustaeder’s et al. [48] taxonomy of low, medium and high volume users and suggests another variable that may need to be accounted for to explain processing approaches. At the very least, the proportion of non-relevant/trivial emails is likely to effect the mental efforts devoted to evaluating each email – and if high levels of non-relevant emails results in shallow processing, may result in increased victimisation as users devote less elaborative processing to evaluating the characteristics of each email. This variable should be operationalised and tested in further attentional-based experiments into phishing victimisation.

7.3 Perceptions of threat, vulnerability and responsibility

Vulnerability and threat severity.

We found participants consistently talked about feeling ‘safe’ and ‘protected’ within the information infrastructure of the bank – and tied these feelings firmly to the perceived emphasis and obvious presence of IS protocols, practices and information in their workplace. An anecdote related to us by a bank security worker involved a staff member who noticed a suspicious email they received in their personal email account at home, and had immediately forwarded it to their professional email address in order to open it at work – rationalising this as the safest thing to do since the security environment at work was ‘safer’ than that which they had access to at home.

These low feelings of vulnerability owing to the perceived presence and visibility of IS practices suggests itself as an important finding since it ties in with the literature around risk homeostasis. This theory suggests that in situations of risk, where controls are implemented to mitigate the risk or the severity of the outcomes, people often either decrease their protective behaviours, or increase risky behaviours in order to subconsciously return to the same level of risk as before the mitigation was put in place [50]. This effect has been seen in examples such as anti-locking brakes, where drivers, once aware of the effect of the new braking system on stopping distance, modified their behaviour to drive closer to cars in front of them – returning the risk to subjectively the same levels as before the application of the protective technology - the anti-locking brakes [51]. An emerging design response to this dynamic has been increasing the subjective feelings of risk and vulnerability in order to encourage users to engage in protective behaviours. This has been implemented in domains as far flung as traffic calming designs to aviation systems [52]. This has immediately actionable implications for staff education in that according to all of the models that include the concept of vulnerability, emphasising this variable in educational efforts is likely to increase protective behaviours deployed by staff. On the other hand, whether this mitigation approach is palatable to organisations’ internal communication values is debatable.

When asked about the difference between practices at home and at work (section 6.3.2) we found indications that perceptions of vulnerability were higher at home than at work while threat severity was lower at home than at work. These assertions were specifically linked to the perception that at home breaches would be centred

around personal loss, but at work would also potentially damage the bank as well.

This understanding of threat severity related to the bank itself was also evident in responses to questions in section 6.3.4 where staff placed a heavy emphasis on the consequences to the reputation and trust of the bank by its clients should it experience a major security breach.

Again, models based in the cognitive literature have been shown to explicate such findings as well as offer avenues for mitigation. Boss's deployment of PMT in a study of virus alert warning messages, showed that perceived threat severity predicts fear, which in turn increases (in the described study by double) the perception of the necessity of taking protective measures. On the other hand, this particular variable is contentious as a predictor since Hanus et al. [17] found that that it was not a significant predictor of security behaviour.

This variance in evidence may be explained by a central dynamic of both the PMT and HBT where they predict that threat severity will only increase protective measures when SE is high. I.e.: regardless of how motivated people are to protect themselves, they will not do so if they believe that are not capable of carrying out the necessary actions to protect themselves.

We found that participants often exhibited very little fear about monetary loss in response to cyber attacks. This was evident in a large number of responses where participants detailed being the victim of cyber-fraud, but with the final outcome of their money being replaced by the bank – sometimes very quickly. This notion is also supported by repeated assertions about faith in the bank to replace lost funds should something go awry.

Furthermore, several participants, when discussing their own fraud victimisation, repeatedly used the term reduced 'limit' to describe the outcome of the event and did not seem to perceive that attack as actually involving any monetary loss. This would suggest that people see losses charged to a credit card as qualitatively different and of far less consequence than that of losses to a savings type account.

These two phenomena together may help explain Yu's [35] finding that victimisation experience significantly increased subsequent fear of cybercrime for cyber bullying, digital piracy and computer viruses – but not for on-line scams.

Locus of Control and responsibility

Staff perceptions around responsibility are particularly interesting where while the primary assertion of responsibility was expressed as lying with the individual, it was then immediately qualified by equally strong assertions of a more collective and dispersed responsibility. While the existent HCI literature does not seem to have engaged with the variable of Locus of Control (LOC) directly, it has been shown, within cognitive studies, to be an important predictor of people engaging in protective activities [53].

However LOC is a complicated construct. Walston showed evidence that it is not a unidimensional continuum but rather two independent constructs [54], and since then a number of researchers have attempted to tease out the proposed multi-dimensional space at the nexus of what has been variously called: Self Efficacy, Perceived behavioural control, Locus of Control and Locus of Responsibility [55].

After Rotter's [56] original formulation of the Locus of Control, Levenson [57] extended the model by proposing three subscales: internality, control by powerful others and control by chance. These variables suggest themselves as being particularly apt to this context since staff seem to put much stock in the presence of existing security systems and protocols – implying awareness of the presence of powerful others.

This is however complicated by the fact that in Levenson's [57] formulation, the presence of powerful others are more likely to be considered agentic in the outcome in question, whereas in our context the presence of powerful others, in the form of effective security systems of the bank, suggests an attribution of less likelihood of the reinforcement – i.e.: falling victim to phishing.

There may also be reason to entertain Terpstra's [58] distinction between moral and actionable responsibility. We saw responses where participants discriminated between taking personal responsibility for engaging protective behaviours, but then relying on the bank to provide technical and material assistance. This suggests that there is some perceived distinction made between the roles of the individual and the bank that may correspond to individual actions being seen as a moral responsibility, but organisational responses as more agentic and actionable.

A further argument for the importance of outcome attribution is presented by Jeuring et al. [59], who deployed an additional variable of Locus of Responsibility (LoR) and showed that internal LoR is associated with higher engagement of coping strategies, but only if it is also accompanied by a perception that the person has the necessary resources to mitigate the risk, i.e.: Self-efficacy.

8. LIMITATIONS

This research, being qualitative, resulted in a number of findings that should be considered not as generalised facts, but rather understandings of processes, in a particular context of a particular group of people in a particular industry. While the nature of this knowledge is richer and deeper than that typically resulting from more quantitative approaches, questions of generalisability remain to be addressed by further more quantitative and larger-*n* work as discussed below [60].

Furthermore, since our study took place within a specific socio-technical system, i.e.: a large bank in Australia and New Zealand, it remains to be seen as to how inter-organisational and inter-cultural differences may affect these findings. Specifically cognitions around punishment for maleficence and IS policies generally are likely to vary from institution to institution and ideas around sharing information and identity may vary across cultures.

9. CONCLUSIONS

Our study resulted in a number of findings that suggest both avenues for future research and intriguing hypotheses to test. We present here a summary of our work related back to the original research questions provided at the outset of this paper.

What cognitive variables may be implicated in staff's behaviour in relation to phishing emails?

We found that self-efficacy may well be a strong determinant of staff sharing stories of negative experiences. This is owing to those staff with a self-image of being less technically literate being embarrassed to admit victimisation while those who saw themselves as technically competent felt motivated to share their

stories to prevent victimisation of their peers. In terms of perceptions of threat and vulnerability, we found a noteworthy lack of perceived vulnerability when within the bank's IT systems that were associated with impressions of confidence in the bank's visible and highly estimated security protocols. Low perceptions of vulnerability within bank networks were often accompanied by stories of falling victim to identity theft but where financial loss was quickly mitigated by the bank – leading to a postulated low threat severity attribution specifically for financial victimisation.

How do staff experience information security within the organisation and how does this differ from their perceptions at home?

We found that mitigating IS risks was perceived as a shared responsibility between the individual and the wider bank security systems. Staff conceptions around security breaches were heavily centred around cognitions of subsequent loss of trust in the bank by the public and was seen as an important and central issue for employees. In contrast to the above finding about low feelings of vulnerability within the banks networks, we found different perceptions around on-line experiences at home where participants felt more vulnerable, but where a wide range of perceptions around threat severity was found.

What environmental and organisational factors affect staff behaviour in relation to phishing attacks and information security more generally?

We found that in relation to education, the existing informational channels seemed to be functioning and well received. However there was opportunity to capitalise on staff self-motivation by providing more avenues for active learning and that participants expressed a clear preference for information presented in brief stories centred around personalised experiences and work contexts. In relation to email practices, we found that some staff receive far more emails than others and that there appears to be much variance in the proportion of non-relevant, or trivial emails that staff receive on a day to day basis which has implications for attentional models of information processing in relation to phishing victimisation. Workload was also found to vary significantly over time for some staff, and that this was associated with perceived differences in practices around scanning and responding to incoming messages.

10. FUTURE WORK

The research presented here describes a number of novel observations pertaining to banking staff cognitions around and experiences of IS. While these findings suggest further investigation, we mention three of the more promising avenues for further research below.

Our finding that staff with a self-image of computer competence and being technically 'savvy' are more likely to communicate to others about their negative experiences (section 6.1.4) should be investigated further. Understanding what factors preclude people from discussing and sharing information about phishing victimisation holds promise for creating organisational cultures with increased normative influence on staff about the correct protective behaviours to deploy.

Our finding in relation to variance in the number of non-relevant emails staff encounter in their inbox may have important implications for attentional based and dual process theories of phishing victimisation. This variable should be deployed in future

work employing theories such as the Heuristic Systematic Model to predict elaborative processing of incoming emails.

Perhaps most interestingly, our findings of low feelings of vulnerability associated with visible organisational security protocols suggests an important avenue for staff education efforts. Manipulating vulnerability in messaging and then validating via behavioural responses may increase protective measures as predicted by risk homeostasis theory.

On a more general note, we suggest that modes of investigation that consider and deploy cognitive variables are likely to be of considerable benefit to the HCI and CSCW communities. Specifically attacks based around social engineering require an understanding of the mental processes that result in victimisation, and in the context of phishing, the factors that lead to elaborative processing; i.e.: users actually deploying the knowledge that they have to evaluate threats. We argue that models based on the fundamental mental constructs that drive behaviour are likely to be increasingly useful in combatting the ever-increasing sophistication of on-line threats and hold promise to transform users from a system weakness to an active line of defence.

11. ACKNOWLEDGMENTS

The authors would like to thank the information security staff at the institution discussed for their support, and the participants for their fearless contributions.

12. REFERENCES

- [1] "IBM Security Services 2014 Cyber Security Intelligence Index," 12-May-2015. [Online]. Available: <http://www.ibm.com/developerworks/library/se-cyberindex2014/index.html>. [Accessed: 06-Mar-2017].
- [2] W. Melicher *et al.*, "Usability and security of text passwords on mobile devices," presented at the Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, 2016, pp. 527–539.
- [3] B. Ur, J. Bees, S. M. Segreti, L. Bauer, N. Christin, and L. F. Cranor, "Do Users' Perceptions of Password Security Match Reality?," presented at the Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, 2016, pp. 3748–3760.
- [4] A. Acquisti and J. Grossklags, "Privacy and rationality in individual decision making," *IEEE Secur. Priv.*, vol. 3, no. 1, pp. 26–33, 2005.
- [5] S. Egelman, "My profile is my password, verify me!: the privacy/convenience tradeoff of facebook connect," presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2013, pp. 2369–2378.
- [6] D. Krawczyk, J. Bartlett, M. Kantarcioglu, K. Hamlen, and B. Thuraisingham, "Measuring expertise and bias in cyber security using cognitive and neuroscience approaches," in *2013 IEEE International Conference on Intelligence and Security Informatics*, 2013, pp. 364–367.
- [7] M. Workman, "Wisecrackers: A theory-grounded investigation of phishing and pretext social engineering threats to information security," *J. Am. Soc. Inf. Sci. Technol.*, vol. 59, no. 4, pp. 662–674, Feb. 2008.
- [8] Z. Xu and W. Zhang, "Victimized by Phishing: A Heuristic-Systematic Perspective," *J. Internet Bank. Commer.*, vol. 17, no. 3, pp. 1–16, Jan. 1970.

- [9] X. (Robert) Luo, W. Zhang, S. Burd, and A. Seazzu, "Investigating phishing victimization with the Heuristic-Systematic Model: A theoretical framework and an exploration," *Comput. Secur.*, vol. 38, pp. 28–38, Oct. 2013.
- [10] Y. Sawaya, M. Sharif, N. Christin, A. Kubota, A. Nakarai, and A. Yamada, "Self-Confidence Trumps Knowledge: A Cross-Cultural Study of Security Behavior," presented at the Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, 2017, pp. 2202–2214.
- [11] S. R. Boss, D. F. Galletta, P. B. Lowry, G. D. Moody, and P. Polak, "What Do Users Have to Fear? Using Fear Appeals to Engender Threats and Fear that Motivate Protective Security Behaviors," Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 2607190, Dec. 2015.
- [12] R. Dhamija, J. D. Tygar, and M. Hearst, "Why phishing works," presented at the Proceedings of the SIGCHI conference on Human Factors in computing systems, 2006, pp. 581–590.
- [13] S. Preibusch, K. Krol, and A. R. Beresford, "The privacy economics of voluntary over-disclosure in Web forms," in *The Economics of Information Security and Privacy*, Springer, 2013, pp. 183–209.
- [14] N. Christin, S. Egelman, T. Vidas, and J. Grossklags, "It's all about the Benjamins: An empirical study on incentivizing users to ignore security advice," presented at the International Conference on Financial Cryptography and Data Security, 2011, pp. 16–30.
- [15] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong, "Teaching Johnny not to fall for phish," *ACM Trans. Internet Technol. TOIT*, vol. 10, no. 2, p. 7, 2010.
- [16] J. Nielsen, "User education is not the answer to security problems," *Alertbox Oct.*, 2004.
- [17] B. Hanus and Y. "Andy" Wu, "Impact of Users' Security Awareness on Desktop Security Behavior: A Protection Motivation Theory Perspective," *Inf. Syst. Manag.*, vol. 33, no. 1, pp. 2–16, Jan. 2016.
- [18] M. Fishbein and I. Ajzen, *Predicting and Changing Behavior: The Reasoned Action Approach*. Taylor & Francis, 2011.
- [19] I. Ajzen, "From Intentions to Actions: A Theory of Planned Behavior," in *Action Control*, P. D. J. Kuhl and D. J. Beckmann, Eds. Springer Berlin Heidelberg, 1985, pp. 11–39.
- [20] S. K. Steginga and S. Occhipinti, "The Application of the Heuristic-Systematic Processing Model to Treatment Decision Making about Prostate Cancer," *Med. Decis. Making*, vol. 24, no. 6, pp. 573–583, Nov. 2004.
- [21] P. Ifinedo, "Information systems security policy compliance: An empirical study of the effects of socialisation, influence, and cognition," *Inf. Manage.*, vol. 51, no. 1, pp. 69–79, Jan. 2014.
- [22] A. Vance, M. Siponen, and S. Pahlila, "Motivating IS security compliance: Insights from Habit and Protection Motivation Theory," *Inf. Manage.*, vol. 49, no. 3–4, pp. 190–198, May 2012.
- [23] D. L. Floyd, S. Prentice-Dunn, and R. W. Rogers, "A Meta-Analysis of Research on Protection Motivation Theory," *J. Appl. Soc. Psychol.*, vol. 30, no. 2, pp. 407–429, Feb. 2000.
- [24] J. M. Davis and B. M. Tuttle, "A heuristic-systematic model of end-user information processing when encountering IS exceptions," *Inf. Manage.*, vol. 50, no. 2, pp. 125–133, 2013.
- [25] M. Butavicius, K. Parsons, M. Pattinson, and A. McCormac, "Breaching the Human Firewall: Social engineering in Phishing and Spear-Phishing Emails," *ArXiv160600887 Cs*, May 2016.
- [26] E. M. Trauth, *Qualitative Research in IS: Issues and Trends*. IGI Global, 2001.
- [27] I. Dey, *Qualitative Data Analysis: A User Friendly Guide for Social Scientists*. Routledge, 2003.
- [28] A. Stephanou, "The impact of information security awareness training on information security behaviour," Thesis, 2009.
- [29] R. T. LaPiere, "Attitudes vs. Actions," *Soc. Forces*, vol. 13, no. 2, pp. 230–237, 1934.
- [30] H. Boer and E. Sydel, "Protection Motivation Theory," *Prot. Motiv. Theory*, pp. 95–120, 1996.
- [31] S. Milne, S. Orbell, and P. Sheeran, "Combining motivational and volitional interventions to promote exercise participation: Protection motivation theory and implementation intentions," *Br. J. Health Psychol.*, vol. 7, no. 2, pp. 163–184, May 2002.
- [32] J. Wang, T. Herath, R. Chen, A. Vishwanath, and H. R. Rao, "Research Article Phishing Susceptibility: An Investigation Into the Processing of a Targeted Spear Phishing Email," *IEEE Trans. Prof. Commun.*, vol. 55, no. 4, pp. 345–362, Dec. 2012.
- [33] A. Vishwanath, T. Herath, R. Chen, J. Wang, and H. R. Rao, "Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model," *Decis. Support Syst.*, vol. 51, no. 3, pp. 576–586, Jun. 2011.
- [34] R. Böhme and T. Moore, "How do consumers react to cybercrime?," in *2012 eCrime Researchers Summit*, 2012, pp. 1–12.
- [35] S. Yu, "Fear of Cyber Crime among College Students in the United States: An Exploratory Study," *Int. J. Cyber Criminol.*, vol. 8, no. 1, p. 36, Jan. 2014.
- [36] G. V. Post and A. Kagan, "Evaluating information security tradeoffs: Restricting access can interfere with user tasks," *Comput. Secur.*, vol. 26, no. 3, pp. 229–237, May 2007.
- [37] P. Kumaraguru, Y. Rhee, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge, "Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2007, pp. 905–914.
- [38] G. Mark, S. T. Iqbal, M. Czerwinski, P. Johns, A. Sano, and Y. Lutchyn, "Email Duration, Batching and Self-interruption: Patterns of Email Use on Productivity and Stress," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2016, pp. 1717–1728.
- [39] Q. Hu, Z. Xu, T. Dinev, and H. Ling, "Does Deterrence Work in Reducing Information Security Policy Abuse by Employees?," *Commun ACM*, vol. 54, no. 6, pp. 54–60, Jun. 2011.
- [40] J. S. Downs, M. Holbrook, and L. F. Cranor, "Behavioral Response to Phishing Risk," in *Proceedings of the Anti-phishing Working Groups 2Nd Annual eCrime Researchers Summit*, New York, NY, USA, 2007, pp. 37–44.
- [41] M. Siponen and A. Vance, "Neutralization: New Insights into the Problem of Employee Information Systems Security Policy Violations," *MIS Q.*, vol. 34, no. 3, pp. 487–502, 2010.
- [42] R. Dodge, K. Coronges, and E. Rovira, "Empirical Benefits of Training to Phishing Susceptibility," in *Information Security and Privacy Research*, 2012, pp. 457–464.

- [43] W. Rocha Flores and M. Ekstedt, "Shaping intention to resist social engineering through transformational leadership, information security culture and awareness," *Comput. Secur.*, vol. 59, pp. 26–44, Jun. 2016.
- [44] I. Ajzen, "The theory of planned behaviour: Reactions and reflections," *Psychol. Health*, vol. 26, no. 9, pp. 1113–1127, Sep. 2011.
- [45] M. Harbach, M. Hettig, S. Weber, and M. Smith, "Using personal examples to improve risk communication for security & privacy decisions," presented at the Proceedings of the 32nd annual ACM conference on Human factors in computing systems, 2014, pp. 2647–2656.
- [46] S. Das, A. D. Kramer, L. A. Dabbish, and J. I. Hong, "Increasing security sensitivity with social proof: A large-scale experimental confirmation," presented at the Proceedings of the 2014 ACM SIGSAC conference on computer and communications security, 2014, pp. 739–749.
- [47] A. Gupta, R. Sharda, and R. A. Greve, "You've got email! Does it really matter to process emails now or later?," *Inf. Syst. Front.*, vol. 13, no. 5, pp. 637–653, 2011.
- [48] C. Neustaedter, A. Brush, and M. A. Smith, "Beyond from and received: Exploring the dynamics of email triage," presented at the CHI'05 extended abstracts on Human factors in computing systems, 2005, pp. 1977–1980.
- [49] K. Parsons, A. McCormac, M. Pattinson, M. Butavicius, and C. Jerram, "The design of phishing studies: Challenges for researchers," *Comput. Secur.*, vol. 52, pp. 194–206, Jul. 2015.
- [50] G. J. S. Wilde, "Risk homeostasis theory: an overview," *Inj. Prev.*, vol. 4, no. 2, pp. 89–91, Jun. 1998.
- [51] C. M. Farmer, A. K. Lund, R. E. Trempel, and E. R. Braver, "Fatal crashes of passenger vehicles before and after adding antilock braking systems," *Accid. Anal. Prev.*, vol. 29, no. 6, pp. 745–757, Nov. 1997.
- [52] K. Malnaca, "Risk Homeostasis Theory in Traffic Safety," *21st ICTCT Workshop*, vol. Session IV: A theoretical approach, pp. 1–7, 1990.
- [53] N. Lalwani and T. S. Duval, "The Moderating Effects of Cognitive Appraisal Processes on Self-Attribution of Responsibility," *J. Appl. Soc. Psychol.*, vol. 30, no. 11, pp. 2233–2245, 2000.
- [54] K. A. Wallston, "Assessment of control in health-care settings," in *Stress, personal control and health*, A. Steptoe and A. Appels, Eds. Oxford, England: John Wiley & Sons, 1989, pp. 85–105.
- [55] K. Wallston, "Control Beliefs: Health Perspectives," *MyScienceWork*.
- [56] J. B. Rotter, *Social learning and clinical psychology*, vol. ix. Englewood Cliffs, NJ, US: Prentice-Hall, Inc, 1954.
- [57] H. Levenson, "Reliability and Validity of the I,P, and C Scales - A Multidimensional View of Locus of Control.," Aug. 1973.
- [58] T. Terpstra, "Flood Preparedness Thoughts, feelings and intentions of the Dutch public," University of Twente, 2010.
- [59] J. Jeuring and S. Becken, "Tourists and severe weather – An exploration of the role of 'Locus of Responsibility' in protective behaviour decisions," *Tour. Manag.*, vol. 37, pp. 193–202, Aug. 2013.
- [60] E. Albrechtsen, "A qualitative study of users' view on information security," *Comput. Secur.*, vol. 26, no. 4, pp. 276–289, Jun. 2007.

APPENDIX

A. COMPLETE INSTRUMENT

Topic: Knowledge

Q1 What do you know about cyber security?

Q1-A What do you think it is?

Q1-B Who is involved?

Q1-C Is it important?

Q1-D Why is it important?

Topic: Attitudes

Q2 What is your role at <bank name redacted>?

Q3 How does Cyber Security affect you?

Q4 How do you feel about Cyber Security?

Q4A Do you feel vulnerable to cyber security threats?

Q4B Do you feel fear about Cyber Security?

Topic: History

Q5 What's your experience with Cyber Security historically?

Q5A What kind of stories have you heard?

Q6 Have you ever clicked on something dodgy? What happened?

Q7 Have you clicked on a phishing email? What happened?

Q7A How severe was it?

Q8 Did that make you change your behaviour?

Q9 How did that make you feel?

Q10 Did you tell anyone about it?

Q11 Has it happened again since?

Q12 What do you think is the worst thing that could happen as a result of a phishing attack on you?

Topic: Practices

Q13 How do you manage passwords?

Q14 Do you connect to free Wi-Fi?

Q15 What makes you suspicious of an email? Discuss.

Q16 How do you deal with emails you are suspicious of?

Q17 Is it getting hard to tell what is suspicious?

Q18 Where do you learn or hear about this stuff?

Q19 Whom do you trust for advice or information on Cyber Security?

Q20 Do you follow their advice?

Q21 Do you think there is enough training/information provided at your work?

Q22 Would you like to learn more?

Q23 How would like this training/information to be provided? (prompt: video, podcasts, intranet pages, workshops, induction?)

Q24 Roughly how many emails would you receive in an average work day?

Topic: Contexts

Q25 Roughly how many emails would you send in an average work day?

Q26 How do you feel about your email practices (is it too much, stressful)?

Q27 In your email practice, do you tend to: Check/Notifications/Batch

Q28 How busy do you feel at work? Do you feel you have enough time in your day to devote to each task you need to do?

Q29 How stressed do you feel at work on an average day?

Q30 Are there any consequences at <bank named redacted> for poor security behaviour?

Q31 Colleagues:

Q31A Do you talk about cyber-security issues?

Q31B Have you learnt things from them?

Q31C Do they seem to care about cyber security?

Q32 Bosses:

Q32A Do they talk about cyber-security issues?

Q32B Have you learnt things from them?

Q32C Do they seem to care about cyber security?

Topic: Identity

Q33 Do you see yourself as being good with computers?

Q34 Are you confident with your use of the internet?

Q35 Do you think you can recognise dodgy emails?

Q36 Do you teach or tell other people about Cyber security matters?

Q37 Whose responsibility is it to prevent Cyber Security attacks?

Q38 Is it important to you to be able to recognise dodgy emails?

Topic: Iterated – Round two interviews only

Q201 So does the way you scan your inbox change according to how busy you are? And if so how?

Q202 So what is the difference between thinking about, or cyber security practices at home compared to at work?

Q203 So how much does your workload and the pace of your workplace vary over time?

Q204 So you work for a bank – does that bring with it any special responsibilities and roles?

Q205 Do you know who the cyber security team are the bank? Or how to find them or contact them?

Q206 Do you think the Cyber-security team are good at what they do?

Q207 Do you get a lot of emails that aren't really relevant to you? Or are trivial?

