

What makes people *click*: Assessing individual differences in susceptibility to email fraud

Helen S. Jones

BA (Hons), MSc

A thesis submitted for the degree of
Doctor of Philosophy

Department of Psychology
Lancaster University
August 2016



Declaration

I declare that this thesis is my own work, and has not been submitted in substantially the same form for the award of a higher degree at this institution or elsewhere.

Name: Helen S. Jones

Signature: _____

Date: _____

Publications

The text from Chapter 1 of this thesis overlaps with text from Jones, H. S, Towse, J., and Race, N. (2015). Susceptibility to email fraud: A review of psychological perspectives, data-collection methods, and ethical considerations. *International Journal of Cyber Behaviour, Psychology and Learning*, 5 (3), 13-29.

An edited version of Chapter 3 is currently under revision, following comments from reviewers, for publication in *Applied Cognitive Psychology*: Jones, H. S., Towse, J., Race, N., and Harrison, T. (under review). Email fraud – the search for psychological markers of susceptibility.

Abstract

Cyber security experts have acknowledged that human users are consistently the most vulnerable part of a computer network, however little psychological research has considered *why*. This thesis focuses on susceptibility to email fraud, and highlights three core approaches to understanding why some users are more likely to respond than others, using a mixed methods approach across seven experiments.

The first approach considers the persuasive techniques employed by the sender to make an email more believable. Qualitative data from Studies 1 and 5 demonstrate that authority, familiarity, and the relevance of a communication are important factors when users are considering the legitimacy of an email.

The second approach focuses on the situational factors that may make users more susceptible under specific circumstances. Findings demonstrate that time pressure (Study 3) and a secondary verbal task (Study 6) can impair accuracy in judging email legitimacy.

Finally, individual differences in cognitive make-up between users are considered, with two distinct tasks used to measure susceptibility. Using a forced-choice email legitimacy task (Study 3) and an office simulation, in which participants were naïve to the purpose of the research (Study 7), cognitive reflection, inhibition, and sensation seeking were found to be influential in the decision-making process.

The findings from this thesis outline key influencing factors, which explain some of the variance in individual differences in susceptibility to email fraud. These provide valuable points for consideration in future efforts to educate users on issues surrounding email fraud. Further to this, the development of two lab-based measures of susceptibility, with findings replicated between the two, provides a

platform for further research in understanding and reducing susceptibility. Variations upon the email legitimacy task demonstrate how this can be used to assess effects of a number of manipulations, such as different proportions of phishing and legitimate stimuli (Study 4) and dual-task paradigms (Study 6). The incorporation of additional qualitative data analysis in the thesis, from the use of focus group discussions (Study 1) and think-aloud protocols (Study 5), also provides convergent evidence for the quantitative research findings reported.

Acknowledgements

First of all, I'd like to express my gratitude to John Towse for his excellent guidance and expertise that have helped shape the thesis into what it is now. It goes without saying that this thesis wouldn't have happened without you, so I thank you for the opportunity to work with you, and I hope that we will be able to continue working together in the future. To Nick Race, for your insights from computer science, especially in putting together the publications we have so far from this thesis. It's been really interesting for me to learn more about the technical aspects of cyber security. To Tim Harrison, for providing a constant source of contact for liaison and collaborations with Dstl, and for your help guiding us through the ethics process.

To my Dad, thank you for always having a spare bed at the country retreat for escapes from Lancaster. For the long walks in the Shropshire hills – there are few better places to clear your mind of scam emails than on the top of the Long Mynd with zero phone signal! But more importantly, thank you for supporting me during the PhD process and for your patience in how long it's taken to finish.

To all the people I have met along the way – I came to Lancaster as an undergraduate and I've seen a lot of people come and go in that time. Maybe now it's finally my turn to move on! I've had an amazing time here, but it would have been nothing without the people. Beth, thank you for many a prosecco and fondue-filled evening away from the PhD, for laughter and adventures on conference frolics, from Enschede to Russia, next stop Sydney. To Diana, for the PhD and non-PhD chats on the snow slope and the climbing wall, we'll get back on there soon I promise. To Dr. Watson, for your words of wisdom and encouragement when I'm suffering from

impostor syndrome. To Rob, for the constant stream of Imgur links that make me laugh even from the deepest depths of thesis writing. And to Jack, you came into my life at the most stressful point to date, so thank you for your patience and support while I've been finishing up my PhD – let the adventures begin!

I don't think I realised until I started this PhD how valuable office mates can be to make the good days better and to keep you motivated through the bad days – so thank you to Sophie, for your no nonsense attitude to life, you are an inspiration! Steven, for always being there to talk things through, not only in the office but also on the dance floor in Hustle... (amongst other places)! To James, for brightening the office with your excellent t-shirt collection and constant supply of biscuits. And to Becky, for listening to many rants, both work related and not, for all our much-needed lunch breaks, and of course for the numerous little adventures outside the office. Without you I would never have met Gino D'Acampo, how empty my life would have been! There are too many other amazing people in the department who I could thank, but I think this thesis is long enough! You know who you are, and I will just say that I am so grateful for the supportive and friendly social group that exists here, this would have been a much more difficult process without you all.

Finally, and most importantly, to my Mum. I wish you were here to see what I have achieved; I know you would have liked to read the thesis (or maybe I would just have liked to make you read it...). Thank you for the love and encouragement that got me where I am today, I know you would be proud, and that's all I needed to get me to this point.

Table of contents

| | |
|--|-------------|
| List of Tables | xiii |
| List of Figures | xvi |
| Chapter 1 | 1 |
| 1.1 Introduction..... | 2 |
| 1.1.1 Demographic influences on susceptibility | 5 |
| 1.1.2 The current thesis | 6 |
| 1.2 Theoretical perspectives..... | 8 |
| 1.2.1 Persuasive techniques employed by the sender | 8 |
| 1.2.2 Situational factors affecting cognitive processing | 10 |
| 1.2.3 Cognitive make-up of the user..... | 14 |
| 1.3 Measuring susceptibility | 17 |
| 1.3.1 Scale measures | 18 |
| 1.3.2 Email legitimacy tasks | 19 |
| 1.3.3 Working with past victims | 23 |
| 1.3.4 Simulated phishing attacks | 27 |
| 1.4 Overview of the thesis | 30 |
| Chapter 2 | 33 |
| 2.1 Introduction..... | 34 |
| 2.2 Study one | 37 |
| 2.2.1 Method | 37 |
| 2.2.1.1 Participants. | 37 |
| 2.2.1.3 Materials. | 37 |
| 2.2.1.4 Procedure..... | 39 |
| 2.2.1.5 Data collation. | 40 |
| 2.2.2 Results | 40 |
| 2.2.2.1 Questionnaire data. | 40 |
| 2.2.2.2 Focus group data..... | 44 |
| 2.2.3 Discussion | 51 |
| 2.3 Study two | 55 |
| 2.3.1 Method | 55 |
| 2.3.1.1 Participants. | 55 |
| 2.3.1.3 Materials. | 56 |
| 2.3.1.4 Procedure..... | 56 |
| 2.3.1.5 Data collation. | 57 |

| | |
|---|------------|
| 2.3.2 Results | 57 |
| 2.3.2.1 Descriptive statistics. | 57 |
| 2.3.2.2 Demographic data..... | 58 |
| 2.3.2.3 Context information..... | 58 |
| 2.3.2.4 Reasons for response decision..... | 58 |
| 2.3.3 Discussion | 59 |
| 2.4 Study 2a..... | 61 |
| 2.4.1 Method | 61 |
| 2.4.1.1 Participants. | 61 |
| 2.4.1.3 Materials. | 62 |
| 2.4.1.4 Procedure..... | 63 |
| 2.4.1.5 Data collation. | 63 |
| 2.4.2 Results | 64 |
| 2.4.2.1 Collective analysis. | 64 |
| 2.4.2.2 Responses from open day 1. | 65 |
| 2.4.2.3 Responses from open days 2 + 3..... | 66 |
| 2.4.2.4 Responses from open day 4..... | 67 |
| 2.4.3 Discussion | 67 |
| 2.5 General discussion | 70 |
| 2.4.1 Conclusion | 72 |
| Chapter 3 | 74 |
| 3 Study three | 75 |
| 3.1 Introduction | 75 |
| 3.2 Method | 83 |
| 3.2.1 Participants. | 83 |
| 3.2.2 Design. | 83 |
| 3.2.3 Materials. | 83 |
| 3.2.4 Procedure..... | 89 |
| 3.2.5 Data collation. | 90 |
| 3.3 Results | 93 |
| 3.3.1 Email legitimacy task..... | 93 |
| 3.3.2 Email usage questionnaire. | 97 |
| 3.3.3 Cognitive measures - Set 1..... | 97 |
| 3.3.4 Set 2. | 101 |
| 3.4 Discussion | 103 |
| 3.4.1 Conclusion..... | 109 |
| Chapter 4 | 110 |
| 4 Study four..... | 111 |
| 4.1 Introduction | 111 |

| | |
|--|------------|
| 4.2 Method | 113 |
| 4.2.1 Participants. | 113 |
| 4.2.2 Design. | 114 |
| 4.2.3 Materials. | 114 |
| 4.2.4 Procedure..... | 115 |
| 4.2.5 Data collation. | 117 |
| 4.3 Results | 117 |
| 4.3.1 Email legitimacy task..... | 117 |
| 4.3.2 Email usage questionnaire. | 119 |
| 4.3.3 Cognitive measures – Set 1. | 120 |
| 4.3.4 Set 2. | 120 |
| 4.4 Discussion | 121 |
| 4.4.1 Conclusion..... | 124 |
| Chapter 5 | 126 |
| 5 Study 4a | 127 |
| 5.1 Introduction | 127 |
| 5.2 Method | 129 |
| 5.2.1 Participants. | 129 |
| 5.2.2 Design. | 129 |
| 5.2.3 Materials. | 130 |
| 5.2.3 Procedure..... | 132 |
| 5.2.4 Data collation. | 133 |
| 5.3 Results | 133 |
| 5.3.1 Reliability analysis. | 133 |
| 5.3.2 Moses illusion. | 135 |
| 5.4 Discussion | 136 |
| 5.4.1 Conclusion..... | 138 |
| Chapter 6 | 140 |
| 6 Study five..... | 141 |
| 6.1 Introduction | 141 |
| 6.2 Method | 146 |
| 6.2.1 Participants. | 146 |
| 6.2.3 Materials. | 147 |
| 6.2.4 Procedure..... | 148 |
| 6.2.5 Data collation. | 149 |
| 6.3 Results | 150 |
| 6.3.1 Email legitimacy task..... | 150 |
| 6.3.2 Cognitive tasks. | 152 |
| 6.3.3 Think-aloud analysis..... | 152 |
| 6.4 Discussion | 161 |

| | |
|--|------------|
| 6.4.1 Email legitimacy task..... | 162 |
| 6.4.2 Cognitive measures..... | 163 |
| 6.4.3 Think-aloud responses. | 164 |
| 6.4.4 Conclusion..... | 168 |
| Chapter 7 | 170 |
| 7 Study six..... | 171 |
| 7.1 Introduction | 171 |
| 7.2 Method | 176 |
| 7.2.1 Participants. | 176 |
| 7.2.2 Design. | 177 |
| 7.2.3 Materials. | 177 |
| 7.2.4 Procedure..... | 178 |
| 7.2.5 Data collation. | 180 |
| 7.3 Results | 180 |
| 7.3.1 Email legitimacy task performance. | 180 |
| 7.3.2 Response times on email task. | 182 |
| 7.3.3 Confidence in recognising phishing emails. | 184 |
| 7.3.4 Cognitive reflection test..... | 184 |
| 7.4 Discussion | 185 |
| 7.4.1 Email task performance and cognitive load..... | 185 |
| 7.4.2 Response times and email task performance. | 187 |
| 7.4.3 Understanding of the term 'phishing'. | 188 |
| 7.4.4 Cognitive reflection and email task performance..... | 189 |
| 7.4.5 Conclusions and future directions. | 190 |
| Chapter 8 | 192 |
| 8 Study seven..... | 193 |
| 8.1 Introduction | 193 |
| 8.2 Method | 200 |
| 8.2.1 Participants. | 200 |
| 8.2.2 Design. | 201 |
| 8.2.3 Materials. | 201 |
| 8.2.4 Procedure..... | 204 |
| 8.2.5 Data collation. | 207 |
| 8.3 Results | 208 |
| 8.3.1 Response likelihood to phishing emails. | 208 |
| 8.3.2 Cognitive measures..... | 213 |
| 8.4 Discussion | 215 |
| 8.4.1 Situational manipulations. | 218 |
| 8.4.2 Cognitive variables..... | 221 |
| 8.4.1 Limitations and future research..... | 224 |
| 8.4.2 Conclusions. | 225 |

| | |
|--|------------|
| Chapter 9 | 227 |
| 9.1 Background and main aims of the thesis..... | 228 |
| 9.2 Understanding the issue of email fraud victimisation | 229 |
| 9.3 Measuring susceptibility | 231 |
| 9.4 Assessing the influence of cognitive make-up on susceptibility | 236 |
| 9.5 Assessing situational influences on susceptibility | 241 |
| 9.6 Assessing persuasive techniques employed in phishing emails | 244 |
| 9.7 Major contributions from the thesis..... | 246 |
| 9.7.1 Theoretical contributions..... | 246 |
| 9.7.2 Methodological contributions..... | 249 |
| 9.7.3 Additional contributions. | 251 |
| 9.8 Limitations of the research | 252 |
| 9.9 Future directions..... | 257 |
| 9.9.1 Development of methodology. | 257 |
| 9.9.2 Theoretical implications..... | 259 |
| 9.9.3 Applications of research findings. | 259 |
| 9.10 Conclusions | 260 |
| References | 262 |
| Appendices | 275 |

List of Tables

| | |
|--|-----|
| Table 2.1 <i>Percentage of participants who reported experiencing offline fraud in each context</i> | 41 |
| Table 2.2 <i>Number of hours spent online each day by participants.....</i> | 43 |
| Table 2.3 <i>Percentage of participants who reported experiencing online fraud in each context</i> | 43 |
| Table 2.4 <i>Means and standard deviations for self-control scale and cognitive reflection test scores by response to fraudulent communication</i> | 44 |
| Table 2.5 <i>Framework of user focus group themes.....</i> | 45 |
| Table 2.6 <i>Percentage of participants who reported having received fraudulent communications through each medium in online and offline environments.....</i> | 59 |
| Table 2.7 <i>Frequency of reasons given for not responding to a fraudulent communication</i> | 59 |
| Table 2.8 <i>Descriptive statistics for each open day</i> | 62 |
| Table 2.9 <i>Frequency of reasons given for not responding to a phishing email.....</i> | 65 |
| Table 2.10 <i>Percentage of participants who correctly recognised each email.....</i> | 66 |
| Table 2.11 <i>Percentage of participants who correctly recognised each email.....</i> | 66 |
| Table 2.12 <i>Percentage of participants who correctly recognised each email.....</i> | 67 |
| Table 3.1 <i>Outline of the tasks to be used in this study</i> | 86 |
| Table 3.2 <i>Scoring system for confidence scores on email judgment task</i> | 91 |
| Table 3.3 <i>Descriptive statistics for measure calculated from email legitimacy task ..</i> | 95 |
| Table 3.4 <i>Correlations between measures from the email legitimacy task</i> | 96 |
| Table 3.5 <i>Percentage responses to questions on the email usage questionnaire.....</i> | 98 |
| Table 3.6 <i>Descriptive statistics for cognitive tasks in set 1</i> | 99 |
| Table 3.7 <i>Correlations between cognitive tasks in set 1 and email legitimacy task scores</i> | 100 |

| | |
|--|-----|
| Table 3.8 <i>Summary of multiple regressions analyses for cognitive variables in set 1 predicting email task behaviour</i> | 101 |
| Table 3.9 <i>Descriptive statistics for cognitive tasks in set 2</i> | 101 |
| Table 3.10 <i>Correlations between cognitive tasks in set 2 and email legitimacy task scores</i> | 102 |
| Table 3.11 <i>Summary of multiple regression analyses for cognitive variables in group 2 predicting email task behaviour</i> | 102 |
| Table 4.1 <i>Sets of cognitive measures</i> | 115 |
| Table 4.2 <i>Outline of conditions</i> | 115 |
| Table 4.4 <i>Descriptive statistics for cognitive measures in set 1</i> | 120 |
| Table 4.5 <i>Descriptive statistics for cognitive measures in set 2</i> | 121 |
| Table 6.1 <i>Correlations between performance measures on the email legitimacy task</i> | 151 |
| Table 7.1 <i>Secondary tasks whilst completing the email task</i> | 178 |
| Table 7.2 <i>Means and standard deviations for performance on the email task</i> | 182 |
| Table 8.1 <i>Summary of phishing email stimuli</i> | 203 |
| Table 8.2 <i>Frequency for reasons given by participants who deleted or did not open each email</i> | 208 |
| Table 8.3 <i>Frequency for each action on each phishing email</i> | 209 |
| Table 8.4 <i>Correlations between response types to phishing emails</i> | 209 |
| Table 8.5 <i>Descriptive statistics for cognitive measures</i> | 213 |
| Table 8.6 <i>Correlations between cognitive measures and responses to phishing emails</i> | 214 |
| Table 8.7 <i>Summary of best-fit regression models for cognitive measures predicting response to phishing emails</i> | 214 |
| Table 9.1 <i>Mean accuracy and confidence score for the email legitimacy task across studies 3 to 6</i> | 233 |

| | |
|---|-----|
| Table 9.2 <i>Mean differences in accuracy on the email legitimacy task for each experiment.....</i> | 233 |
| Table 9.3 <i>Summary of best-fit regression models for cognitive measures predicting response accuracy across experiments</i> | 238 |

List of Figures

| | |
|---|-----|
| Figure 2.1 <i>The Cognitive Reflection Test</i> | 38 |
| Figure 2.3 <i>Reasons given for choosing not to respond to fraudulent communications in offline and online environments</i> | 42 |
| Figure 3.1 <i>Diagram to show mean rating across participants for each email stimulus</i> | 94 |
| Figure 5.1. <i>Questions included in the Moses illusion task</i> | 131 |
| Figure 5.2 <i>Number of emails correctly identified at test session 1 and test session 2</i> | 134 |
| Figure 5.3 <i>Confidence score at test session 1 and test session 2</i> | 135 |
| Figure 6.1 <i>Frequency of reference to cue types in email decision-making</i> | 153 |
| Figure 7.1 <i>Finger tapping sequence for complex condition of secondary task</i> | 179 |
| Figure 7.2 <i>Mean response times on email legitimacy task by secondary task condition</i> | 183 |
| Figure 8.1 <i>Graph to show mean number of emails and standard deviation for each response type by time pressure condition</i> | 211 |
| Figure 8.2 <i>Graph to show mean number of emails and standard deviation for each response time by priming condition</i> | 212 |
| Figure 9.1 <i>Diagram of evidence-based theoretical approaches to understanding individual differences in susceptibility to email fraud</i> | 247 |

Chapter 1

General introduction

Chapter summary

This chapter provides a rationale for the research to be reported in the thesis, highlighting previous research from psychology and computer science. Three core theoretical approaches are considered in addressing the issue of individual differences in email fraud susceptibility. These are: the psychology of persuasion, situational influences on cognitive processing, and the cognitive make-up of the user. In addition to the theoretical constructs to be examined in the thesis, methodological approaches to assessing susceptibility are discussed in terms of practical benefits in line with ethical restrictions. Finally, this chapter provides an overview of the studies that will be reported in the current thesis.

1.1 Introduction

The internet provides an ever-expanding, valuable resource for entertainment, communication, and commerce. However, along with this comes the ever more sophisticated threat of cyber attack, with over two and a half million incidents of computer misuse reported in 2015 (covering infection by viruses and account hacking, often a result of response to phishing emails; Office for National Statistics, 2015). Such incidents have obvious implications on a personal and commercial level, as well as within the criminal justice system. However, psychologically, they also offer an intriguing arena for the understanding of the decision-making processes leading to online fraud victimisation. In this chapter, previous research in this area will be discussed from both a theoretical and methodological perspective, followed by an overview of the empirical research that follows in the thesis to address key unanswered questions.

The thesis will provide a psychological analysis of decision-making surrounding email management and phishing emails, focusing on the role of cognitive variables on detection accuracy. A definition of phishing is not straightforward given the multitude of formats that these communications can take. Nonetheless, one broad and useful description is offered by Myers (2007):

“Phishing: A form of social engineering in which an attacker, also known as a phisher, attempts to fraudulently retrieve legitimate users’ confidential or sensitive credentials by mimicking electronic communications from a trustworthy or public organisation in an automated fashion.” (p. 1)

Most phishing emails are sent out to thousands of internet users, with only a small response rate necessary to make it worthwhile (economically) for the attacker. On average, successful phishing attempts have around a five per cent response rate (Norton, 2014). This makes phishing a potentially more sustainable fraud than traditional formats, such as postal and telephone fraud, which incur higher financial and time costs for the fraudster.

Computer science research is continually developing algorithms to detect phishing emails before they reach the user's inbox in both traditional network-based systems (e.g. Fette, Sadeh, & Tomasic, 2007; Bergholz et al., 2010; Islam, & Abawajy, 2013), but also more recently in cloud-based systems which aim to detect and eradicate phishing attacks in the cloud before they even reach the network (Salah, Alcaraz Calero, Zeadally, Al-Mulla, & Alzaabi, 2013). However, a simultaneous increase in sophistication of the emails themselves and the volume being sent means that the benefits of newly developed approaches are often short-lived; advances in the technology developed to protect against phishing attacks are often quickly mirrored in the methods used by the fraudsters to circumvent such detection algorithms. Similarly, efforts to block the phishing websites that emails direct users to, through automated heuristic filters which detect machine learned patterns (e.g. in words used on the webpage - Abu-Nimeh, Nappa, Wang, & Nair, 2007; or in URLs - Garera, Provos, Chew, & Rubin, 2007), or through manual blacklisting, face the same issues with continual technological advancement on the part of the fraudsters in line with that of the researchers. Moreover, encouraging reliance on security software may result in users developing a false sense of security. If they believe (erroneously) that software can reliably capture phish, then they may treat all messages that reach

their inbox undetected, and accessible linked websites, as being genuine. The required sophistication in filtering software also means that false positives occur – with legitimate emails being filtered out of inboxes as spam. An overreliance on these filters means that users are at risk of missing important messages, as highlighted by participants during focus group discussions in Study 1. The inaccuracy in these filtering efforts means that it is left to the user to recognise and manage potential phishing attempts.

It is acknowledged by experts in the field of cyber security that human susceptibility is often the biggest problem though (Mitnick, & Simon, 2002; Schneier, 2000a), with a well-known quote from Schneier to suggest that *“only amateurs attacks machines; professionals target people”* (Schneier, 2000b). He goes so far as to suggest that even if absolute computer security were achieved technologically, the reliance on user interaction to manage the software means the system would still be vulnerable. In relation to email management, biases in human decision-making lead some users to respond to phishing emails whilst thousands of others will receive the same email and ignore it or delete it straight away. This raises the questions of why certain users make these poor decisions but others do not: is this just the luck of the draw, or is there at least some level of systematic group differentiation? For example, there is a common assumption that advancing age leads to increased susceptibility, when in fact the research surrounding this is either inconclusive or suggests a more complex explanation.

1.1.1 Demographic influences on susceptibility

Some research supports the assumption that the average age of fraud victims is significantly higher than the general population (Pak & Shadel, 2011; Shadel & Pak, 2007). However, other research has demonstrated, in contrast, that older internet users are actually *less* susceptible than younger users (Sheng, Holbrook, Kumaraguru, Cranor, & Downs, 2010; Pratt, Holtfreter, & Reisig, 2010). The gender of users has also been highlighted as a potential factor influencing susceptibility, but again the findings are inconsistent, with some research demonstrating that males show more caution in assessing fraudulent communications (Jagatic, Johnson, Jakobsson, & Menczer, 2007; Bailey, Mitchell, & Jensen, 2008; Hong, Kelley, Tembe, Murphy-Hill, & Mayhorn, 2013), whilst other research reports no difference in susceptibility based on gender (Parsons, McCormac, Pattinson, Butacicius, & Jerram, 2013).

Prior web experience and usage habits have been considered in some studies, with a general assumption that users who spend more time online or who have experience with cyber security would be less susceptible to phishing attacks. Jagatic et al. (2007) demonstrated that student participants who were majoring in a technology related subject, such as computer science or informatics, were less vulnerable than students from other subject areas such as business and liberal arts. Friedman, Hurley, Howe, Felten, and Nissenbaum (2002) supported this, with findings that showed participants from a high-technology community were better able to recognise unsecure connections than participants from a typical suburban or rural community. However, Vishwanath, Herath, Chen, Wang, and Rao (2011) report surprising results, which demonstrate that users who receive more email traffic are also more likely to respond to a phishing email. It is thought that persuasive

techniques in some phishing emails, such as a sense of urgency, make them stand out to users against the normal, relevant emails that they are used to receiving, making them more likely to react. Further to this, self-report data regarding email usage habits demonstrated that higher usage increased the likelihood of victimisation in a simulated phishing attack (Vishwanath, 2015).

The inconsistent nature of past research addressing demographic differences in susceptibility to email fraud, suggests that these alone cannot be relied upon to detect the most at-risk users. The National Fraud Authority (2011) produced a report which outlines key demographics for different victim typologies, emphasising that victims come from a range of age groups and backgrounds, suggesting that there is no specific demographic to focus attention on when identifying and addressing real-world incidents of susceptibility. Therefore, alternative explanations for individual differences in susceptibility should be considered, which are addressed later in the chapter.

1.1.2 The current thesis

Given the relatively novelty of the issue of susceptibility to email fraud in research, the literature remains sparse, especially in relation to psychological approaches. As yet, there is no comprehensive psychological model to explain susceptibility through human error. A number of studies report findings regarding individual variables, but have not considered how these interact with one another, as well as with situational influences.

Throughout the thesis, three main approaches will be considered in order to develop such an explanation of individual differences to susceptibility. These are

outlined in Figure 1.1, along with potential examples of each based on past research. Perspectives from social psychology suggest that poor decisions result from persuasive techniques employed by *the sender*, i.e. the perpetrator. On the other hand, cognitive approaches naturally focus on the fallible mental architecture of *the recipient*. This can be considered both in a variable sense, with behaviours that change based on situational, and contextual factors, as well as more concrete individual differences in the cognitive make-up of the user.

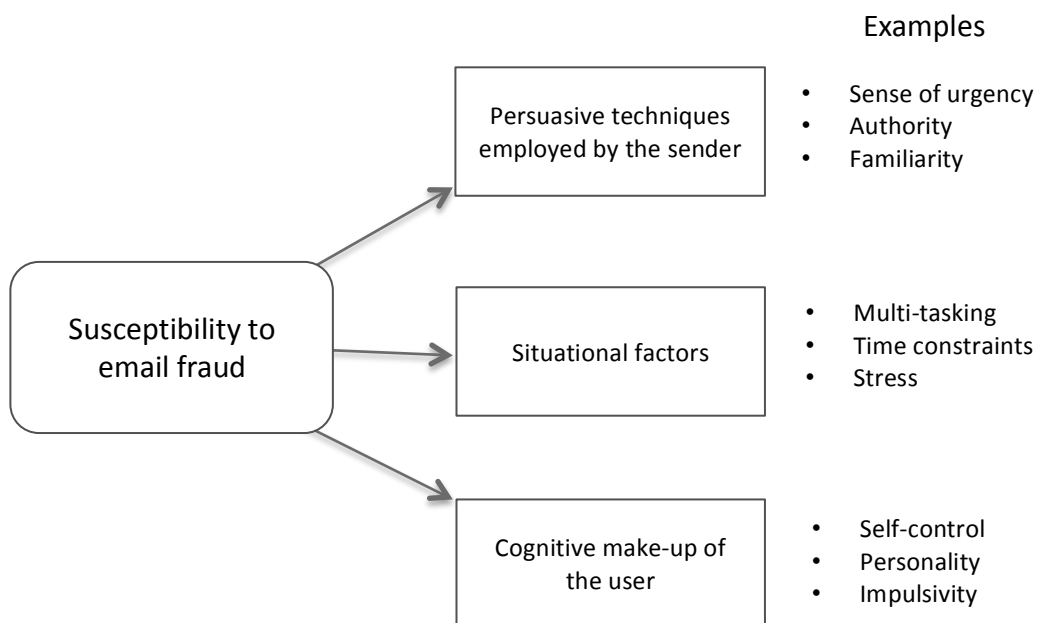


Figure 1.1 *Diagram of theoretical perspectives on susceptibility to email fraud* - based on Figure 1 in Jones, Towse, and Race (2015)

Although this thesis will focus mainly on the psychological approaches, some chapters will also touch on approaches from computer science that have considered human-computer interaction (HCI) and the influences that can increase or decrease users' trust in online systems and communications. These influences include the appearance of the email or website, the perceived quality of information provided,

and the degree of transparency with regard to how information will be used once shared (Karat, Karat, & Brodie, 2009). In order to fully understand the decision-making process, it is crucial to consider the techniques employed on *both* sides of the interaction. There may be emergent properties from the dynamic that exists between message content and message interpreter that are important in the outcome of the email response decision-making process.

1.2 Theoretical perspectives

1.2.1 Persuasive techniques employed by the sender

General theoretical work surrounding the psychology of persuasion points to the relevance of factors such as authority, scarcity, and social conformity (Cialdini, 1993). These persuasive influences can be present in fraudulent communications (as well as legitimate messages) and have been demonstrated to lead to more successful phishing attempts. For example, spear phishing is a technique that uses information collected about the victim, usually from publicly available sources such as social networks, to make the communications more personal, and thus more believable. Work by Jagatic, et al. (2005) used a simulated phishing attack to demonstrate that participants were more likely to respond to an email purporting to be from a friend than an email that came from an unknown sender. This provides evidence of the influence of social conformity – as a user is more likely to conform to an email from a known friend who they may wish to maintain social acceptance from by responding.

Further to this, the level of authority in an email seems to be an important factor in persuading the user to respond. Historically, ground-breaking research such

as that conducted by Stanley Milgram in the 1960s has shown how, regardless of the severity of the consequences, people are generally submissive to the instructions of an authority figure. Workman (2007) demonstrated that self-report measures of obedience to authority were predictive of response behaviour to a number of simulated phishing emails. Further to this, Guéguen and Jacob (2002) used a simulated phishing attack to demonstrate that participants were more likely to respond to an email asking them to complete a survey when the sender was a scientific researcher - labelled by the researchers as the more authoritative figure - in comparison to the same email when it came from an undergraduate student. There may be a number of additional factors (such as history of past messages, awareness of what to expect from the sender, etc.) that are also influencing this decision-making process and variation in response likelihood, but regardless, the effect is noteworthy.

In addition, other visceral influences might be manipulated by fraudsters to emulate scarcity as a way of increasing the persuasive power of an email. Such influences include greed and fear, for example via limited time offers of money, or threat of loss (e.g. access to online accounts, or impending fines) as part of the email message. Higher levels of visceral influence are thought to lead users to overlook the importance of cues that might otherwise trigger suspicion (Langenderfer & Shimp, 2001).

Although these persuasive techniques provide a partial explanation of why some phishing emails are more successful than others, they do not explain how thousands of users can receive the same email and only a small proportion respond.

Therefore, alternative approaches that address situational and individual differences between users are discussed below.

1.2.2 Situational factors affecting cognitive processing

In a real life scenario, users are often required to manage emails whilst pre-occupied with other tasks or under limited time constraints, for example in a work environment where they have deadlines to meet, but 'urgent' emails are being received simultaneously. Therefore, this section considers how individual cognitive capacities and processing might affect decision-making in these scenarios. Such considerations also apply to behaviour surrounding emails that emulate a sense of urgency on the user, through the use of tactics such as threat of loss, thus requiring an immediate response.

In these circumstances users will often experience an impulsive reaction to a phishing email, portrayed by the perpetrator as the rational response, when this is actually a sub-optimal decision (Dong, Clarke, & Jacob, 2008). In order to avoid these poor decisions, users must engage rational decision-making processes in order to suppress their initial intuition. Yan and Gozu (2012) examined susceptibility to email fraud by distinguishing rational and intuitive decision-making conditions in an email legitimacy task. Participants were told to either give rapid responses upon a first look at the email (intuitive), or told to take their time, and read the email carefully before deciding on their final response (rational). In the rational decision-making condition, participants accurately identified more emails as scams than when intuitive decision-making was employed. Further to this, research based on self-reports regarding the use of rational and intuitive decision-making strategies after receiving a simulated

phishing email demonstrated that higher reliance on rational processing predicted lower trust in the legitimacy of the email (Harrison, Vishwanath, & Rao, 2016).

One explanation for these findings could be a dual-systems reasoning approach to decision-making. Such theories of reasoning (e.g. Stanovich, 1999; Kahneman, 2000; Stanovich & West, 2002; and Evans, 2003) propose two psychological systems for generating behavioural responses, the deployment of which depends on the nature of the individual situation. System 1 relies on intuitive, immediate, and emotional responses to make decisions. Sometimes termed the heuristic approach, this approach to reasoning makes use of prior knowledge and experiences to make decisions in unfamiliar situations. This allows decisions to be made more rapidly, with less information processing, and thus is usually the dominant approach taken in decision-making scenarios. On the other hand, system 2 reasoning requires suppressing the intuitive response, in order to gather the necessary information to allow hypothetical thinking. In taking a more rational approach, this allows consideration of the future consequences of a given behaviour, thus allowing a more considered decision to be made. However, this approach is slower, requiring greater attention and information processing, and is therefore not employed as often.

The reliance on generalisations of prior knowledge means that system 1 decisions can lead to systematic errors (Tversky & Kahneman, 1975), as demonstrated in the research outlined above relating to email response behaviour. This is consistent with the idea of such dual-system approaches being employed in relation to email management and susceptibility, suggesting that users are at higher

risk of victimisation when they are reliant on intuitive, system 1 decisions about email legitimacy.

It has been argued that the ability to engage more rational, system 2 decision-making strategies correlates with working memory capacity (Kyllonen & Christal, 1990; Markovits, Doyon, & Simoneau, 2002). Limitations in working memory capacity can affect a wide variety of cognitive situations, such as the ability to maintain multiple interpretations of ambiguous sentences (e.g. Miyake, Just, & Carpenter, 1994). By analogy, such limitations may make it more difficult for a person with lower working memory capacity to engage rational decision-making strategies in detecting phishing emails. Processing the content of an email whilst assessing its veracity may be cognitively taxing, especially in situations where time is limited or the user is completing multiple tasks simultaneously. As a result, users in these situations, especially those with lower working memory capacity, may fall back on impulsive decision-making strategies.

An alternative explanation of those with higher working memory capacity performing better under cognitive load may be because they are better able to efficiently divide attention. With the development of modern technology, users are often engaged with multiple devices simultaneously, working on numerous tasks, thus requiring continual shifts in attention between these (Roda, 2011). Efforts to divide attention normally result in effective distribution of the cognitive processing capabilities required to complete tasks simultaneously. However, under certain circumstances, dividing attention can result in erroneous decisions. Theories of divided attention suggest that individual differences, in factors such as working memory capacity, as well as situational factors can influence how well a person is

able to perform multiple tasks at once (Kane, Bleckley, Conway & Engle, 2001). Colflesh and Conway (2007) demonstrated that participants with a higher working memory capacity performed better on a selective attention, dichotic listening task (participants must concentrate on a task based on an auditory stimulus, whilst simultaneously listening to another different auditory stimulus) than those with a lower working memory capacity. In relation to online decision-making, it may be that a person's ability to detect phishing emails is, in part, influenced by any other tasks requiring their attention at the same time, combined with their ability to effectively divide attention.

In an experimental context, divided attention is often measured using dual-task paradigms, in which participants are given a primary task to complete, whilst simultaneously completing a secondary task that adds to their cognitive load, such as counting out loud, or remembering a letter string. More recent literature has used these types of tasks in applied setting, for example, looking at participants' attention to driving when using a mobile phone. Beede and Kass (2006) gave participants a simulated driving task, whilst also engaging in a conversation on a hands-free mobile device and found that driving performance was negatively affected by simultaneous use of a mobile device. The dual-task paradigm may be transferred to situations involving email management, whereby a user is often attempting to interleave between multiple tasks. This would provide evidence to establish whether dividing attention, and an inability to do so effectively, is detrimental to the decision-making process.

Aside from the situational influences on cognitive processing, individual differences in a construct such as working memory capacity may affect the

propensity of individuals to process emails sub-optimally. That is, working memory demands can vary by situation for a person, but working memory capacity differs across individuals too (Conway, Jarrold, Kane, Miyake, & Towse, 2007). In this regard, Cokely and Kelley (2009) found that participants who demonstrated a higher working memory capacity were less likely to engage in risk taking behaviour, which is likely to be because they were able to engage more rational decision-making strategies than those participants with lower working memory spans.

1.2.3 Cognitive make-up of the user

In addition to working memory capacity, other psychological variables can be considered in terms of how variability between users acts as an influencing factor on susceptibility. A combination of relevant psychological factors may contribute to a sort of 'cognitive profile', representative of those users who are most at risk to online fraud victimisation. Research has shown links between working memory capacity and inhibition (Engle, 1996; Redick, Heitz, & Engle, 2007), demonstrating that lower working memory capacity relates to impaired performance in tasks measuring inhibition, such as the Flanker task (Redick & Engle, 2006). Inhibition describes a cognitive function requiring suppression of surrounding information to allow a person to successfully complete the task in question. In relation to email management, increased inhibitory capacity would allow a user to suppress their intuitive response to a phishing email, in order to contemplate all cues available to them and make a more informed decision. This relates back to the dual-systems theories of reasoning described above, with inhibitory capacity improving ability to

engage more rational decision-making processes, as demonstrated with working memory capacity.

A further variable, cognitive reflection (Frederick, 2005), describes ability to reflect upon a problem in order to engage rational decision-making and reach a more accurate, reasoned decision, rather than relying upon an intuitive response. There are clear parallels between this and inhibitory capacity in explaining behavioural responses during decision-making. When considered in relation to email management, users with higher levels of cognitive reflection may be better at recognising more subtle cues to deception through rational contemplation, and thus reduce their susceptibility. However, some evidence has demonstrated a link between cognitive reflection and risk preferences – with participants who have higher levels of cognitive reflection being more likely to take a calculated risk in a gambling task in order to achieve a larger pay out at a later date, rather than taking the immediate, smaller pay out. This may mean that those users who have higher levels of cognitive reflection would be more likely to rationalise the risk involved with responding to a phishing email, thus *increasing* their response likelihood.

Other research looking at decisions surrounding risk using gambling / risk pay-off scenarios has directly considered the relevance of dual-system theories of reasoning in this scenario (e.g. Brand, Heinze, Labudda, & Markowitsch, 2008; Porcelli & Delgado, 2009). There are strong parallels between these types of gambling tasks and fraud detection. The prevalent methodologies all rely on the same basic principle of exploring whether people would take a gamble or not. In employing intuitive decision-making processes, a person would likely choose the option which gave them the biggest reward, rather than thinking through the

decision in terms of the long term benefits or risks. If the same processes are in fact employed in email decision-making, then it is likely that psychological variables found to influence risky decision-making in these gambling scenarios will also influence decisions surrounding email management (e.g. sensation seeking: Hoyle, Stephenson, Palmgreen, Lorch, & Donohew, 2002).

Previous research looking specifically at fraud victimisation has considered personality and self-control as indicators of susceptibility. Modic and Lea (2011) demonstrated that higher levels of agreeableness and lower levels of extraversion, measured using the International Personality Item Pool (IPIP; Goldberg, 1999), were associated with higher levels of susceptibility to fraud victimisation, as measured by self-reports of past behaviour. They suggest that users who show higher levels of agreeableness are likely to be more trusting generally and so believe what they are told by those they are communicating with online, whilst those who show lower levels of extraversion are likely to seek and build stronger relationships online than they are in person as they are not as comfortable in offline social situations, which again may lead them to be more trusting of people who they are interacting with online.

In an adaptation of Gottfredson and Hirschi's self-control theory (1990), Schreck (1999) suggests that low self-control is a strong predictor of crime victimisation. Although this theory relates more broadly to all types of victimisation, elements of Schreck's explanation of this theory relate clearly to fraud. In parallel with theories of dual-system reasoning (discussed above), it is suggested that victims engage in intuitive, rather than rational, decision-making processes, with no consideration for the negative consequences of their actions. Instead the decision is

based on the proximity of gain or loss. Further, Schreck proposes that those with lower self-control demonstrate less diligence in terms of security related behaviour, thus leaving themselves more at risk. This theory has been applied to fraud victimisation - with lower levels of self-control found to relate to fraud susceptibility based on response to a set of hypothetical written scenarios (Holtfreter, Reisig, Piquero, & Piquero, 2010). In addition, research looking at victims of internet consumer fraud (i.e. paying for an item online that never arrived), based on self-reports of victimisation, demonstrated that participants with lower self-control were more likely to have been victimised (van Wilsem, 2013). This study focuses on one specific online fraud scenario though, and so there is no evidence as yet that this can be generalised to responses to phishing emails.

The research highlighted in this section demonstrates a number of potential links between cognitive variables and the decision-making process surrounding fraud victimisation. In the thesis, these cognitive variables, including those from broader research such as that on risky decision-making and working memory capacity, will be assessed in terms of their relationship with susceptibility.

1.3 Measuring susceptibility

Whilst the psychological constructs considered above in relation to susceptibility, such as personality, self-control, and working memory, can be assessed using reliable, well-specified measures, and the persuasive techniques manipulated through content shown to participants, the methods used so far to assess susceptibility are less clear-cut. The following discussion outlines four key methodologies that have been used to date in research surrounding email decision-

making, each of which has practical benefits. However, none have been extensively replicated to assess reliability or validity. Therefore, the benefits and limitations of each will be considered in contemplation of assessing susceptibility in the thesis.

1.3.1 Scale measures

The Human Aspects of Information Security Questionnaire (HAIS-Q; Parsons, McCormac, Butavicius, Pattinson, & Jerram, 2014) measures the extent to which individuals are at risk from security threats, based on the interactions between their knowledge of policy and procedures, attitude towards these, and computer behaviour. Although designed to cover a broad spectrum of security threats such as password management and social networking use, this questionnaire does include questions about email behaviours, including distributing emails and opening attachments. The scale was reported to have good internal consistency and test-retest reliability (McCormac, et al., 2016), although no measures have yet been reported on the ecological validity in relation to how well the questionnaire measures actual information security behaviour.

Modic and Anderson (2014) have recently developed a self-report scale measure of susceptibility to persuasion, which is linked more directly to behaviour surrounding email management. This incorporates psychological mechanisms from a range of contexts, with a focus on measuring factors that influence scam compliance, such as social influence, sensation seeking, self-control, and risk preferences. Example items from the scale include: *'It is important to me that those who know me can predict what I will do'*, *'I have a hard time breaking bad habits'*, and *'In general, I work better when I'm under pressure'*. Validity testing on an earlier version of the

scale (Modic & Lea, 2013) found the factors measured in the scale, such as self-control and sensation seeking, were all related to susceptibility as measured by self-reported compliance to a set of written scenarios. However, since developing the second version of the scale, and including a number of added variables, this validity testing does not seem to have been repeated so we cannot know for sure the extent to which this updated scale is actually measuring susceptibility.

Both of these measures were developed after the current thesis was designed, and so were not incorporated to assess susceptibility. With further development and validity testing, they do offer an ethically sound alternative methodology for future research though. However, as with all self-report measures, these scales may be subject to demand characteristics, i.e. the participant responds to the scale in a way that they think is socially desirable rather than responding truthfully (Nederhof, 1985; Paulhus, 1991; King & Bruner, 2000). Therefore, validity testing would need to ensure that the scale measures were predictive of actual behavioural response to email attacks, rather than self-reported victimisation, before these could be considered as an alternative to behavioural measures, such as those described below.

1.3.2 Email legitimacy tasks

A more commonly used measure of susceptibility involves asking participants to rate their likelihood to respond in given situations. This method has been used in research looking at both online and offline fraud (e.g. Holtfreter et al., 2010), but this section will focus on those studies that have investigated email fraud.

Yan and Gozu (2012) used this method with email screenshots, showing participants either the subject line of an email or the entire body of text, as a measure of the importance of availability of information. Overall, there were 36 emails in this task, all of which were genuine examples of unsolicited phishing emails. Participants were asked to report whether they would 'read' or 'delete' each email. Results demonstrated that participants reported they would 'delete' significantly more emails when they were shown the entire body of text, as opposed to when they saw only the subject line of an email, demonstrating the importance of considering all available information when managing emails. Such tasks provide an arguably more valid measure of susceptibility than the questionnaire approach, as they measure actual behaviour, although this behaviour is still laboratory based. Participants may not be interacting with the emails in the same way as they would were they to receive the given emails to their own inbox. In addition, Yan and Gozu's task uses only phishing emails as stimuli, which may influence participants' responses. To the extent that participants expect to differentiate between a mixture of phishing and legitimate emails when they are given a forced-choice task, decisions may show an expectancy bias. Further to this, the binary response choice of 'read' or 'delete' does not necessarily reflect the range of attitudes or beliefs a user could have concerning an email. Choosing to read an email does not necessarily demonstrate susceptibility as the user might then disregard the email. Adaptations of this task to include legitimate emails and more response options may provide a more accurate measure of email management behaviour, as participants would be discriminating phishing from legitimate emails, as in their own inbox.

One additional limitation of a forced-choice email task such as that employed by Yan and Gozu (2012) is that participants are actively seeking to distinguish between phishing and legitimate emails, which they are not doing in day-to-day email management. Parsons et al. (2013) reported that participants performed better in an email judgment task when they were forewarned about the nature of the research. This suggests that participants may perform with higher accuracy when actively seeking to recognise phishing emails.

Variations on this type of methodology have attempted to address this issue through the use of role-play scenarios, in which participants are asked to access the account of a fictional character and report how they would deal with a number of emails in the inbox of this account. Downs, Holbrook, and Cranor (2007) employed this method when assessing how knowledge of cues, such as security icons, affected phishing susceptibility. Participants, in the role of the fictional 'Pat Jones' were asked how they would respond to a set of five emails, each of which contained a URL link - with no mention that the study looked at ability to detect phishing emails. Possible response options for each email included, *'reply by email'*, *'click on the link'*, and *'type the URL into a browser window'*. Those who indicated that they would click on the link were then shown the associated webpage and asked how they would respond faced with this. Participants who could correctly define 'phishing' and recognised incorrect security lock images showed lower susceptibility to phishing, whilst knowledge of other risks such as spyware and viruses did not affect susceptibility. This type of task is informative in the sense that it can assess susceptibility in a controlled environment whilst not alerting participants to the nature of the task, thus reducing expectancy effects (Parsons et al., 2013). However,

the way in which this specific task was constructed may still have prompted socially desirable responses. For example when given the option '*type the URL into a browser window*', this may alert participants that this is the most sensible option compared to the other options such as '*click on the link*'.

Hong et al. (2013) employed a similar methodology, with participants taking on the role of 'Bob Jones' and being asked to categorise 14 emails messages, of which 9 were illegitimate. Two response options were available – marking the email as important, or deleting the email. This study identified trust, extraversion, and openness as factors correlated with likelihood to delete legitimate emails. However, no significant findings were reported in relation to psychological constructs and responses to the phishing email stimuli. In this study, the majority of the email stimuli were designed to be illegitimate, which is not representative of a real inbox and may generate some bias in how participants respond. In addition, the binary nature of the response options generates the same problem as Yan and Gozu's study. The categories are very broad and marking an email as 'important' does not necessarily demonstrate susceptibility.

These email legitimacy tasks may also be limited in the emulation of threat to the participants, as there is nothing to risk in their participation. If participants were to judge these emails in their own inboxes then it would be their personal information or money at stake if they chose to reply, whereas in the lab situation, participants have nothing to lose whether they perform well or not. In order to encourage participants to perform realistically, it would be necessary to provide them with some incentive to perform well. This may encourage participants to put more effort into the task, although the level of risk still does not mirror that faced

when managing phishing emails in real life. As yet, there is no evidence to suggest how much this affects study validity.

Although there are still some unanswered questions in relation to how ecologically valid email legitimacy tasks are as a measure of susceptibility, these do provide an ethically sound method for assessing how well users can distinguish between a set of phishing and legitimate emails. This design allows for assessment of a variety of email types, as well as providing the controlled environment necessary for manipulating other variables of interest. Examples include time pressure and completion of simultaneous tasks, which allow for the assessment of situational predictors of performance on the email task. To date, research has not considered how this type of task relates to real world susceptibility, but if a strong correlation were demonstrated between accuracy on email legitimacy tasks and susceptibility on more ecologically valid measures of susceptibility, this would imply that these tasks could act as a complementary alternative to more ethically restricted and time-consuming methodologies.

1.3.3 Working with past victims

Modic and Lea (2011) used an alternative to the judgment tasks above, by asking participants whether they had ever responded to each of a set of fraudulent scenarios outlined in written descriptions. Data from those participants who reported responding to a fraudulent communication in the past were used in further analysis. Of 506 participants that were screened, only 67 claimed to have responded, meaning that the sample size used in the remainder of the analysis was drastically reduced. In this study, reports of victimisation were limited to the written scenarios

given to participants. Whilst these cover a range of scam types, the list is not comprehensive, and interpretation of the scenarios described may differ between participants. Therefore, the scenarios may underestimate victimisation, with some participants being victims of fraudulent communications that do not fit into any of the scenarios described.

Whitty and Buchanan (2012) recruited victims of romance scams in order to compare performance measures of loneliness, extraversion, agreeableness, neuroticism, romantic beliefs, and sensation seeking, between them and a control group to establish differences. The only significant finding in this study was that victims were more inclined to idealise romantic partners. This work looks solely at victims of online romance scams though, which is a specific focus and it cannot be assumed that findings would generalise to victims of other types of scams online.

Shadel and Pak (2007) also worked with past victims, looking at differences in demographic, and psychological characteristics - such as impulsivity, self-reliance, and optimism - between offline fraud victims and a control group. These victims were grouped depending on the type of fraud perpetrated, in this case either lottery, or investment victims. Some differences were found between the victim and the control groups on the psychological measures. For example, lottery victims were found to demonstrate higher impulsivity than both investment victims and the control group, whilst the victim group as a whole demonstrated more self-reliance than the control group. However, each of the variables discussed in this study was only measured by one question, meaning that its accuracy in assessing psychological variables may be fairly limited.

On one hand, past victims form a sample population who have self-evidently demonstrated susceptibility to online fraud attempts in a real-world setting. However, there are still a number of challenges faced in using such a sample for this kind of research. One of these is that fraud victimisation is a quasi-experimental variable; such individuals have not been assigned at random to a 'victim' group, meaning that the contextual circumstances surrounding their victimisation need to be established. Being a past event outside of experimental control, this can be difficult as gathering this information is reliant upon the victim's own recall of the event. The likelihood of a victim being able to recall the situation they were in or external constraints on them upon receiving the email and as a consequence of reading it – were they tired, distracted, busy – is minimal. The reliance on self-reported information about situational factors affecting susceptibility means that validity is limited. It is also possible that the individual who has been a victim will have changed as a result of the incident in terms of their cognitive make-up; after all, they may have been financially ruined, or they may be embarrassed by what now appears to them to be a demonstration of gullibility. They may also have read about the scam or other similar scams, and may have been part of training programmes with respect to online security. As a result, the responses they give to questions and psychological measures may differ from the way they would have responded prior to victimisation.

In addition, this method relies on the researcher's ability to establish a comparable control group so that differences can be measured between the susceptible group and a group of users who do not demonstrate susceptibility. This may prove challenging for a number of reasons – the most significant being

establishing who is not susceptible. It may be the case that a user has never responded to a fraudulent email communication, but this may be because they have not received a sufficiently convincing phishing email. Matching control group participants and past victims on their demographic characteristics (age, gender, educational background), internet experience, and email usage habits would go some way to ensure a comparable sample. However, it is not realistic to establish the exact emails which participants have received in their time as internet users, so is not possible to know whether members of the control group are less susceptible or whether they have simply not received the same phishing emails as victims.

Finally, it is important to consider how the method of recruitment for work with past victims may influence the validity of the data. Cybercrime goes heavily underreported, with substantially lower reporting rates than other crimes (Copes, Kerley, Mason, & Van Wyk, 2001), so there may be some bias in the sample of victims who are willing to participate in research relating to their experiences. Some may not report their victimisation due to embarrassment or lack of belief that it will help in any way, so it is unlikely that these people would be willing to discuss their experiences with a researcher either. There are also ethical considerations in working with those victims who have reported their victimisation, as this is a sensitive and traumatic experience for some. Reliving the experience may be difficult and as a result, a sample of past victims who choose to volunteer to take part in research on the topic may not be representative of all victims – as those who have been affected more severely may be less willing to participate.

1.3.4 Simulated phishing attacks

The most ecologically valid, and yet ethically restricted method for assessing susceptibility is to simulate a genuine phishing attack by sending a fake phishing email to participants and recording whether or not they respond. Such an approach addresses a number of the limitations outlined with other methods of measuring susceptibility discussed above. A number of studies have used this methodology to measure real world susceptibility. Wright and Marett (2010) measured the importance of different behavioural factors to recipients' likelihood to respond to a simulated phishing email. In this study, participants signed up to take part in research with the generic title 'security research' and were then given a unique ID code, which they were told was to be used to access course materials, assessments, and grades. These students then studied a module on internet security and privacy as part of their course, and completed a number of questionnaires assessing web experience and risk taking behaviour. After this, the simulated phishing attack occurred, asking students to disclose their unique ID code due to lost information within the information technology database. Of the 299 participants included in the final analyses, 32% of these responded to the email and it was reported that success of the phishing attack was related to less web experience, less security knowledge, less suspicion of humanity, and lower computer self-efficacy.

Although participants in this study were not told explicitly about the phishing attack that would occur, they gave consent to take part in 'security research'. Therefore, they may have been more suspicious of an email coming into their inbox, knowing that it may be part of the research. Preliminary evidence from an undergraduate student project at Lancaster University demonstrated that when

participants were forewarned that they might receive a phishing email as part of the experiment, their response likelihood to a simulated attack was lower (Mack, 2014). In addition, the completion of a module on internet security and privacy means that participants in Wright and Marett's study would have been more aware of security risks, having just studied them. The role of information about the existence of such threats is not simple though, with some evidence to suggest that domain-specific experience and knowledge of cues to legitimacy does not always reduce susceptibility (Downs, Holbrook, & Cranor, 2006; Vishwanath et al., 2011)

Guéguen and Jacob (2002) also used a simulated attack, to measure the influence of the authority level of the sender on likelihood to respond. An email containing a HTML form for a survey on food habits was sent to two samples – students at a university, and users who were listed on internet software designed to procure lists of email addresses. Emails that purported to come from a scientific researcher (deemed to be of higher authority) gained a higher response rate than those purporting to come from an undergraduate student.

A further example comes from Jagatic et al. (2005), who used a simulated attack on University students in their research assessing the influence of familiarity on response likelihood. They found a significantly higher response rate to the emails that purported to come from senders familiar to the participant. For this study, Jagatic et al. used social networking information publicly available on the internet about participants to generate the 'familiar' senders. Although this information is freely available online, some participants may have been uncomfortable with this being gathered and used for research purposes. Finn and Jakobsson (2007) report that this study received 30 complaints from over 1700 participants, and 7 requested

their data be removed. Although these numbers seem relatively low, given the large sample size, upset from any number of participants should be taken into account in the consideration of developing measures of susceptibility to be used in the future. It is also worth noting that there may have been other participants who were also unhappy or distressed following the study, but did not raise a formal complaint with the researchers.

Although the use of simulated phishing attacks allows assessment of real-world susceptibility, a number of limitations may restrict the interpretation of results. The content and the context of the emails is very important to response likelihood, as discussed previously with regard to persuasive techniques employed by senders and situational factors that affect decision-making processes. Whilst the content of the simulated emails can be manipulated, the context in which it is received cannot. This means that, similar to working with past victims, any data about situational predictors of susceptibility is reliant on self-report measures. Further to this, the effect of different persuasive techniques employed in phishing emails, as discussed above, may differ between individuals. Although emails can be designed to emulate specific persuasive techniques, delivering only one email to a participant cannot capture all aspects of susceptibility. Rather it focuses on users who are susceptible to the specific techniques used in that example. The only way to deal with this limitation would be to send multiple emails using different techniques to the same users, but from a practical viewpoint this is not effective, as participants would undoubtedly become suspicious and be less likely to respond upon receipt of numerous simulated emails.

Finally, the ethical constraints of a simulated phishing attack must be considered. In sending out simulated phishing attacks, researchers must be sensitive to the embarrassment and upset that may be caused by a participant's decision to respond, in the same way as working with past victims. Further to this, the deception required to obtain personal information to conduct an attack without consent may be considered as an invasion of privacy. In summary, although it is the most ecologically valid measure, providing a real-world assessment of susceptibility, a simulated phishing attack compromises the experimental control gained in a lab-based environment. Therefore, the most appropriate methodology may depend upon the factors being assessed in a specific experiment, dependent on the level of control and experimental manipulation required.

1.4 Overview of the thesis

Through a number of experiments, this thesis aims to develop an understanding of individual differences in susceptibility to online fraud victimisation. Building on previous research, the three theoretical strands outlined in Figure 1.1 will be considered both independently and together in order to understand how these interact to influence susceptibility. The methodological considerations outlined above are also taken into account throughout, with a mixed methods approach to measuring susceptibility across the different experiments.

Study 1 describes a number of focus group discussions, aimed to gather qualitative data on awareness of email fraud, and of persuasive techniques employed by the fraudsters. Studies 2 and 2a provide exploratory data about user awareness and experience of email fraud. Further to this, Study 2a introduces a small

set of email stimuli, with participants asked to judge the legitimacy of each. This provided some initial data on the stimuli to be used later in the thesis.

Study 3 introduces these stimuli in a more extensive email legitimacy task, as a measure of susceptibility, and assesses this in relation to a battery of cognitive tasks in order to begin building a profile of the most at-risk users. This study also considered how an induced time pressure on the email task affects the decision-making process. In Study 4, an adaptation of the email legitimacy task assesses how varying the proportion of phishing and legitimate emails in the stimuli set (with majority legitimate emails being more representative of a genuine inbox) affects performance, and the predictive nature of the cognitive variables outlined in Study 3. Study 4a assesses the reliability of this varied proportions version of the email legitimacy task, as this was a novel task developed for the purpose of this thesis and therefore has not been previously tested. Study 5 employs a think-aloud protocol to gain more insights into the cues used by participants, and the persuasive techniques they are aware of, in identifying phishing emails during the email legitimacy task. This also allows for qualitative evidence to be gathered demonstrating the presence of the cognitive predictors outlined in Study 3. Study 6 also uses the email legitimacy task, but includes secondary tasks to be completed simultaneously. In emulating a situation where a user has multiple tasks to complete at once, this study demonstrates how increased cognitive load can affect the decision-making process when assessing email legitimacy.

Finally, Study 7 outlines an alternative methodology for assessing susceptibility. In this study an office simulation was designed, whereby participants were naïve to the nature of the study, and completed a number of office based tasks,

including email management. A number of emails were sent over the duration of the simulation, some designed to emulate phishing emails, and response likelihood to these was recorded. The cognitive variables outlined as predictive in Study 3 were also included in this experiment, to assess whether findings were replicated with a more ecologically valid measure of susceptibility. Chapter 9 provides an overview of the findings from these experiments, as well as a comprehensive discussion of the theoretical and practical implications of these.

Chapter 2

Exploring personal experience and understanding of fraud

Chapter summary

This chapter reports findings from three initial exploratory experiments, which aimed to elicit insight on participants' personal experiences of fraud. Study one describes qualitative data from a number of focus group discussions, which highlight a limited understanding of more sophisticated phishing attacks and a tendency to rely on out-dated cues in spotting fraud. Study two provides further insight into the cues employed in recognising email fraud, with data from two distinct age groups – prospective students and their parents at University open days. Study 2a reports further data about personal experience of fraud, and also pilots a set of email stimuli for use later in the thesis. The exploratory data collected in these studies provide grounding for further research to understand why there are apparent differences in the way users respond to phishing emails, and how this information might be utilised to reduce susceptibility.

2.1 Introduction

Shocking statistics about victimisation and monetary loss are commonplace in adverts and warning notices regarding online fraud; however there is little research that considers the personal experiences of every day users or how effective media warnings are. This chapter reports two initial studies about just that – asking users about their experience of fraud in both online and offline scenarios, as well as more in depth information about the cues that they use to detect deceptive communications. In order to understand how best to tackle the issue of online fraud susceptibility, it is important to understand how users manage their emails on a daily basis and how familiar they are with the issues they face regarding online security. In a relatively new area of research interest, exploratory analysis such as that described in this chapter provides a valuable insight into how the topic should best be approached in a way that can directly help the user and work towards reducing susceptibility.

In Study 1, a short questionnaire was constructed, incorporating both multiple choice and open-ended questions, to gain information about whether participants had experienced fraud in both online and offline environments, and about the cues that allowed them to recognise (or not recognise) the fraudulent nature of the communication. In addition, this questionnaire included measures of self-control and cognitive reflection to examine the relationship between these and participants' history of responding to fraudulent communications. These variables will be explored in more detail later in the thesis, but were included here to provide an insight into their relationship with self-reported victimisation. Holtfreter et al. (2010) reported lower levels of self-control being predictive of higher susceptibility

to a telemarketing fraud scenario. It was suggested this might be because users with lower self-control are less likely to consider the negative consequences that may result from their choice to respond to the fraudulent communication. Cognitive reflection (Frederick, 2005) was also assessed here, as a measure of a participant's ability to rationally reflect upon a problem rather than relying on an intuitive response. Based on previous research and the nature of these variables, the following hypotheses were generated:

H¹: Participants who demonstrate lower levels of self-control are more likely to report having previously responded to a fraudulent communication.

H²: Participants who demonstrate lower cognitive reflection are more likely to report having previously responded to a fraudulent communication.

The main part of this first study though was a focus group discussion that followed the questionnaire. As outlined by Onwuegbuzie, Dickinson, Leech, and Zoran (2009), focus groups are a valuable method for collecting rich data about users' experiences and opinions. The group scenario provides a less threatening environment for participants to share and interact with one another, which will hopefully encourage discussion of personal experiences with fraud. In addition, the focus group discussion will elicit information about the tactics participants use to protect themselves, including the cues used to recognise a fraudulent communication.

Study 2 consisted of a short questionnaire completed by attendees at a number of University open days, following a research talk about the current project. The open days provide an opportunity for prospective students and their parents to

visit the campus and hear about the psychology department at Lancaster. This means that a different demographic to the usual undergraduate student sample was accessible, as the parents were also able to take part. Given the relatively recent growth of the internet as a platform for communication, entertainment, and knowledge, this alternative demographic is an interesting one as these users have learnt to use the internet as adults, whilst the students themselves are likely to have grown up with the internet being readily available to them. As outlined in Chapter 1, there is no consistent evidence for the effect of age on susceptibility. The elements of this study that look at how people experience fraudulent communications in terms of frequency and response to these may provide insightful evidence though when compared between different age groups. Similar to Study 1, this questionnaire asked participants about personal experiences of fraudulent communications in offline and online environments, as well as the cues that they used to recognise suspicious emails as such.

Study 2a focuses on online fraud, as this will be the focus of the remainder of the thesis. This study also sampled attendees of university open days, who were asked to respond to a number of short questions and judge the legitimacy of a set of email screenshots displayed to them. These email stimuli make up part of a stimuli set to be used in the email legitimacy task employed in Studies 3-6. The data collected in this study will allow measurement of the typical responses and accuracy of users in judging these stimuli, before developing the extended version of the task for further studies. Four open days are described in this study, with different email stimuli being used across these in order to test a wider range of stimuli for the further studies.

Given the exploratory nature of Studies 2 and 2a no specific empirical hypotheses were generated. However, the general aim of the studies reported in this chapter is to establish how familiar users are with the issues and risks surrounding fraudulent communications, and whether there are certain common misconceptions that could potentially make users more vulnerable. Further to this, we aim to establish whether there are any links between users internet usage habits, and their understanding or awareness of email fraud.

2.2 Study one

2.2.1 Method

2.2.1.1 Participants. A sample of 68 participants, consisting of 57 females and 11 males, were all first year Psychology students at Lancaster University. Participants were aged between 18 and 29 years, with a mean age of 19.12 (SD = 1.58).

2.2.1.3 Materials.

Fraud experience questionnaire. All participants were asked to complete a questionnaire (shown in Appendix A [all appendices available at <https://dx.doi.org/10.17635/lancaster/researchdata/117>]) that incorporated a set of bespoke questions, developed to understand participant experiences and awareness of fraud in both online (email) and offline (face-to-face, telephone, and postal) environments. As well as asking for some basic demographic information, the questionnaire includes a combination of both multiple choice questions about the

type of fraud experienced and internet usage, as well as open-ended questions about cues used to make response decisions.

Self-control scale. Originally reported by Tangney et al. (2004), the brief self-control scale consists of 13 items, some of which include: *I am good at resisting temptation; I say inappropriate things; and People would say that I have iron self-discipline* (scale can be seen in Appendix B). A self-control score was obtained by summing the response values, after reverse scoring as appropriate.

Cognitive Reflection Test. The cognitive reflection test (CRT; Frederick, 2005) is a short test involving three problems, each of which has an intuitive response, which is incorrect. The test, shown in Figure 2.1, was originally developed using American currency in question 1, so this was changed to pounds sterling for the purpose of this experiment. The test is scored based on the number of correct responses that a participant gives, so each participant will receive a score between 0 and 3.

- 1) A bat and a ball cost £1.10 in total. The bat costs £1.00 more than the ball. How much does the ball cost? Five pence.
- 2) If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? Five minutes.
- 3) In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? 47 days.

Figure 2.1 The Cognitive Reflection Test

Focus group discussion. Participants were asked to contribute to a group discussion, which was structured by the researcher around a list of core questions. These were formulated to instigate conversation within the group, but whilst also loosely controlling the topics of conversation which were covered. An initial group

discussion with 7 participants allowed assessment of the quality of question content. Following this, the questions were adjusted in order to maximise the amount of relevant information elicited from the groups. The data from this initial session will not be used in the final analysis. The questions used to instigate discussion can be seen in Appendix C. Examples include '*What cues do you rely on to recognise a scam email?*' and '*When you think about victims of fraud, are there certain types of people who you imagine as the victims?*'. A total of eight focus group discussions were conducted, with between 6 and 11 participants in each.

2.2.1.4 Procedure. This study took place in a large computer lab within the Psychology department, so that participants could complete the online questionnaire and then the focus group session in one place. Once consent was gained, participants completed the questionnaire part of the study on an Apple iMac via Google Docs, with each participant using an independent computer.

Once participants had completed the questionnaire they were then asked to join a group for the discussion. The discussion was led by the researcher, based upon the list of questions mentioned above. The questions were adjusted to fit around what was being said by that particular group in order to gain as much insight as possible, whilst ensuring that all of the key areas of interest were addressed during the discussion. Participants were also given the opportunity to ask any questions at the end. The discussions were voice-recorded for subsequent transcription and analysis. Finally, participants were debriefed once all of the tasks were completed. Following the study, the researcher transcribed each audio recording and the original audio files deleted (transcripts can be viewed in Appendix D).

2.2.1.5 Data collation. Questionnaire data was collated, with closed ended questions inputted accordingly and open-ended questions given a single code that was deemed most representative for the response. Consistency in responses meant that a set of codes could be established that were representative of all participants. In order to analyse the effect of self-control and cognitive reflection on past response behaviour, cases where the participant reported never having received a fraudulent communication were not included, as the focus was on response behaviour.

A data-driven approach was taken in analysing data from the focus group discussions, given the exploratory nature of this research. Thematic analysis was used to highlight key areas of interest in order to gain insight into user experiences of fraud. Large sections of each discussion were coded under these broader themes to begin with, before subthemes were generated during a second round of coding. Analysis was partially influenced by the semi-structured nature of the questions developed for the discussions, so the main themes extracted followed a similar pattern to these. Subthemes were established within each of the main themes, based on coding of each participant response during the sessions. The specific themes constructed are described in more detail below.

2.2.2 Results

2.2.2.1 Questionnaire data. Of the 68 participants who took part in the experiment, 56 reported having received a fraudulent communication of some sort. The type of communications participants reported having received are shown below in Figure 2.2.




Figure 2.2 Percentage of participants who reported received online, offline, or both types of communication

Offline fraud. The contexts in which participants' experiences occurred are outlined in Table 2.1 below. Of those who had experienced some form of offline fraud, 29% responded to the communication. The low response rate means that most participants gave a unique response when asked their reason for responding, with one participant citing each of the following reasons: curiosity, communication seemed relevant, naivety, sympathy for the sender, a good cause. Two participants reported threatening behaviour within the communication as their reason for responding. The reasons participants gave for choosing not to respond to a fraudulent communication are shown in Figure 2.3.

Table 2.1 Percentage of participants who reported experiencing offline fraud in each context

| Received this type of communication | Face-to-face | Telephone | Post |
|-------------------------------------|--------------|-----------|------|
| Yes | 46 | 71 | 29 |
| No | 54 | 29 | 71 |

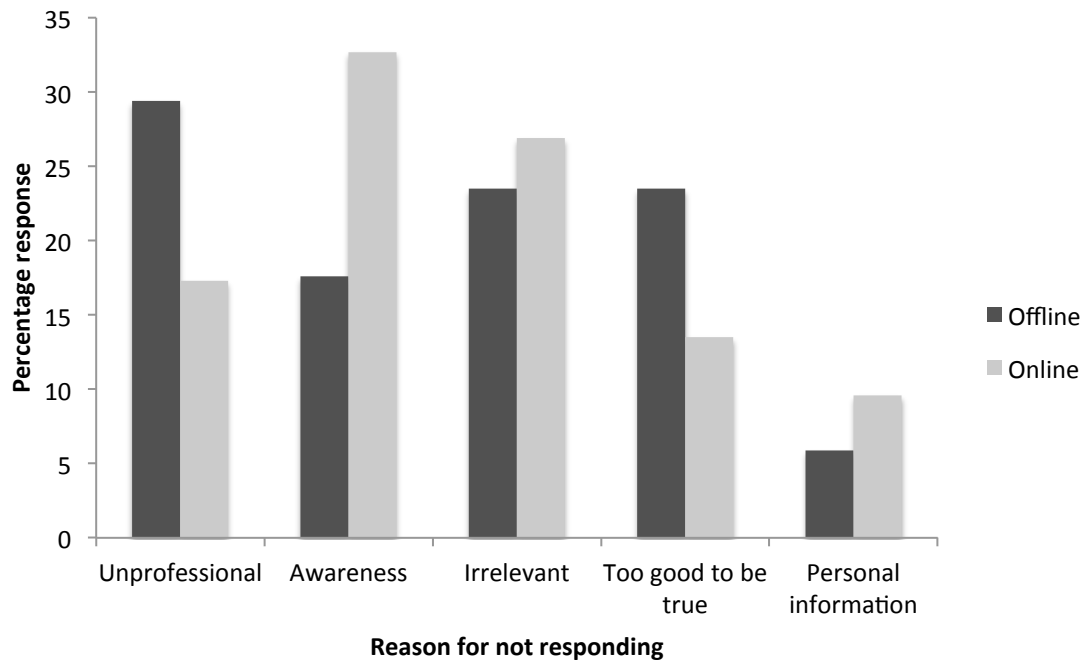


Figure 2.3 Reasons given for choosing not to respond to fraudulent communications in offline and online environments

Internet use. As shown below in Table 2.2, most participants (54%) reported spending between 3 and 6 hours per day on the internet. Participants were also given a multiple-choice question about what they used the internet for. The most common responses were social networking (88%) and university work (79%), whilst the least common were news (11.8%) and shopping (10.3%). When asked how long they had actively been using the internet, the majority of participants (62%) reported using the internet for over 6 years, whilst 32% had been using it for 3 to 6 years and only 6% for between 1 and 3 years. No participants reported using the internet for less than one year or not regularly using the internet.

Online fraud. The contexts in which online fraudulent experiences reported by participants occurred are outlined in Table 2.3 below. Of those who had experienced fraud online, only 4% reported responding to the communication. Accounting for those who have never received a fraudulent online communication,

Table 2.2 *Number of hours spent online each day by participants*

| Time spent on the internet per day | % participants |
|--|----------------|
| 0-1 hours | 0 |
| 1-3 hours | 22 |
| 3-6 hours | 54 |
| 6+ hours | 24 |
| I do not use the internet on a daily basis | 0 |

this meant that only 2 participants chose to respond, and the reasons given for this were financial benefit and the official manner of the communication. The reasons that participants gave for choosing not to respond are shown above in Figure 2.2.

Table 2.3 *Percentage of participants who reported experiencing online fraud in each context*

| Received | Computer at home | Computer at work | Tablet at home | Tablet at work | Mobile at home | Mobile at work |
|----------|------------------|------------------|----------------|----------------|----------------|----------------|
| Yes | 96 | 9 | 11 | 0 | 40 | 7 |
| No | 4 | 91 | 89 | 100 | 60 | 93 |

Cognitive measures. After removing those cases in which participants reported never having received a fraudulent communication, independent samples t-tests were conducted to establish whether there was a difference in self-control and cognitive reflection scores between those who had and had not responded to fraudulent communications in offline or online environments. The mean scores and standard deviations for each response in each environment are shown in Table 2.4. No significant difference in self-control scores between those who responded and those who did not was found in either offline, $t(22) = -0.07, p = .94, d = 0.03$, or online environments, $t(53) = -1.92, p = .06, d = 1.24$.

Further to this, no significant difference was found between scores on the cognitive reflection test and response to fraudulent communications in an offline, $t(22) = 0.33, p = .75, d = 0.14$, or online environment, $t(53) = 0.06, p = .96, d = 0.03$.

Table 2.4 Means and standard deviations for self-control scale and cognitive reflection test scores by response to fraudulent communication

| | Offline | | | | Online | | | |
|----------|--------------------|------|----------------------------|------|--------------------|------|----------------------------|------|
| | Self-control score | | Cognitive Reflection score | | Self-control score | | Cognitive Reflection score | |
| Response | Mean | SD | Mean | SD | Mean | SD | Mean | SD |
| Yes | 35.14 | 8.23 | 1.14 | 1.07 | 27.50 | 9.20 | 1.00 | 1.41 |
| No | 35.41 | 8.20 | 1.00 | 0.94 | 37.85 | 7.45 | 0.96 | 0.92 |

2.2.2.2 Focus group data.

Construction of themes. Given the semi-structured nature of the focus group schedule that was developed prior to the commencement of the group discussions, the themes constructed from the resulting data follows a similar pattern. Across the eight focus group discussion conducted, four themes were extracted, with a total of 11 subthemes. These are illustrated in Table 2.5 and are described in more detail below.

Knowledge and understanding of scams. Participants across all of the focus group sessions demonstrated an awareness of the most common types of scams in both online and offline settings. Offline examples include fake phone numbers for companies that are listed as the top hit on a Google search for the genuine company phone number, and key loggers on ATM machines that collect information from the victim's credit card and can then be used to make online purchases. Online examples include emails "pretending to be your bank, saying like, you need to put your bank details in and change your password" (P24, Session 6, Line

Table 2.5 *Framework of user focus group themes*

| Theme | Subtheme |
|---|--|
| 1. Knowledge and understanding of scams | <ul style="list-style-type: none"> 1.1 Types of scam 1.2 Knowledge from the media 1.3 Experience of family and friends |
| 2. Perception of 'typical' victims | <ul style="list-style-type: none"> 2.1 Age 2.2 Experience |
| 3. Personal experience | <ul style="list-style-type: none"> 3.1 Differences between university and personal email accounts 3.2 Emotional reaction |
| 4. Detecting deceptive emails | <ul style="list-style-type: none"> 4.1 Spam filtering 4.2 Authority level of the sender 4.3 Familiarity of the sender 4.4 Relevance of communication |

5), romance scams, “a lot that say I’ve won the lottery in like a different country” (P9, Session 2, Line 8), and emails that induce a sense of panic, for example by threatening the user with loss of access to their online account for a given company. Further to this, some participants demonstrated an understanding of the concept of spear phishing, which involves collecting personal information about the user from sources such as social media accounts, in order to make an email more targeted and believable. Participants reported that they may be more likely to trust such an email if it contained personal information, “because if they can tap into your interests and your hobbies they can tap into that to send you emails that you think ‘oh, this might be genuine, I might be interested in this’” (P33, Session 7, Line 80).

When asked to consider the difference between how they would react to online versus offline fraudulent communications, participants were split in their

opinions. Some argued that they would be better at detecting the fraudulent nature of these communications in an online scenario, saying that there are more cues available in an email rather than, for example, over the phone, “because it’s like a visual, you’re seeing it rather than just hearing it” (P1, Session 1, Line 66). In addition, these participants felt that they had a better understanding and awareness of online scams. On the other hand, some participants felt that offline fraud is easier to recognise, arguing that “over the phone, I don’t know, in general you get a feeling for whether they’re actually genuine or not” (P36, Session 8, Line 50), and that “you can hide a lot over the internet” (P36, Session 8, Line 49).

There was a general consensus amongst participants that the media provides useful information to users about the danger of scams. Some participants highlighted the value of news reports about recent scams and advice on how to avoid being conned, as well as the information provided in television programmes such as “the Real Hustle, that’s quite good” (P17, Session 4, Line 21), which demonstrate to the public how fraudsters plan and execute scams. A number of participants did however report that their best source of knowledge about fraud comes “from friends and family, like, people that have been affected by it” (P21, Session 5, Line 45). For example, some participants indicated that they are reliant upon their parents for advice and warnings about fraud. However, it seems that this knowledge sometimes only develops following a personal experience of the scam – meaning that in many cases it comes too late for the victim, even if it does mean that they are able to pass on a warning to those around them, as one participant points out:

Probably need to be kept more up to date so people are actually aware of new scams that are around because you quite often don't know about them until they're actually out there and you may have fallen for it. (P29, Session 7, Line 131)

Perception of 'typical' victims. There were some mixed responses amongst participants about who a 'typical' fraud victim might be. Most respondents commented on the age of the user, with many feeling that older users would be most likely to become victims, possibly because they are grateful for the company and attention that communication with another individual – in both online and offline scenarios - can provide:

like my Nan is always getting those kinds of phone calls and she's 90 so she's like completely oblivious to things like that and they're ringing up saying 'congratulations' and she feels 'aww, thank you', she's really happy, and she doesn't understand the concept at all that this person, that she hasn't won anything but she thinks she has. (P11, Session 3, Line 146)

However, some other participants reported that they considered younger internet users to be at higher risk. In both cases, participants gave similar responses when considering why it is that these groups may be more likely to be victims – which was due to them being “less accustomed to using the internet” (P6, Session 2, Line 106). Participants felt that those who have not been using the internet for as long may not be aware of the numerous scams that exist and how to identify them.

Personal experience. All participants reported having both a university and a personal email account, and most agreed that they were more wary of the emails that they received to the personal account. Many said that they use their personal email accounts to sign up for newsletters and websites, so they receive a lot more traffic – including that from third parties to whom their details have been sold – in that account. Given the smaller level of traffic, and the people who emails are from, participants seemed to be more trusting of emails that came into their university accounts as these were usually all relevant to the recipient. Participants did report that if an email from a trusted source came into their university account containing information that is “nothing related to the course then you’d be a bit more suspicious”. It seems a common assumption amongst our participants that lecturers’ accounts are more secure and can be trusted, in comparison to a friend’s Hotmail or Gmail account, which is deemed much more likely to be hacked:

I think you just automatically think, ‘oh it’s from my lecturer, it’s going to be fine’, whereas you know quite a few friends have been hacked before and you get quite a few spam ones, you just don’t think it would happen to a lecturer.

(P28, Session 7, Line 49)

Some participants reported feelings that they would be embarrassed if they did become a victim of fraud. It seemed amongst the participants that there was a kind of stigma attached to fraud victimisation, with one participant even suggesting that “I think you’ve got to be a bit stupid [to fall victim]” (P11, Session 3, Line 135).

Detecting deceptive emails. When asked what cues normally make them aware of the fraudulent nature of an email, participants gave a number of responses that typify the 'stereotypical' view of what a phishing attack would look like – elements such as “overuse of exclamation marks” (P6, Session 2, Line 5), spelling and punctuation errors, “a logo which is not exactly the same as the real people, but quite similar” (P12, Session 3, Line 15), and an impersonal approach, for example “in bank emails they refer to me by email address and not my actual name” (P18, Session 4, Line 12). Whilst these indicators are beneficial in spotting the more obvious phishing emails, the development of sophisticated approaches to phishing mean that they are not always reliable in spotting the more subtle attacks.

Overall, it seemed that most participants are not entirely trusting of the in-built spam filters that email providers use, with many emphasising that “there’s lots of stuff that it does stop, but like you said, there’s stuff that goes in there that shouldn’t and obviously there’s stuff that goes into your inbox that shouldn’t be there” (P6, Session 2, Line 30). This means that users are required to check their spam folder on a regular basis to ensure that no genuine emails have ended up in there, which most participants reported doing.

A lot of participants reported that they always open emails that come from lecturers, with most agreeing that they considered lecturers’ accounts more trustworthy and presumed “that they wouldn’t have viruses because they’re probably not like being fooled on the internet, not like your friends” (P11, Session 3, Line 105). When asked why they thought they were more trusting of emails from lecturers, participants suggested that this was likely to be due to the familiarity of the scenario as well as the authority level of the sender:

P11: *I don't know, if I got an email off the Police and you opened it then I'd still be suspicious...I don't think it's to do with that.*

P12: *It depends how realistic it is, because realistically we are quite likely to get an email from a lecturer saying 'oh, could you open this and fill it in before you come' or something like that, it 'we're conducting this kind of survey' because they do it all the time, but the Police, like with your example, I would be wary because the Police don't just go around emailing you. (Session 3, Line 112)*

Therefore, it seems there is an interaction between the authority level and the familiarity of the sender to the user in the believability of a fraudulent email. Whilst some participants commented that they would “probably click on it if it was my friend” (P21, Session 5, Line 75), compared to something coming from an unknown company, many participants commented in reference to emails from friends, “it’s different, like, a lot of my friends, the link is via Facebook, I wouldn’t normally get emails from them so I’d find that a bit strange” (P5, Session 2, Line 65). Therefore, it seemed that they would be more likely to immediately question the veracity of an email from a peer. Most participants agreed that in this situation they would contact the friend who the email came from “because then you’d have to let them know that their account has been hacked” (P5, Session 2, Line 75). Others highlighted a feature within Hotmail, which is likely available from other email providers as well, whereby there is an option to click ‘My friend has been hacked’. This allows the email provider to address the issue, as well as contact the account holder and provide them with relevant advice.

Finally, it was clear from all participants that the relevance of any given communication is a key factor in their likelihood to pay attention or respond to it. Participants reported that they would feel more suspicious “if it’s an email [address] you just don’t recognise, that you haven’t been involved with before” (P37, Session 8, Line 12). It seems that the success of a given phishing attack is partly dependent, on a user-by-user basis, upon the relevance of the communication to the given recipient.

2.2.3 Discussion

This study aimed to gain an insight into the personal experiences of users with regard to online and offline fraud. With two sections – a short questionnaire, and a follow-up focus group discussion - exploratory information was elicited about users’ familiarity with fraudulent communications, and their approaches to dealing with these.

Considering how commonplace phishing emails are these days, accounting for one in every 392 email sent in 2013 (Symantec, 2014), it was surprising to find that 18 per cent of the sample claimed to have never received a fraudulent communication, either online or offline. However, this may be due to the limited demographic of the sample, with a narrow age range, and a low mean age. These participants may only have had email accounts for a short period of time and be less likely to receive much mail through the post or through other mediums of offline fraud. Alternatively, it may be that these participants have received such a communication, but were unaware of its fraudulent nature. Of those participants who did report having received a fraudulent communication, it was apparent that

this was more common in online scenarios than offline. Due to the ease for fraudsters to distribute their scams online, as well as being more economically viable in the costs to send out the scam related to the returns obtained, internet based scams such as phishing emails are becoming more commonplace than alternative methods, such as telephone and postal fraud.

When asked about the cues used to recognise a fraudulent communication as such, participants gave many of the same responses when referring to offline and online scams. Across both environments, the unprofessional appearance, and irrelevance of communications were in the top three reasons for participants choosing not to respond. For offline fraud, 'too good to be true' scenarios were also frequently given as a reason to not respond, whilst for online fraud, awareness of the scam or scams in general was given as the top reason for recognising the fraudulent nature of the communication. Given the higher prevalence of email fraud, it might be that participants are more aware of particular online than offline scams due to increased media coverage, or discussion amongst peers.

Hypotheses 1 and 2 predicted that lower levels of self-control and cognitive reflection, respectively, would be related to an increased likelihood to have responded to a fraudulent communication in the past. However, neither of these variables was found to be associated with likelihood to respond. This may be due to the small response rate within the sample though. In a young demographic, past response may underestimate susceptibility, as it may be that users just have not received as much email traffic, and so have not yet received a convincing enough phishing email. With an alternative measure of potential susceptibility, a difference in performance may become apparent.

During the focus group discussions, participants demonstrated a knowledge of various types of scams, in both online and offline environments. Despite this, some admitted that they would find a spear phishing attacks, incorporating personal information about them, to be more believable. This suggests that awareness of scams does not necessarily mean immunity to them, which may be important in the attempts made to educate users and reduce their susceptibility. Further to this, most users commented that they would trust anything that came from a lecturer's email account because they would be more trusting of it, and felt that those in more senior positions would not be hacked. This reinforces findings from Guéguen and Jacob (2002), which highlight the influence of authority on response likelihood to a phishing email. It seems that although participants are aware of online fraud and how prevalent it is, they do not realise that anyone could be susceptible to hacking, regardless of their position or level of authority.

This was emphasised when participants were asked who they perceived to be more vulnerable to fraudulent communications. There seemed to be two extremes in participants responses: either older users who have had less experience in using computers and are less familiar with technology; or younger users who have not had online accounts for as long and so would not get the same volume of traffic into their accounts. Whilst there is empirical evidence to support both of these opinions (as discussed in Chapter 1), the contradictory nature of past research findings suggest that there is no specific demographic profile of who is most susceptible to fraud, and this is an important point in itself for users to understand – that anyone can become a victim in a certain situation or frame of mind.

When asked about the cues that they would use to recognise fraudulent communications, participants listed numerous stereotypical cues, such as poor spelling and grammar, which are often out-dated when it comes to the more sophisticated phishing emails that are becoming more commonplace and pose a higher risk to the user. Whilst users felt that most media coverage of scams was relatively useful, it is apparent that what is being learnt from this is a very basic understanding of the threat of fraudulent communications, which does not advance at the same rate as the sophistication of the scams. Certain approaches can be valuable though, for example television shows like the Real Hustle, which educate users on the methods used by the fraudsters so that users can understand how the scams are developed rather than simply how a message might look. It seems however, that the most effective method of education currently comes from personal experience, when it is too late and the user has already become a victim. This emphasises the need to understand users' susceptibility and attempt to develop ways to educate users that remain effective in the long term, even with the advancement of the techniques used within the scams making them more believable.

The use of focus groups in this study encouraged discussion between participants, eliciting a valuable data set that provides insight into the extent to which our sample has experienced fraud, and their perceptions of this. However, as discussed in previous literature (Smithson, 2000; Leung & Savithiri, 2009), one limitation to this methodology should be noted. In each of the group discussions, some members of the group were more dominant in responding to questions and discussing experiences. The group setting means that more introverted participants

or those with little experience of fraud are able to avoid speaking aloud in front of a group. However, this does mean that the data collected is limited to a smaller number of more outspoken group members and may have masked information withheld by those more introverted participants. There were multiple participants engaging in conversation in each focus group though, even if this did not involve all participants, providing a sufficient sample of data for this initial exploratory study.

From this study, in particular the focus group discussions, it is clear that participants demonstrate a basic, but not a comprehensive understanding of fraudulent communications. Although they are able to identify a number of types of scams, there is little appreciation for the fact that anyone, of any age or professional status, might become a victim to fraud, or account hacking. It might be that long-term educational approaches need to address the theoretical logic behind how scams are developed and executed in order for users to be able to recognise novel, more sophisticated scams as they emerge. However, the data collected in this study was reliant on participants recalling from memory the cues that they utilise when judge the legitimacy of communications. Therefore, a real time assessment of cues utilised by participants in decision-making surrounding such communications may be beneficial.

2.3 Study two

2.3.1 Method

2.3.1.1 Participants. A total of 108 participants were recruited at a departmental UCAS open day at Lancaster University for the Psychology department. The sample consisted of 36 males and 72 females, aged between 17 and 69 years as

both applicants and their families were present at the open day. The mean age was 32.28 years ($SD = 16.15$).

2.3.1.3 Materials. A questionnaire was developed for the purpose of this experiment (this can be seen in Appendix E). This consisted of demographic questions, as well as questions about the contexts in which participants had received fraudulent communications, and their reasons for responding or not responding. These questions were designed to elicit information about the frequency of scam occurrence, as well as an overview of the cues used in detecting fraud. The nature of the questions included meant that some were multiple-choice, allowing simple quantitative analysis, whilst open-ended questions about reasons for responding or not responding to a phishing email required coding for the purpose of content analysis.

2.3.1.4 Procedure. At the UCAS open day, Dr. John Towse gave a talk to all prospective students and their parents, about the work being carried out on this project. This highlighted the relevance of psychological processes underlying fraud victimisation and why this is an important area to study. At the end of the talk, a consent slide was displayed as part of the PowerPoint display from the talk and this was discussed with the participants to ensure that they understood that participation was completely voluntary and they had the right to withdraw their data should they wish to do so. Whilst this talk was being given, a copy of the information sheet and the short questionnaire were given to each member of the audience. Participants

were then walked through the questions in front of them whilst they filled in the questionnaire.

2.3.1.5 Data collation. Responses to closed-ended questions from this study were inputted accordingly, whilst responses to open-ended questions were allocated a single code, deemed to be the most representative for the response. Frequency data about each of the coded responses was then calculated in order to demonstrate the cues that users are most reliant on in email management.

In order to assess the distinction between the two age groups within the sample (ie. students and parents), participants were categorised as being either in the 'younger' or the 'older' age group. Based on the distribution of participant ages within the sample, a threshold of 30 years was used to distinguish between members of the sample. This resulted in 59 participants categorised in the 'younger' age group, with an age range between 17 and 28 years, and 49 participants in the 'older' age group, with an age range from 39 to 69 years.

2.3.2 Results

2.3.2.1 Descriptive statistics. When asked whether they had personally received a fraudulent communication, 100 per cent of participants reported that they had. Of these participants, 88.9 per cent reported that they had received either an online communication, or both online and offline communications, whilst only 11.1 per cent reported having received just offline fraudulent communications. Participants were then asked whether they had responded to such a communication, and 8.3 per cent reported that they had.

2.3.2.2 Demographic data. There were two distinct age groups in this sample, given the nature of the open day where data was collected. A chi-square analysis to compare self-reporting of past response to fraud found no difference between the two age categories, $\chi^2 (1, N = 108) = 1.80, p = .18, V = .13$. A further chi-square analysis was run to establish whether gender was influential on past response, but no significant difference was found, $\chi^2 (1, N = 108) = 0.55, p = .46, V = .07$.

2.3.2.3 Context information. Table 2.6 below shows the descriptive statistics obtained when participants were asked about the contexts in which they have received both online and offline communications. These demonstrate that participants report receiving more fraudulent communications, both online and offline, in the home environment than in a work environment. It is also apparent that online communications are most commonly received via computer, whilst offline communications mostly by phone.

2.3.2.4 Reasons for response decision. Those participants who reported having not responded to a fraudulent communication were asked how they identified the malicious nature of these. Their responses were coded, and percentage responses are shown below in Table 2.7. The sample of participants who reported that they had previously responded to a fraudulent communication was small ($N = 9$), so most reasons given were only given by one participant. These included: financial need, negative consequences of not responding, and being unfamiliar with issues surrounding fraud. There were also five participants who did not give a reason for their choice to respond.

Table 2.6 *Percentage of participants who reported having received fraudulent communications through each medium in online and offline environments*

| Context | Percentage |
|----------------------|------------|
| Online | |
| Computer at home | 84.3 |
| Computer at work | 16.7 |
| Tablet at home | 3.7 |
| Tablet at work | 0 |
| Smartphone at home | 21.3 |
| Smartphone at work | 1.9 |
| Offline | |
| Phone at home | 30.6 |
| Phone at work | 5.6 |
| Post at home | 12 |
| Post at work | 2.8 |
| Face-to-face at home | 2.8 |
| Face-to-face at work | 1.9 |

Table 2.7 *Frequency of reasons given for not responding to a fraudulent communication*

| Reason for response decision | Percentage |
|---------------------------------|------------|
| Awareness of scams | 55.6 |
| Irrelevant communication | 17.2 |
| Appearance/wording of email | 10.1 |
| 'Too good to be true' | 8.1 |
| Asking for too much information | 3.0 |
| Unclassifiable | 6.1 |

2.3.3 Discussion

During this experiment, basic information was collected about the situations in which users receive fraudulent communications, and the cues that they use to detect these as such and prevent victimisation. All participants in this study reported that they had received a fraudulent communication, with a vast majority having had such a communication online at some point. The self-reported response rate from

participants was slightly higher than expected at 8.3 per cent, where the reported average rate is around 5 per cent (Norton, 2014). Fraud victimisation, especially online fraud, is known to go heavily underreported though (Copes, Kerley, Mason, & Van Wyk, 2001), and so it might be that when asked directly and anonymously, participants were more likely to acknowledge their victimisation.

There were two distinct age groups involved in this study, given the nature of the event where data was collected -with students attending the open day with their parents. Previous research has provided contradictory evidence about the age group that might be most at risk, as discussed earlier in Chapter 1, with some indicating that older users are more at risk (Pak & Shadel, 2011; Shadel & Pak, 2007) whilst others suggest that younger users are more susceptible (Sheng, Holbrook, Kumaraguru, Cranor, & Downs, 2010; Pratt, Holtfreter, & Reisig, 2010). In line with this, no difference was found between these two age groups in their likelihood to have responded to a fraudulent communication. As with Study 1 though, the small response rate found within our participants means that the comparison between those who did and did not respond is not necessarily representative given the vastly different group sizes.

In all scenarios, across both online and offline environments, participants reported receiving more fraudulent communications whilst at home than whilst at work. This may simply be because participants receive a higher volume of emails to personal accounts and addresses than at work, or because security software is more effective in an organisational setting. In this short questionnaire, participants who had responded to fraudulent communications were not asked about the scenario in which this occurred, so it cannot be reported whether a higher volume of traffic to

certain mediums makes users more susceptible. Although Vishwanath et al. (2011) provide evidence that increased email load makes users more susceptible to phishing emails, there is not enough evidence from this questionnaire to support or refute this. Situational factors should also be considered – for example, at work a user might be under increased pressure, and thus be more prone to errors in decision-making, whilst at home they may have more time to contemplate emails. Such influences will be considered in more detail later in the thesis.

As shown in Study 1, awareness of scams was the most frequently reported cue to recognising scams as such. Whilst this is a positive finding, in that it suggests that users are able to recognise some scams, it does mean that their reliance in recognising a scam based on familiarity might leave them more susceptible when faced with a novel, more sophisticated scam. The second most frequent response was that the communication received was irrelevant to them, which again is a useful cue for more generic phishing. However, more sophisticated spear phishing attempts often use personal information to make the communication more relevant to the user. Therefore, users seem well prepared to protect themselves against basic phishing attacks. However, users are often unable to apply knowledge to unfamiliar risks (Downs et al., 2006), which raises concern for ability to protect against novel attack strategies.

2.4 Study 2a

2.4.1 Method

2.4.1.1 Participants. A total of 371 participants completed the study, across four different UCAS open days at Lancaster University. Participants were visiting the

university either as prospective students, or as the parents/guardians of these prospective students. Information about participants' gender was only recorded for open days 1, 2, and 3 (with 12 participants who did not report gender across open days 2 and 3), but the age of participants across all open days ranged from 16 to 72 years ($M = 32.24$, $SD = 16.41$). Participant information for each of the open days individually, including participant gender for those where this was recorded, is shown in Table 2.8.

Table 2.8 *Descriptive statistics for each open day*

| Open day | N | Age | | Age group | | Gender | |
|----------|-----|-------|-------|-----------|--------|--------|--------|
| | | Mean | SD | Student | Parent | Male | Female |
| 1 | 145 | 31.73 | 16.65 | 79 | 63 | 34 | 101 |
| 2 and 3 | 145 | 32.56 | 16.15 | 71 | 65 | 41 | 92 |
| 4 | 81 | 32.60 | 16.58 | 43 | 38 | - | - |

2.4.1.3 Materials. A similar set of questions to the two previous studies were developed for the purpose of this study to gain insight into participants' perception of email fraud and previous history of responding to phishing. In addition, this study incorporated a small set of email stimuli that participants were asked to judge on legitimacy, in order to assess their ability to recognise phishing emails. At each open day, four emails were displayed, each of which was obtained from either the researcher's own inbox, from friends and colleagues, or from online articles relating to phishing. The emails varied across the open days, but were all stimuli that will also be used in later studies in this thesis. Therefore, data permit an analysis of the difficulty of the legitimacy judgment task. Questions were answered in paper form at the first of the open days, but personal response units were used for the

remainder. These devices allow participants to respond to questions displayed on a screen at the front of the lecture theatre.

2.4.1.4 Procedure. At each of the UCAS open days, a talk was given to prospective students and their parents about the current research project looking at individual differences in susceptibility to fraudulent emails. At the first open day, participants completed a pen and paper version of the questionnaire (see Appendix F), handed out during the talk. At the remainder of the open days, each attendant was handed a Promethean ActivExpression response unit on entering the lecture theatre, which would allow him or her to respond to the questions following the talk (shown in Appendix G), via an ActivHub device plugged into the computer. This allows for responses to be collected wirelessly. At the end of the talk, a consent slide was displayed on the screen at the front of the lecture theatre, which outlined the voluntary nature of participation in the questionnaire and the participant's right to withdraw their data. The questions and email stimuli were then displayed one at a time on the screen, through the Flipchart software that is designed to work with the response units. Each question had a 30 second time limit, after which the next question would be displayed. Responses were recorded with an anonymous response unit number associated with each, and were then downloaded and saved to an encrypted hard drive following the talk.

2.4.1.5 Data collation. Data for questions that were displayed at all of the open days were collated for analysis. However, the emails presented at the open days were different each time, so it may be that the type of emails presented are

also having an effect on performance in this task. As a result it was decided that the variables affecting accuracy in judging emails would be analysed separately for each open day.

As in Study 2, participant age was categorised into a 'younger' and 'older' age group to distinguish between student and parent sample groups. Using the threshold of 30 years again, the sample consisted of 193 participants in the 'younger' age category, with an age range of 16 to 25 years, and 166 participants in the 'older' age category, with a range from 34 to 72 years.

2.4.2 Results

2.4.2.1 Collective analysis. The collective analysis demonstrated an overall rate of 5.9 per cent for participants who reported that they had previously responded to a phishing email (although there were 75 participants who did not respond to this question). Of these participants, a higher percentage (78%) were older participants than younger participants (22%), $\chi^2 (1, N = 296) = 5.63, p < .05$, however the effect size for this finding was small, $V = .02$. For the open days which included a question about gender, no significant difference was found in previous response likelihood between males and females, $\chi^2 (1, N = 208) = 0.18, p = .67$. When asked how at risk they felt to online fraud (on a 5-point scale), the parent age group reported feeling at significantly higher risk ($M = 3.75, SD = 1.11$) than the student age group ($M = 3.17, SD = 1.07$), $t (356) = 5.02, p < .001, d = 0.53$. No difference in risk perception was found between male ($M = 3.32, SD = 1.14$) and female ($M = 3.58, SD = 1.06$) participants, $t (265) = 1.73, p = .08, d = 0.24$.

Backward multiple regression analysis on the collated data was conducted to establish whether the number of emails correctly identified could be predicted by any of the following variables – age, perceived risk, and previous response. The best-fit model from this analysis produced an R^2 value of .01 (adjusted $R^2 = .01$, $F(1, 294) = 3.37$, $p = .07$), and included age only, but this was not found to significantly predict accuracy on judging email legitimacy ($\beta = -.01$, $p = .07$).

2.4.2.2 Responses from open day 1. When asked whether they had ever received a phishing email themselves, only 57.2% of participants reported that they had, 15.9% reported that they had not, and 39% reported that they did not know. Table 2.9 shows a summary of the reasons given by participants who had received a phishing email, for choosing not to respond.

Table 2.9 *Frequency of reasons given for not responding to a phishing email*

| Reason for response decision | Percentage |
|-------------------------------------|------------|
| Appearance/wording of email | 34.9 |
| Sender was unfamiliar | 21.7 |
| Requested personal information | 20.5 |
| Communication was irrelevant | 10.8 |
| Recognised the scam | 3.6 |
| Offer seemed too good to be true | 1.2 |
| Unclassifiable (or did not respond) | 7.2 |

The percentage of correct responses for each of the emails which participants were asked to judge are shown below in Table 2.10. The NatWest phishing email was found to be the most easily identified, whilst the Amazon email proved the most difficult for participants to recognise as a phishing email.

Table 2.10 *Percentage of participants who correctly recognised each email*

| Email | Phishing/Legitimate | Percentage correct |
|---------|---------------------|--------------------|
| NatWest | Phishing | 79.3 |
| Paypal | Phishing | 53.1 |
| eBay | Legitimate | 44.1 |
| Amazon | Phishing | 24.1 |

Backward multiple linear regression analysis was conducted to establish whether the variables addressed in the questionnaire were predictive of participants' ability to correctly identify the four emails presented to them. Age, gender, perceived risk, and previous response were included as predictor variables, but the best fit model included only gender ($\beta = -.34, p = .11$) and perceived risk ($\beta = .11, p = .20$) as predictors, with an R^2 value of .01 (adjusted $R^2 = .01, F(1, 140) = 1.74, p = .19$), but neither were significant predictors.

2.4.2.3 Responses from open days 2 + 3. The percentage correct response for each of the four email stimuli presented to participants is shown below in Table 2.11. This demonstrates that again, the NatWest phishing email is the most easily identifiable, as at open day 1, whilst the legitimate emails included in this set were more often judged incorrectly.

Table 2.11 *Percentage of participants who correctly recognised each email*

| Email | Phishing / legitimate | Percentage correct |
|----------|-----------------------|--------------------|
| NatWest | Phishing | 82.1 |
| Barclays | Phishing | 73.8 |
| Dropbox | Legitimate | 47.6 |
| Paypal | Legitimate | 33.8 |

Backward multiple linear regression analysis was conducted with four predictor variables – age, gender, perceived risk, and previous response. The best-fit model included age ($\beta = -.01, p = .07$) and perceived risk ($\beta = -.12, p = .13$) as predictors, although neither was significant. This model produced an R^2 value of .06 (adjusted $R^2 = .04, F(2, 120) = 3.67, p < .05$).

2.4.2.4 Responses from open day 4. Table 2.12 shows the percentage correct response for each email presented to participants. In this set, the Caxton FX legitimate email was the most accurately judged, whilst the Facebook phishing email proved the most difficult for participants to correctly judge.

Table 2.12 *Percentage of participants who correctly recognised each email*

| Email | Phishing / legitimate | Percentage correct |
|-----------|-----------------------|--------------------|
| Caxton FX | Legitimate | 71.6 |
| DirectGov | Phishing | 69.1 |
| Facebook | Legitimate | 45.7 |
| Facebook | Phishing | 23.5 |

Backward multiple linear regression analysis was conducted, which demonstrated that all three variables included - age ($\beta = -.01, p = .18$), perceived risk ($\beta = .13, p = .15$), and previous response ($\beta = .63, p = .18$) - produced the best-fit model, although none of the variables were significant predictors. This model produced an R^2 value of .06 (adjusted $R^2 = .02, F(3, 77) = 1.63, p = .19$).

2.4.3 Discussion

Across the four open days at which we collected data for this study, questions remained mostly the same, but with different email stimuli being used to assess

participants' ability to differentiate between phishing and legitimate emails. This was part of the development of a set of stimuli for the email legitimacy task to be used in studies 3 to 6. Using these stimuli at the open days allowed us to get an idea of how 'easy' or 'difficult' categorisation was. The additional questionnaire items provided information about the relationship between perceived risk, age, response likelihood, and gender (open days 1, 2, and 3 only).

For those questions that were the same across the open days, older participants, namely the parents of prospective students, were found to be more likely to have responded to a phishing email. This may be due to the age of our younger participants - they may have spent less time actively using emails as a form of communication and so have had less exposure to email fraud. An alternative possibility is that the younger participants are just as likely to have fallen victim, but are unaware that they have, or are more embarrassed to admit this. The effect size for this age difference in previous response was small, which suggests that there are additional factors influencing response likelihood, as well as age.

Participants in the older, parent age group also reported feeling more at risk to online fraud, which may be because of greater exposure to cyber security issues, through immersion in the use of digital technology across a wider range of environments – from the office to the home. The data from open days 2 and 3 showed that those participants who felt more at risk to online fraud performed worse when asked to judge the legitimacy of the email stimuli. Although this finding was only present in a portion of our data, it is noteworthy, as it demonstrates that a lack of confidence in their own ability may result in a user being more susceptible to phishing emails, something that could be investigated further.

When looking at the email stimuli that participants were better or worse at identifying, there are certain characteristics that stand out and may be important in influencing a users' judgement. At the first three open days, participants were most accurate in identifying a NatWest phishing email. This email contained the correct company colours, however was poorly laid out and did not mirror the format of emails that you would expect to receive from the company. On the contrary, the email that participants most commonly misjudged at the first open day was a phishing email from Amazon, which incorporated the company logo and exact layout and formatting as a genuine email from the company, with an incorrect email address being the only cue to its illegitimacy. At the second and third open days, the least accuracy was with the two legitimate emails shown. Although both of these emails incorporated the company logos and were laid out in a typical style for that company, participants were obviously suspicious of them. This may be due to the nature of the task, in that participants knew they were looking to identify phishing emails, and so may have been over cautious. Or it may demonstrate a genuine over cautiousness within our participant sample. On the fourth open day, a legitimate email was the most easily identified, as this again had the appearance that would be expected in an email from that company and was not asking for unusual personal information from the user. The least accurately identified email was a phishing email from Facebook, which contained all of the graphics and formatting that a genuine email would have, making it difficult for the user to recognise it as a phishing email.

These findings suggest that participants are reliant on the appearance of an email as an identifier of veracity, showing that more sophisticated and realistic looking emails are more commonly misidentified. In the phishing emails that were

commonly mistaken for legitimate, a small cue such as the email address of the sender, was often the only identifiable factor that would have allowed a user to recognise it as phishing. Given the setting in which this data was collected, where participants were reading the emails on a large screen at the front of a lecture hall, it is possible that these details may be less noticeable than if they had the email in front of them with unlimited time to assess its veracity. Further investigation would be worthwhile in order to establish whether these elements are commonly missed in other environments as well.

2.5 General discussion

Through a number of exploratory studies, this chapter aimed to outline how every day users perceive email fraud, how they manage it, and how well they are able to identify phishing and legitimate emails. It was hoped that this would outline the need for further research to understand why some users demonstrate higher susceptibility than others and how this might be effectively reduced.

A number of key findings are outlined across these three experiments, which helped to address these aims. Previous research has demonstrated that phishing emails have an average response rate of five per cent (Norton, 2014). Data from Studies 2 and 2a indicate similar response rates in participants, both being below 10 per cent. This suggests that the samples are representative of a normal population of internet users in this respect.

During the focus group discussions in Study 1, participants demonstrated opinions that certain users might be more vulnerable to fraud victimisation than others. Some participants believed older users who are unfamiliar with the internet

are more vulnerable, whilst others believed that younger users who have not been using the internet for as long would be likely to be more at risk. Other common misconceptions were also demonstrated; such as the opinion that lecturers email accounts were trustworthy because it was unlikely that someone in this position would have their account hacked. Whilst it may be true that some people are more vulnerable than others, previous research suggests that these differences are due to more than basic demographic factors such as age and status (as outlined in Chapter 1). It is therefore important to ensure in future efforts to educate users, that it is emphasised that *anyone* can become a victim of fraud, regardless of factors such as those discussed here.

In Study 2a, participants were asked how at risk they felt to email fraud, and results demonstrated that this differed between the two distinct age groups in our sample – the prospective students and their parents. The older age group reported feeling more at risk to email fraud. This does not necessarily mean that they *are* more at risk, but rather that they less confident in their ability to recognise a phishing email. There was no difference in ability to recognise the emails that we displayed to participants depending on age. This was only a small set of stimuli, and in an artificial environment, however it would suggest that the differences in perceived risk are unfounded, and instead may suggest that the younger age group demonstrate an overconfidence in their ability to protect themselves against email fraud.

In each of the studies described in this chapter, participants were asked about their reasons for not responding to a fraudulent communication. Across all studies, similar cues were outlined, suggesting that these are likely to be used

commonly by the majority of users. These include the communication being irrelevant to the recipient, awareness or recognition of the scam, and the unprofessional appearance of the communication. Unfortunately, more sophisticated, spear phishing emails are able to surpass such cues, with exact copies of genuine emails from a company, often including information personal to the user to make them seem more relevant.

When shown the sets of email stimuli in Study 2a, participants were also better able to identify those emails with clear deviations in appearance from an email sent legitimately from the company in question. Those phishing emails that were better counterfeits of the original company emails were found to be more convincing, with a lower recognition rate. In comparison to the average response rate, mentioned above, the overall recognition rates in this task were very low. It is possible that the simulated nature of this task reduced participants' motivation to correctly identify the emails. In addition to this, given that the emails were displayed on a screen at the front of the lecture theatre, the more intricate details that could indicate phishing may have been less obvious.

2.4.1 Conclusion

The findings across studies in this chapter suggest not only that users create unsupported assumptions about who is more likely to become victim of fraud, but also that their own ability to recognise phishing emails is reliant upon generic cues that are not always present in more sophisticated attacks. Therefore, it is felt that these studies emphasise a need to focus more long-term educational attempts towards training users to understand the foundations of scams, how they are

constructed and executed. Rather than learning to identify specific scams, this would help users to recognise scams based upon the underlying techniques used to execute them rather than the superficial appearance. In this case, when the fraudsters develop new, more sophisticated techniques, users are in a better position to recognise these as fraudulent, rather than relying on out-dated cues.

Further to this, although a difference was found in previous response between parents and students from the open day data, this was a small effect and indicates that there is much more to the variance in response likelihood than age. The lack of evidence to suggest demographic influences on susceptibility supports the need for further research into a more comprehensive explanation of individual differences in response likelihood. In turn, this could provide insight for educational materials, in targeting them at the most susceptible users, and incorporating additional materials that focus on any cognitive or situational factors found to be influential.

Chapter 3

The cognitive make-up of the user as a predictor of email fraud susceptibility

Chapter summary

The study reported in this chapter outlines a set of cognitive variables to be examined in relation to susceptibility to email fraud. These variables are taken from previous fraud research, or related research from fields such as risky decision-making, and will provide an initial insight into how the cognitive make-up of a user might predict susceptibility. Using an email legitimacy task, developed for the purpose of the experiment, three cognitive variables were highlighted as predictors of accuracy in differentiating between phishing and legitimate email. These were: cognitive reflection, inhibition, and sensation seeking. In addition to this, a time pressure manipulation was included as part of the email task, with some participants given a limited amount of time to complete the task. Results demonstrated that accuracy on the email task was impaired for participants in this time pressure condition.

3 Study three

3.1 Introduction

As discussed in Chapter 1, there are a number of approaches to be considered when addressing the question of why people become victims of online fraud, drawing from computer science and psychology. Research from the field of human computer interaction (HCI) has considered the elements within an email itself that may convince the user of its 'legitimacy', which can also be combined with research in the psychology of persuasion demonstrating the social influence of techniques such as manipulating the authority level (Guéguen & Jacob, 2002) and familiarity (Jagatic et al., 2005) of the sender. The influence of such factors was highlighted in Study 1, with users in a focus group discussion commenting that they would be more likely to trust an email from a familiar person of higher authority, as they felt this person would be less likely to have been hacked. However this may only provide a partial explanation of susceptibility, as it does not help to explain how thousands of users can receive the same email (including these persuasive techniques), and yet only a few will respond. Therefore it is not sufficient to conclude that the content of the email itself is the sole explanation for likelihood to respond. Whilst these factors may be influential to some extent, it is important to consider these in combination with the individual differences in the cognitive processing capacity and cognitive make-up of the users that may lead to a response, whilst other users will immediately click 'delete'. Thus, this study aims to combine broader research topics in an attempt to understand susceptibility to fraudulent emails.

In Chapter 1, the potential relevance of dual-system reasoning in email decision-making is discussed. There is some evidence to suggest that when users are

reliant on intuitive, immediate responses to email stimuli, they are more likely to make errors in legitimacy judgments (Yan & Gozu, 2012; Harrison et al., 2016). Dong et al. (2008) note that the intuitive response users often have to a phishing email can be encouraged by the fraudster, making the recipient believe that this is the trustworthy, rational response. This reliance on intuition may be especially prevalent in situations where the user is under limited time constraints, or increased cognitive load. In a real world scenario, such pressures may come from an email requiring an immediate response, or a need to complete multiple tasks simultaneously. In order to assess how this affects accuracy in legitimacy judgments, a time pressure condition was added to the email task used in this study to measure susceptibility. Based on previous findings from Yan & Gozu (2012) and dual-system theories of reasoning, which suggest that intuitive reasoning can lead to errors in decision making (Tversky & Kahneman, 1975), the following hypothesis was generated:

H¹: Participants in the induced time pressure condition will perform with lower accuracy on the email legitimacy task.

One explanation for the likelihood to engage more rational decision-making processes, aside from situational influences, relates to working memory capacity (Kyllonen & Christal, 1990; Markovits, Doyon, & Simoneau, 2002). Hinson, Jameson, and Whitney (2002) report that participants demonstrate more impulsive, risk-taking behaviour, associated with system 1 reasoning, when under increased working memory load through implementation of secondary tasks. As well as load varying due to situational factors, such as increased cognitive load, working memory capacity can also vary between individuals (Conway et al., 2007). In this sense, it is possible

that users with a greater working memory capacity may be consistently better at detecting phishing emails, as they are better able to engage rational decision-making strategies. Based on this, the following hypothesis was generated:

H²: Participants with a higher working memory capacity will perform with higher accuracy on the email legitimacy task.

The consideration of factors such as working memory introduces a novel area of interest, by considering how differences in cognitive make-up between users may influence susceptibility to phishing emails. Chapter 1 outlined a number of additional variables, working in parallel with working memory capacity, which may contribute to an explanation of susceptibility based on their association either with intuitive decision-making, or risk-taking behaviour. These will be assessed in this study alongside variables that have shown direct links to fraud victimisation in past research.

Inhibitory capacity has been linked to working memory capacity in previous research (Engle, 1996; Redick, Heitz, & Engle, 2007). This ability to suppress an intuitive response to a task may therefore relate to email decision-making under principles of dual-system processing. If a user is able to inhibit a response, by looking past distracting information designed to make an email look legitimate, they are potentially more likely to recognise the cues available that identify an email as phishing, thus leaving them less susceptible.

Similarly, the concept of cognitive reflection (Frederick, 2005) is considered in relation to susceptibility. As discussed in Chapter 1, research has related this concept to decision-making surrounding risk preferences. That is, participants with higher

levels of cognitive reflection were more likely to take a calculated risk in order to receive a larger pay out at a later date than take an immediate, smaller pay out. By analogy, internet users who demonstrate higher cognitive reflection may be more likely to engage rational decision-making strategies in order to consider the risk and consequences of replying to an email, rather than basing judgments on an intuitive, reward-based, response. Based on their relation to dual-system approaches to decision-making, and previous research, the following hypotheses were generated about these two variables:

H³: Higher levels of inhibition will predict higher accuracy on the email legitimacy task.

H⁴: Participants who demonstrate higher levels of cognitive reflection will also demonstrate higher accuracy on the email legitimacy task.

The notion of intuitive behaviour resulting in a decision to take an immediate pay out over a more worthwhile long-term decision in the gambling task above highlights additional variables that may have links with email response behaviour. As discussed in Chapter 1, theories of self-control suggest that crime victims are more likely to engage intuitive, rather than rational decision-making strategies (Gottfredson & Hirschi, 1990; Schreck, 1999). It is suggested that this relationship comes from a lack of consideration of the negative consequences of a decision from those with lower self-control. Holtfreter et al. (2010) supported this in research demonstrating that lower self-control related to higher susceptibility to offline fraud. This research has clear parallels to online fraud, and findings were mirrored in relation to internet consumer fraud (van Wilsem, 2013). A lack of consideration for

the consequences of a decision also brings attention to a concept known as need for cognitive closure, which describes a person's desire to reach a conclusion to a problem or decision as soon as possible (Kruglanski, 1990). Similar to theories surrounding self-control, this would suggest that a user may rely on system 1 decision-making processes so that they can reach a conclusion quicker, rather than taking the time to consider potential consequences rationally. In both of these variables, the desire for an immediate behavioural response may make it difficult to employ rational decision-making strategies, similar to the effects of inhibitory capacity and cognitive reflection. Based on this, the following hypotheses were generated:

H⁵: Participants who demonstrate higher levels of self-control will perform with higher accuracy on the email legitimacy task.

H⁶: Lower need for cognitive closure will predict higher accuracy on the email legitimacy task.

An alternative approach to considering the links demonstrated in gambling task scenarios such as the one described above, is to consider the intuitive response as a demonstration of risk-taking behaviour. In this sense, it is possible that the desire to engage with such risks could lead to increased susceptibility. Hoyle et al. (2002) report sensation seeking as being dispositional to risk-taking behaviour due to a desire for novel experiences and sensations. In relation to email fraud, if there were a relationship, it may be due to awareness that responding to an email involves taking a risk and this is something that appeals to some users. Alternatively, it may be that some users have an unconscious disposition to more impulsive behavior.

Sensation seeking has been linked to impulsive behaviour (Whiteside, Lynam, Miller, & Reynolds, 2005), and therefore may also be considered as a predictor of susceptibility in the same manner as other variables associated with reliance on intuitive, impulsive responses.

Sensation seeking and impulsivity have been linked to increased extraversion (Whiteside & Lynam, 2001), one of the big five factors of personality. However, previous research has demonstrated that in fact *lower* levels of extraversion and higher levels of agreeableness predict past response behaviour relating to email fraud (Modic & Lea, 2011). It is suggested that individuals with higher levels of agreeableness are more trusting and likely to believe what they are told, whilst those with lower levels of extraversion are likely to build stronger relationships online, leading them to be more trusting of people interacting with them in this environment. This research contradicts findings relating to sensation seeking, but the direct link to email fraud provides a more grounded foundation for predictions about future findings. Based on previous research both from the wider field and relating directly to email fraud, the following hypotheses were generated:

H⁷: Participants who demonstrate lower levels of sensation seeking behaviour will perform with higher accuracy on the email legitimacy task.

H⁸: Lower levels of agreeableness and higher levels of extraversion will predict higher accuracy on the email legitimacy task.

In summary, a suite of psychological measures are administered in this experiment, to establish a profile of the cognitive make-up of the most susceptible users, as well as beginning to consider the situational factors and email content that

may affect judgments. It is felt that these variables cover a range of perspectives in considering how approaches to email decision-making may affect accuracy, including risk-taking and control over impulsive responses. All of the cognitive variables discussed above are measured using established procedures. However, there are no comparable standardised protocols for measuring email fraud susceptibility.

Different methodologies have been employed to measure susceptibility in past work – such as scale measures, email judgment tasks, and work with past victims - each of which, as noted in Chapter 1, has its own advantages and disadvantages.

A judgment task was developed for this experiment, in which participants rate how confident they are that each email they see is either phishing or legitimate, from a corpus of 36 emails (18 phishing, 18 legitimate). There are broad similarities with Yan and Gozu (2012) in using an email judgment task. However, in this experiment, both phishing and legitimate emails are presented (rather than just phishing emails). This provides a degree of uncertainty about authenticity, which is surely true for everyday behavior, and provides a point of comparison between the emails. In addition, we ask participants to make judgments on a six-point scale, rather than a binary choice (as in Yan & Gozu, 2012). Potentially this better reflects real-world thinking – where someone may be either cautious or definitive about authenticity – and provides a richer, more graded set of data on performance.

In order to measure susceptibility, two approaches will be taken in this experiment, as this is the first time that the task is being used. An accuracy score will be calculated incorporating ability to detect both phishing and legitimate emails. Whilst phishing emails are deemed more risky than legitimate emails, the consequences of dismissing a legitimate email as phishing should not be

underestimated. Legitimate emails can contain important information about system and account updates that may leave personal information vulnerable if ignored. In addition to this approach, signal detection theory is considered as an alternative assessment of accuracy. Using this method, two scores are calculated - one which considers the distribution of 'hits' (when a phishing email is judged as such) and 'false alarms' (when a legitimate email is judged as phishing), and another which estimates the amount of bias in a participants responses. The calculation of bias will demonstrate whether participants' judgments are affected by their knowledge that they are looking to identify phishing emails. That is, are participants more biased towards the 'phishing' response as they are acting more cautiously in their decision-making than they would in real life? Signal detection theory also provides a useful assessment when considered manipulations such as time pressure, as it provides insight into the decision-making process employed by participants (MacMillan, 2002). This will highlight whether time pressure leads to participants making more mistakes generally or whether this has a greater impact on ability to detect phishing emails rather than legitimate emails.

Email decision-making, as measured by the legitimacy task, will be considered with respect to self-control, working memory, cognitive reflection, inhibition, personality, sensation seeking, and cognitive closure given the direct or indirect empirical and conceptual justifications for each, as discussed above. These variables have not been considered together, nor have they all been evaluated using a common performance task, making the current study unique.

3.2 Method

3.2.1 Participants. A sample of 224 participants, consisting of Lancaster University students and staff, comprised 67 males and 156 females (in one case, gender was not specified). Participants ranged in age from 18 to 30 years, with a mean age of 19.37 ($SD = 1.69$). All participants completed the email legitimacy task and the email usage questionnaire; they also completed one of two sets of cognitive measures - either set 1 ($N = 120$) or set 2 ($N = 101$), which are detailed below.

3.2.2 Design. A between-participants manipulation of time pressure (time pressure vs. no time pressure) was included on the email legitimacy task. Outcome measures from this task (number correct, confidence score, D-prime) act as the dependent variables in this study, whilst the cognitive measures (Cognitive Reflection Test, International Personality Item Pool, Brief Need for Closure scale, Flanker Task, Brief Self-Control scale, Stroop task, Brief Sensation Seeking scale, and Reading Span task) and time pressure manipulation are the independent variables.

3.2.3 Materials.

Email usage questionnaire. To establish participants' familiarity with email management, a questionnaire was developed (shown in Appendix H) to identify some general demographic information, as well as details of time spent on the internet, email volume, and phishing email frequency. As part of the questionnaire, participants were asked to report whether they had ever responded to a phishing email.

Email legitimacy task. Participants judged 36 emails, of which 18 were legitimate and 18 were phishing. Most of these emails were received by the research team, with a few harvested from internet sites describing examples of phishing emails. Examples of the stimuli used can be found in Appendix I. Attempts were made, where possible, to match the phishing and legitimate emails by including examples of each from the same types of companies with similar requests, such as password resets or transaction confirmations. Direct comparison was not possible for all emails in the stimuli set, so 30 individuals were recruited to rate the legitimacy of the emails and how obvious each was as an example of phishing or a genuine email, prior to the experiment. This ensured that the stimulus set incorporated both more and less clear-cut examples of legitimate and phishing messages, so all of these stimuli were incorporated into the final version of the task.

The recipient details were modified on all emails so that the format of the address and subject information was the same throughout. The content of the emails themselves, the subject title, and the sender address were not modified. Participants received the following instructions regarding the relevance of the email stimuli:

During this part of the study you will be asked to indicate, on the piece of paper in front of you, whether you think each of the emails shown is a legitimate or a phishing email.

A phishing email refers to a malicious communication sent from somebody posing as someone they are not which aims to elicit personal information from the user such as usernames, passwords, or bank account details, or to encourage the user to download a file which contains a virus that will infect their computer and could lead to personal data being stolen.

For each of the emails that you will see during this task, you are to assume that all communications are relevant to the recipient unless it is obvious otherwise. For example, assume that the recipient does have an account at the bank that an email is from, that they do hold an eBay account, etc. Please indicate on the scale in front of you how confident you are that the email on screen is either phishing or legitimate.

Fifteen order presentations were randomly generated for the emails so that this could be varied between participants. For each email, participants were asked to identify, on a 6-point Likert scale, how confident they were that it was phishing or legitimate, with 1 being 'definitely phishing' and 6 being 'definitely legitimate'. Emails were presented using Microsoft PowerPoint, with responses made on a printed scale in front of the participant, which had 4 stimuli per page (to limit answer review strategies or response pattern detection).

Time constraints were varied between-participants. Participants were either told that they only had five minutes to complete the entire task (time pressure condition), or that they should complete the task at their own pace (no time pressure condition). To provide motivation and incentivise performance during the email task, all participants were also informed of an optional prize draw for three cash prizes based on the most accurate scores.

Cognitive tasks. Each participant completed one of two sets of cognitive tasks, shown in Table 3.1, as it was not feasible to administer all tasks to participants without inducing fatigue effects. Each set comprised four tasks, including both questionnaire and behavioural measures, covering a spectrum of potential psychological constructs, with each set being of equivalent overall duration. All

questionnaire tasks involved a 5-point Likert scale, allowing participants to give a neutral response (3) if they were unsure of how to respond or how the statement related to them. All measures were programmed and administered through PsychoPy (Peirce, 2009).

Table 3.1 *Outline of the tasks to be used in this study*

| Cognitive task | Factor being measured |
|-------------------------------------|---------------------------------|
| <i>Set 1</i> | |
| Cognitive Reflection Test | Cognitive reflection |
| International Personality Item Pool | Big Five factors of personality |
| Brief need for closure scale | Need for closure |
| Flanker test | Inhibition |
| <i>Set 2</i> | |
| Brief self-control scale | Self-control |
| Stroop test | Inhibition |
| Brief sensation seeking scale | Sensation seeking |
| Reading span task | Working memory span |

The Cognitive Reflection Test (Frederick, 2005) is a 3-item measure of impulsivity and cognitive reflection, outlined in more detail in Study 1.

The 50-item version of the International Personality Item Pool (IPIP; Goldberg, 1999; shown in Appendix J) provides a widely used measure of the big-five factors of personality. Example items for each factor are: extraversion, e.g. *I don't talk a lot*; agreeableness, e.g. *I feel little concern for others*; conscientiousness, e.g. *I leave my belongings around*; neuroticism, e.g. *I get stressed out easily*; and intellect, e.g. *I am quick to understand things*. This break down of personality factors means that specific predictors of susceptibility can be assessed. Once specified items are reversed, a sum of the participant's responses is taken as a measure of personality.

This can then be broken down into the five factors of personality, again with a sum of the responses to questions in each subset being taken as the score.

The brief need for closure scale (shown in Appendix K) consists of a 15-item questionnaire measure of cognitive closure developed by Roets and van Hiel (2011) based on the original need for closure scale (Webster & Kruglanski, 1994). By using the brief version of the scale we are reducing the time needed to complete the study, whilst still using an empirically validated measure. Examples of items from the scale include: *I dislike questions which could be answered in many different ways; I feel uncomfortable when I don't understand the reason why an event occurred in my life; and When I am confronted with a problem, I'm dying to reach a solution very quickly.* A mean score across all items is derived as a measure of cognitive closure.

The Flanker test measures inhibition by assessing a participant's ability to suppress distracting response cues. This task was originally published by Eriksen and Eriksen (1974) and has been used in many varieties since then. For this study we used arrows stimuli (pointing left or right), with the direction of the central arrow being either congruent or incongruent with the surrounding four arrows (as outlined by Fan, McCandliss, Sommer, Raz, & Posner, 2002). Neutral trials were also included, in which the central arrow was surrounded by four square shapes. There were six possible combinations of stimuli: two congruent, two incongruent, and two neutral. Participants completed an initial 24 practise trials to familiarise them with the task, following which a further 192 trials were completed over four blocks, separated by short breaks. Response times were recorded, with a mean difference between response time on congruent and incongruent trials, where the target item was correctly identified, being taken as a measure of the participant's performance.

The brief self-control scale (Tangney et al. 2004; shown in Appendix B) is outlined in more detail in Study 1. The Stroop test (Stroop, 1935) provides a behavioural measure of inhibition. In this version of the task, participants saw a colour word on screen and responded to the colour of the font of the word, rather than the colour described by the word itself. In doing this, participants are required to suppress the word stimuli that the image in front of them induces in order to respond with the font colour. The task comprised congruent (word and font colour match) and incongruent (word and font colour are different) trials. There were a total of nine possible word and colour combinations. Font and word colours were red, green, and blue, with the corresponding keys being r, g, and b, all of which are close together on the keyboard so should not have caused confusion for participants trying to find the correct key. Participants completed an initial 27 practise trials to provide task familiarisation, followed by 144 trials, split into two blocks with a short break in between. As with the Flanker test, a mean difference score between response times on the congruent and incongruent trials with correct responses was taken as a measure of performance. A higher difference score would indicate that a participant takes longer to differentiate between congruent and incongruent stimuli, thus suggesting lower levels of inhibition.

The brief sensation-seeking scale (Hoyle et al., 2002; shown in Appendix L) provides an 8-item measure of risk-taking behaviour. This brief version of the scale is derived from the original 40-item sensation-seeking scale developed by Zuckerman, Eysenck, and Eysenck (1978). Some examples of items from the scale include: *I would like to take off on a trip with no pre-planned routes or timetables; I like to do*

frightening things; and I get restless when I spend too much time at home. A mean score of responses across the eight items yields a scale score.

Finally, the reading span task (Daneman & Carpenter, 1980) assesses working memory span. Participants read a set of independent sentences and afterwards recall the last word from each. Testing followed a conventional approach. Initially, participants saw three trials with two sentences and subsequently recalled the two sentence end words. Participants also indicated after reading each sentence whether they thought it made sense or not; some sentences had the final word switched with another sentence to encourage comprehension. Providing the participant recalled two out of three trials correctly then sentence length increased by one. Testing cycled in this way until recall dropped below two out of three trials correct, or reached the maximum level of six sentence trials. A corpus of 60 sentences taken from Daneman and Carpenter (1980) was used as stimuli. Participants also had a practice phase comprising three sets of two sentences for familiarisation. The span size for each participant was calculated based on the highest number of sentence end words that participants were able to recall. For example, if a participant was able to completely recall two out of three sets of four words or more, but were unable to when shown five sentence end words in a trial, then their span size would be four.

3.2.4 Procedure. Participants signed up for this study on the University Sona system. After giving consent, participants were first asked to complete the email usage questionnaire, followed by the email legitimacy task either with an induced time pressure or without. In the time pressure condition participants saw on-screen instructions, and were also reminded by the researcher that they only had five

minutes to complete the task, and thus should work through the emails as quickly as possible. Those who were not in the time pressure condition were told to complete the task at their own pace. All participants were given task instructions, and information about the cash prizes available for those participants who performed best on the task. Inclusion in the prize draw was optional - participants who wished to participate were required to leave a contact email address at the end of the email task. Emails were presented on an Apple iMac using Microsoft PowerPoint, with responses recorded on paper.

Upon completing the email legitimacy task, participants were then asked to complete one of the two sets of cognitive tasks. Task order was varied 4-ways between participants to avoid order effects. Instructions were displayed for each task before the participant was asked to complete it and, where appropriate, practise trials were performed to help participants understand the tasks. Once all of the tasks were completed, participants were debriefed and thanked for their time. The data was saved with the participant's unique identification number, and transferred from the computer to an encrypted hard drive in preparation for analysis.

3.2.5 Data collation. Responses to the email usage questionnaire were close-ended and so inputted accordingly. In order to score the email legitimacy task, each email decision was identified as correct or incorrect, and correct responses were summed to derive a total number correct for each participant. In order to generate this binary response, the 6-point scale was divided at the mid-point, meaning that responses between 1-3 were classed as a judgment of 'phishing' and

responses between 4-6 as 'legitimate'. A confidence score was also generated, which comprised the sum of the difference between scale rating and the true veracity for each email, thus reflecting extreme vs. hedged responses. This score provides a complimentary perspective to the binary scores, demonstrating a participant's self-confidence in their ability to accurately judge email legitimacy, rather than simply their overall accuracy. Table 3.2 provides an overview of the scoring system used to generate the overall confidence score, which includes negative scoring for incorrect responses. This scoring method means that confidence is assessed in relation to whether it is validated in the accuracy of the response given. When generating the confidence score for each email type separately, only the positive end of the scoring system was employed. The maximum possible score for overall confidence was 108, whilst the maximum possible score for phishing and legitimate emails individually was 54.

Table 3.2 *Scoring system for confidence scores on email judgment task*

| | Scale response | | | | | Definitely legitimate |
|-------------------|------------------------|----|----|----|----|--------------------------|
| | Definitely phishing | | | | | |
| | 1 | 2 | 3 | 4 | 5 | 6 |
| Phishing emails | 3 | 2 | 1 | -1 | -2 | -3 |
| Legitimate emails | -3 | -2 | -1 | 1 | 2 | 3 |

In addition, an alternative measure of accuracy was calculated on the email legitimacy task – using signal detection parameters. This provides a measure of sensitivity to demonstrate how well a participant can discriminate phishing from legitimate emails, and whether they demonstrate a bias in their responses. Using the binary responses from the email task, D-prime scores (d') and response bias (c) were calculated using formulae outlined by Macmillan (1993). This uses the 'hit' rate (the

proportion of phishing emails correctly identified) and the 'false alarm' rate (the proportion of legitimate emails incorrectly judged as phishing). A high, positive d' score suggests a participant is discriminating well between the two stimuli types, whilst a low, negative number suggests they are poor at discriminating and show low accuracy. For response bias, a value greater than 1 indicates bias towards 'legitimate', whilst a value less than -1 indicates bias towards 'phishing'. Inclusion of this analysis approach will allow for comparison with the number correct and confidence score calculations outlined above, so the most useful technique for this newly developed task can be established for later use in the thesis.

Data from each of the cognitive measures was collated according to previous literature to establish a score for each. In the cognitive reflection test, there were eight cases of data failure, so for this task, $N = 112$. Performance on the Flanker and Stroop tasks is based on the mean difference between (correct) response time on congruent and incongruent trials. To screen for outliers, z scores were generated for trial response times and any values below -2.56 or above 2.56 were removed before the mean values were extracted. When testing began, a technical error meant that data was not collected for the Flanker task for the first 18 participants, so for this task, $N = 102$. In addition, one further case was removed from the Flanker task analysis, as this was an extreme outlier compared with other participants' response times ($z = -8.51$). Reading span was calculated based on the highest sequence of sentence-end words correctly recalled. For example, if a participant was able to completely recall two out of three sets of four words or more, but were unable to when shown five words at a time, then their span size would be four. A technical error made it impossible to distinguish participants who reached span size six or just

size five as the last set of sentences were not recorded. However, as only five per cent of participants reached at least span size five, this score coarseness should not be a major issue.

3.3 Results

3.3.1 Email legitimacy task. Since the email legitimacy task was generated for the purpose of this thesis, patterns in detection ability are assessed first. This will highlight any emails in the stimuli set that were difficult for participants to judge. In addition, the manipulation of time pressure in the email task will be examined to assess whether increased pressure to complete the task in a given time frame impacts performance.

Figure 3.1 demonstrates performance on each email from the stimuli set. Participants judged each email on a scale from 1 (definitely phishing) to 6 (definitely legitimate) and the means displayed show the average score given to each email across all 224 participants. From this, three phishing and two legitimate emails were identified with mean response judgments that lie on the wrong side of the mid-point (as identified by the dashed line). There are also a number of emails which have mean scores near to the extremes of the scale and very small standard deviation values, suggesting that they were easily identified as either phishing or legitimate.

Each participant's data were also examined for any evidence of response bias, for example participants being over cautious and marking the majority of emails as phishing, or being under cautious and marking most as legitimate. There were an equal number of phishing and legitimate emails in the task, so based on scale responses we would expect that someone performing well on the task would have a




Figure 3.1 Diagram to show mean rating across participants for each email stimulus

mean between 3 and 4, with a large standard deviation to show that they were using the full scale when judging emails. We judged a mean value below 2.5 or above 4.5 as evidence of bias. Using these criteria, 8.48% of participants showed a bias. Of these, 42% were biased towards ‘definitely phishing’. Based on response bias calculated as part of the signal detection analysis, which uses binary response data, 4.9% of participants were found to demonstrate a bias, of which 36.4% were biased towards ‘phishing’ responses.

Table 3.3 shows the means and standard deviations for each of the performance measures on the email task, including binary correct or incorrect response, and a sensitivity score generated using signal detection theory, as well as

the confidence score. A number of one-sample t tests showed that participants performed better than chance on the total number of correct emails ($t(223) = 28.66$, $p < .001$, $d = 1.97$), on phishing emails ($t(223) = 15.32$, $p < .001$, $d = 1.02$) and on legitimate emails ($t(223) = 15.41$, $p < .001$, $d = 1.03$).

Table 3.3 *Descriptive statistics for measure calculated from email legitimacy task*

| Measure | Mean | SD |
|---|-------|-------|
| Number correct (/36) | 24.57 | 3.43 |
| Number of phishing emails correct (/18) | 12.09 | 3.02 |
| Number of legitimate emails correct (/18) | 12.48 | 3.38 |
| Confidence score | 33.08 | 17.14 |
| Confidence score on phishing emails | 28.91 | 9.06 |
| Confidence score on legitimate emails | 27.23 | 9.85 |
| D-prime | 1.09 | 0.60 |

Table 3.4 reports the correlations between all email legitimacy measures. It is evident that confidence strongly correlates with accuracy, as measured by both the number correct and D-prime. This pattern is also observed more modestly, for phishing and legitimate emails separately. The other notable finding is that all measures which separate the phishing and legitimate emails produced negative correlations between the two email categories. This would suggest that people perform better at detecting phishing *or* legitimate emails, rather than accuracy on one type predicting high accuracy on the other type.

A series of independent t -tests assessed the effect of time pressure on email decision accuracy. Participants in the time pressure condition ($M = 24.12$, $SD = 3.33$) accurately identified fewer emails than those without time pressure ($M = 25.06$, $SD = 3.49$), $t(222) = 2.05$, $p < .05$, $d = 0.28$. A significant difference was also found for confidence scores, $t(222) = 2.39$, $p < .05$, $d = 0.32$. Participants in the time pressure

Table 3.4 *Correlations between measures from the email legitimacy task*

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|-----------------------------------|-------|--------|--------|-------|--------|-------|------|
| 1. No. Correct | 1.00 | | | | | | |
| 2. No. phishing correct | .46** | 1.00 | | | | | |
| 3. No. legitimate correct | .61** | -.43** | 1.00 | | | | |
| 4. Confidence score | .94** | .50** | .51** | 1.00 | | | |
| 5. Confidence score on phishing | .45** | .91** | -.36** | .58** | 1.00 | | |
| 6. Confidence score on legitimate | .53** | -.35** | .85** | .57** | -.18** | 1.00 | |
| 7. D-prime | .97** | .40** | .63** | .92** | .41** | .56** | 1.00 |

Note. ** $p < .01$ (two-tailed)

condition ($M = 30.47$, $SD = 16.69$) performed with less confidence than those who completed the task with no time pressure ($M = 35.88$, $SD = 17.25$). Signal detection analysis also demonstrated that participants in the time pressure condition were worse at discriminating phishing from legitimate emails ($M = 1.00$, $SD = 0.56$) compared to participants with no time pressure ($M = 1.19$, $SD = 0.63$), $t(222) = 2.38$, $p < .05$, $d = 0.32$.

In order to assess whether time pressure affected accuracy in judging phishing and legitimate emails differently, a repeated measures ANOVA was run on the number of emails accurately identified in each of these categories, with time pressure as a between subjects factor. This showed no significant difference between the number of phishing and legitimate emails correctly identified, $F(1, 222) = 1.15$, $p = .29$, $\eta^2 = .01$, and no interaction between performance on each set of stimuli and the time pressure element, $F(1, 222) = 0.09$, $p = .76$, $\eta^2 = .00$. Similarly, no significant difference was found between confidence on phishing and legitimate emails, $F(1, 222) = 2.93$, $p = .09$, $\eta^2 = .01$, and there was no interaction between confidence and the time pressure element, $F(1, 222) = 0.05$, $p = .83$, $\eta^2 = .00$.

3.3.2 Email usage questionnaire. Percentage responses for each of the questions included in the email usage questionnaire are shown in Table 3.5. A multiple linear regression analysis evaluated whether internet behaviour predicted how at risk people reported feeling to online fraud. The best-fit model was found to include only Question 1, 'How many hours do you spend actively using the internet each day?' ($\beta = 0.20, p < .05$), with more hours spent online indicating higher feelings of risk. This model produced an R^2 value of .02 (adjusted $R^2 = .02, F(1, 222) = 4.90, p < .05$).

A further multiple linear regression analysis was conducted to test whether responses to the email usage questionnaire predicted performance on the email legitimacy task. None of the questions were found to significantly predict performance based on number of emails correctly identified, with the best-fit model including only Question 3, 'How many emails do you receive on an average day?' ($\beta = -0.28, p = .22$). This model produced an R^2 value of .01 (adjusted $R^2 = .00, F(1, 222) = 1.51, p = .22$).

3.3.3 Cognitive measures - Set 1. Before assessing performance on each of the cognitive task sets individually, independent samples t-tests explored whether there was underlying difference in email performance between the sets. No significant differences emerged for the number correct ($t(218) = 1.82, p = .07, d = 0.25$) the confidence score ($t(218) = 1.47, p = .14, d = 0.20$).

Descriptive statistics for each of the cognitive tasks in this group are shown in Table 3.6. Additionally, correlations between the tasks and the measures of performance on the email legitimacy task are shown in Table 3.7. These correlations

Table 3.5 *Percentage responses to questions on the email usage questionnaire*

| Question | 0-1 hours | 1-3 hours | 3-6 hours | 6+ hours | I do not use the internet on a daily basis |
|--|-----------|-----------|-----------|----------|--|
| 1 - How many hours do you spend actively using the internet each day? | 0.4 | 30.8 | 50.4 | 18.3 | 0 |
| | 0-20% | 20-40% | 40-60% | 60-80% | 80-100% |
| 2 - What proportion of that time is spent reading and responding to email correspondences? | 68.8 | 28.6 | 1.8 | 0.9 | 0 |
| | 0-5 | 5-10 | 10-15 | 15-20 | 20+ |
| 3 - How many emails do you receive on an average day? | 13.4 | 46.4 | 27.2 | 6.7 | 6.3 |
| 4 - How many phishing emails do you receive in an average week? | 64.7 | 15.6 | 6.3 | 7.1 | 5.4 |

suggest a significant relationship between scores on the cognitive reflection test and performance on the email legitimacy task, as well as between agreeableness and confidence scores.

Backward multiple regression analysis established whether the cognitive variables predicted email judgments. Regressions were run separately for number correct, confidence, and D-prime score and are summarised in Table 3.8. Using the number of emails correctly identified as the dependent variable, the model that explained the most variance included the cognitive reflection test ($\beta = 0.97, p < .01$), Flanker task ($\beta = 0.01, p < .05$), and agreeableness ($\beta = -0.08, p = .21$) as predictors of

Table 3.6 *Descriptive statistics for cognitive tasks in set 1*

| Cognitive measure | N | Mean | SD |
|-------------------------------------|-----|--------|--------|
| Cognitive reflection test | 112 | 1.22 | 1.18 |
| International personality item pool | 120 | 172.44 | 18.75 |
| Extraversion | 120 | 32.35 | 7.67 |
| Agreeableness | 120 | 40.58 | 6.00 |
| Conscientiousness | 120 | 33.87 | 6.54 |
| Neuroticism | 120 | 29.63 | 8.66 |
| Intellect | 120 | 36.03 | 5.29 |
| Flanker test | 101 | 4.58 | 104.77 |
| Need for closure scale | 120 | 3.13 | 0.55 |

susceptibility. This model produced an R^2 value of .16 for the model (adjusted R^2 = .14, $F(3, 91) = 5.96$, $p < .01$). These predictors demonstrate that a higher level of cognitive reflection and inhibition predicted better performance on the email task.

When confidence score was taken as the dependent variable, the same three variables produced the best model - cognitive reflection test score ($\beta = 3.68$, $p < .05$), Flanker test score ($\beta = 0.04$, $p < .05$), and agreeableness ($\beta = -0.56$, $p = .07$). This model produced an R^2 value of .16 (adjusted R^2 = .13, $F(3, 91) = 5.70$, $p < .01$). Again, this demonstrates that higher cognitive reflection and inhibition predicted better performance.

When D-prime was taken as the dependent variable, similar patterns were found, with the best-fit model including the cognitive reflection test score ($\beta = 0.15$, $p < .01$), Flanker test score ($\beta = 0.00$, $p < .05$), and agreeableness ($\beta = -0.02$, $p = .17$). This model produced an R^2 value of .14 (adjusted R^2 = .11, $F(3, 91) = 5.03$, $p < .01$). This indicates that participants with higher cognitive reflection were better at discriminating between phishing and legitimate stimuli.

Table 3.7 Correlations between cognitive tasks in set 1 and email legitimacy task scores

| Task | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
|------------------------------|-------|-------|-------|--------|--------|------|------|--------|------|------|------|
| 1. No. Correct | 1.00 | | | | | | | | | | |
| 2. Confidence Score | .93** | 1.00 | | | | | | | | | |
| 3. D-prime | .97** | .91** | | | | | | | | | |
| 4. Cognitive reflection test | .28** | .21* | .26** | 1.00 | | | | | | | |
| 5. Extraversion | -.12 | -.12 | -.13 | -.07 | 1.00 | | | | | | |
| 6. Agreeableness | -.15 | -.20* | -.17 | -.27** | .31** | 1.00 | | | | | |
| 7. Conscientiousness | -.05 | -.01 | -.07 | -.14 | .01 | .18* | 1.00 | | | | |
| 8. Neuroticism | .08 | .04 | .06 | .06 | .18 | -.02 | .07 | 1.00 | | | |
| 9. Intellect | .01 | -.01 | .00 | .03 | .12 | .12 | .20* | .07 | 1.00 | | |
| 10. Flanker test | -.17 | -.20* | .14 | -.12 | .17 | .12 | -.01 | -.09 | .02 | 1.00 | |
| 11. Need for closure | -.03 | .01 | -.06 | -.10 | -.27** | -.15 | .12* | -.30** | -.03 | -.02 | 1.00 |

Note. **p < .01, *p < .05 (two-tailed)

Table 3.8 *Summary of multiple regressions analyses for cognitive variables in set 1 predicting email task behaviour*

| | Number correct | Confidence score | D-prime |
|---------------------------|----------------|------------------|---------|
| Cognitive reflection test | .31** | .25* | .28** |
| Flanker test | .23* | .25* | .20* |
| Agreeableness | -.13 | -.18 | -.14 |
| Adjusted R ² | .14 | .13 | .11 |

Note. Parameter estimates are standardised coefficients.

* $p < .05$, ** $p < .01$ (one-tailed)

3.3.4 Set 2. Table 3.9 reports descriptive statistics for each of the cognitive tasks, while Table 3.10 reports correlations between the tasks and email judgments. These show a significant negative relationship between email accuracy and sensation-seeking scores.

Table 3.9 *Descriptive statistics for cognitive tasks in set 2*

| Cognitive measure | N | Mean | SD |
|-------------------|----|-------|-------|
| Self-control | 97 | 38.87 | 8.33 |
| Sensation seeking | 97 | 3.19 | 0.74 |
| Stroop task | 99 | 93.81 | 77.28 |
| Reading span | 99 | 2.12 | 1.53 |

Again, backward multiple regression analysis was employed to identify which variables best predict email task performance. These models are summarised in Table 3.11. When the dependent variable was the number of accurate email judgments, the best-fit model included sensation seeking ($\beta = -1.01$, $p < .05$) as the only predictor. This model produced an R^2 value of .05 (adjusted $R^2 = .04$, $F(1, 99) = 5.37$, $p < .05$). Sensation seeking was also the only predictor in the best-fit model for confidence score, ($\beta = -4.67$, $p < .05$). This model produced an R^2 value of .04 (adjusted $R^2 = .03$, $F(1, 99) = 4.09$, $p < .05$). Both analyses demonstrate that those

Table 3.10 *Correlations between cognitive tasks in set 2 and email legitimacy task scores*

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|----------------------------|-------|-------|-------|------|------|------|------|
| 1. No correct | 1.00 | | | | | | |
| 2. Confidence score | .95** | 1.00 | | | | | |
| 3. D-prime | .98** | .93** | 1.00 | | | | |
| 4. Self-control score | .09 | .08 | .08 | 1.00 | | | |
| 5. Sensation-seeking score | -.23* | -.20* | -.21* | -.11 | 1.00 | | |
| 6. Stroop test | .02 | .01 | .00 | -.00 | -.08 | 1.00 | |
| 7. Reading span | .01 | .03 | .00 | .09 | .09 | -.09 | 1.00 |

Note. ** $p < .01$, * $p < .05$ (two-tailed)

who performed better on the email task demonstrated less risk-taking behaviour on the sensation seeking scale.

When D-prime was considered as the dependent variable, the best model included only sensation seeking, but this was not found to be a significant predictor ($\beta = -0.13$, $p = .09$). The model produced an R^2 value of .03 (adjusted $R^2 = .02$, $F(1, 95) = 2.86$, $p = .09$). The direction of this finding indicates that participants who demonstrated higher sensation seeking were poor at discriminating between phishing and legitimate stimuli.

Table 3.11 *Summary of multiple regression analyses for cognitive variables in group 2 predicting email task behaviour*

| | Number correct | Confidence score | D-prime |
|-------------------------|----------------|------------------|---------|
| Sensation seeking score | -0.23* | -0.20* | -0.17 |
| Adjusted R^2 | .04 | .03 | .02 |

Note. Parameter estimates are standardised coefficients.

* $p < .05$ (one-tailed)

3.4 Discussion

This study set out to establish whether there are links between individual differences in email judgments and a range of psychological constructs. The findings indicate that sensation seeking, inhibition, and cognitive reflection are significant predictors of accuracy in email judgments, which was used as a measure of susceptibility. In other words, data offer evidence that there are psychological signatures of who is susceptible to making erroneous email judgments, though this at best offers a partial explanation of performance. These markers are considered in more detail in the discussion that follows.

Based on research linking sensation seeking with risky decision-making (Hoyle et al., 2002) and impulsivity (Whiteside et al., 2005), Hypothesis 7 predicted that higher sensation seeking scores would correlate with lower accuracy in the email task. The results confirm this. In relation to risky decision-making, it may be argued that at the point users become victims of phishing they are not necessarily demonstrating risky behaviour - if they genuinely believe that an email is from the sender it purports to come from (i.e. riskiness may be an attribution given with hindsight). Therefore, risky behaviour, and sensation seeking in relation to this, may be more relevant to certain types of phishing emails than others – specifically, those in which the user is given the chance to take part in a prize draw or other form of competition, compared to those which purport to come from a bank or reputable company regarding account access. Although formal analyses would be required to investigate this, an initial informal observation suggests that differences in sensation seeking do not predict ability to recognise one example of a more risky, prize draw email ($p = .93$) any better than a more subtle, transaction confirmation email from a

reputable company ($p = .87$). Impulsivity may pose a more substantial explanation of sensation seeking and susceptibility, with those participants demonstrating higher impulsivity (through the measure of sensation seeking) performing with lower accuracy. In line with approaches discussed below regarding intuitive and rational responses, impulsivity may demonstrate a further reliance on the initial, intuitive response to an email, with little consideration of the consequences.

Inhibitory capacity was measured with the Flanker and Stroop tests. Both require the participant to suppress an initial impulsive response in order to reach the correct response regarding a set of visual stimuli with congruent and incongruent trials. Hypothesis 3 proposed that demonstrations of lower inhibition would predict lower accuracy on the email task. This was supported to some extent, with performance on the Flanker task predicting response accuracy. In addition, Hypothesis 4 proposed that better performance on the Cognitive Reflection Test, a measure of impulsivity through the requirement to suppress an intuitive response in order to identify the correct answer, would predict increased accuracy on the email task. This was supported in the current experiment. Although evidence is not conclusive about the exact processes behind these relationships, these findings support the relevance of dual-system processing theories in email decision-making. As discussed earlier, the factors highlighted as predictors of susceptibility have been associated with a reliance on intuitive processing styles. This may suggest that those participants with lower inhibitory capacity and cognitive reflection were more reliant on immediate, intuitive responses during the decision-making process, which can lead to more erroneous judgments due to missed information (Tversky & Kahneman, 1975).

It is important to note that in both task sets, the amount of variance explained by the regression models and its predictors was relatively modest. This indicates that there is more to explaining vulnerability than these constructs alone. It could be that these measures only partially, and imperfectly, capture the underlying constructs being targeted. However, it is important to also consider additional factors, outlined in the introduction (section 3.1), which may influence susceptibility – such as email content and situation. As demonstrated in the emails that were easily identified, one such additional explanation may be the familiarity of the email to the user – either with the sender or with the format and content of the email.

This study manipulated time pressure on the email task as a situational constraint, with participants either working under limited time constraints or told to complete the task at their own pace. This manipulation aimed to emulate a scenario where users are under pressure to complete tasks in a limited time frame whilst also managing emails, or where a sense of urgency is induced in an email to encourage an urgent response from the recipient. Hypothesis 1 predicted that the induced time pressure would impair performance on the email task. This was supported, as participants in the time pressure condition demonstrated lower accuracy compared to those participants who were given no time limit. Yet, whilst significant, the effect size is relatively small, implying a modest influence at this level of contrast. In addition, this effect may vary between emails, with some requiring more consideration and contemplation than others. Further data about the response times to each email would allow detailed examination of this, to establish how time spent on the email affects response accuracy. In addition, the consideration of alternative

situational influences would be beneficial to further understand decision-making processes. Such influences will be explored more in Study 7.

Given that the email legitimacy task was developed for the purpose of this thesis, some initial consideration of its feasibility as a measure of susceptibility was conducted. Examination of responses to the email task found that most participants made good use of the full response scale whilst completing the task. In addition, analysis of the individual emails supports the initial pilot data in Study 2a that our stimuli set contained a range of obvious and less obvious examples of both phishing and legitimate emails. This analysis also highlighted three phishing and two legitimate emails that were consistently judged on the wrong side of the mid-point by participants. The phishing emails were from well-known companies with all of the same logos and formatting that would appear on legitimate emails. The two legitimate emails were poorly formatted and unprofessional looking emails from well-known organisations. Whilst incorrect identification of legitimate emails (a false positive) is potentially less damaging to the user in real-world scenarios, it remains noteworthy as an index of judgment difficulty.

As well as those emails that were commonly misjudged, there were also some emails that were consistently judged correctly by participants. The phishing emails that were easily identified (email numbers 4, 7, and 14) were either advance fee fraud, or prize draw scams. The legitimate emails which were easily identified (email numbers 23, 25, 27, and 29) were either from well-known companies familiar to most people, or internal emails at Lancaster University, which the participants could have received personally and recognised during the task. In both groups, the common factor was the familiarity of the email – either familiarity with the company

that it purported to come from, or with the type of scam and format of the email.

Whilst there are some consistent patterns in how participants judged certain emails, it may inevitably be the case that interpretation of stimuli and motivation during the task differed between participants, and thus affected the study outcomes.

In order to measure susceptibility, this study incorporated both raw accuracy scores, as well as signal detection parameters. Whilst the signal detection theory approach concentrates on ability to recognise phishing emails, with legitimate emails judged as phishing considered 'false alarms', this does not necessarily reflect the complex nature of email management. The consequences of misidentifying a phishing email can be substantial, both economically and psychologically, but it should also be considered that dismissing a legitimate email as phishing can also have negative consequences. For example, if an important email about a security update is dismissed as phishing, this may result in a users system becoming unsecure – leaving them more vulnerable to attack. The findings from this study demonstrate that the two measures of performance on the email task – accuracy and D-prime (taken from signal detection theory) – produce similar findings, with the same cognitive predictors being demonstrated for each. Therefore, given the additional value of correct judgments of legitimate emails in email management, it was decided that further analysis of the email legitimacy task throughout the thesis will focus on the overall accuracy and confidence measurements, rather than signal detection theory.

Whilst no extensive literature on the psychological traits associated with susceptibility has been found whilst completing this project, the results from this study do not replicate findings from the few previous studies that are related. Modic

and Lea (2011) concluded that agreeableness and extraversion were linked to decision-making surrounding fraudulent communications, whilst Holtfreter et al. (2010) found self-control to be a significant predictor of susceptibility. Although agreeableness was included in one regression model, this was not a significant predictor – and so none of these findings were replicated in the current experiment. Methodological differences may partially explain these discrepancies, as alternative judgment tasks were used in each – with Holtfreter et al. focusing on offline fraud scenarios. The traits that lead to increased chance of victimisation may vary depending on the medium in which the fraud is delivered. Furthermore, constant changes in the techniques being employed by perpetrators mean that alternative psychological variables may be relevant to the different decision-making processes required to detect these.

The nature of the email legitimacy task means that there are a number of methodological limitations associated with this, which will be discussed at various points throughout the thesis. The main issue though is that the task is a lab-based method, which requires participants to actively differentiate between phishing and legitimate emails. This is not necessarily representative of real world email management behaviour, where users interact with their inbox, are familiar with the companies they usually have contact with, and will know whether an email is relevant to them. However, the task provides an ethically sound measure of ability to distinguish between phishing and legitimate emails, and it is hoped that the inclusion of a prize draw for performance provided some incentive for participants to do well at the task. Analysis demonstrated that most participants were utilising the whole

response scale, suggesting that they were engaging with the task, rather than simply clicking the same response for every answer.

3.4.1 Conclusion. This study identifies a number of psychological markers of susceptibility – both cognitive and situational. However, there remain several key research questions or issues to address. In particular, further investigation into the utility of the email legitimacy task as a measure of real-world susceptibility would be beneficial both to the validity of results in this project, but also in developing this as a tool for other researchers. Moreover, given the level of unexplained predictive variance within the data from both sets of cognitive tasks, it is important to investigate where this variance might be sourced. For example, further analysis of additional psychological constructs and alternative measures of susceptibility may both offer complementary perspectives. Some constructs were not predictive of susceptibility; it is not clear whether they are simply not relevant, or whether their contribution may be masked by other variables. Therefore, further assessment of their relevance would help advance explanations of susceptibility.

Nonetheless, this study gives an outline of a cognitive profile of those who are most at risk alongside other factors that may be relevant, such as time pressure, and the content of the emails themselves. This provides an initial insight into individual differences in email fraud susceptibility, but also brings up a number of other research questions, which the following studies in the thesis will aim to address.

Chapter 4

The impact of methodological variations on decision-making in the email legitimacy task

Chapter summary

The study reported in this chapter outlines a variation on the email legitimacy task used in Study 3, in order to explore how different proportions of phishing and legitimate emails in the stimuli set affect accuracy. In part, this is an attempt to emulate a real inbox, in which there would be more legitimate than phishing emails. The same battery of cognitive variables was included here as in Study 3, in order to demonstrate whether prior findings would be replicated with this alternative version of the email task. Findings demonstrated that performance across both varied proportion (majority phishing and majority legitimate) conditions was lower than Study 3. In addition, findings regarding the cognitive variables predicting email task performance in Study 3 were not replicated here.

4 Study four

4.1 Introduction

In order to assess susceptibility to email fraud, an email legitimacy task was employed in Study 3 (described in the previous chapter), which contained an equal mix of phishing and legitimate email stimuli. In the current chapter, a study is reported that aims to understand how manipulating this proportion of phishing to legitimate emails affects accuracy in judging the legitimacy of the stimuli. That is, are judgments of email stimuli affected by the other emails that are presented alongside them? Previous studies using a similar judgment task to measure susceptibility have taken different approaches, either using equal proportions (Parsons et al., 2013), only phishing emails (Yan & Gozu, 2012), or a majority of phishing emails (Hong et al., 2013) in their stimuli set. The different methodologies employed mean that it is not possible to compare participant performance between these studies though.

In day-to-day email management, users are unlikely to receive fifty or one hundred per cent phishing emails. Despite this, previous research has not considered how methodological decisions regarding the sample of stimuli included in an email judgment task affect the decision-making process. In order to develop the most effective, and yet practical, method of measuring susceptibility in a lab setting it is important to understand how participants respond when the task content is manipulated.

This study includes two variations on email proportions – one where the participant sees a majority of legitimate emails (designed to be more akin to a real inbox), and one where the participant sees a majority of phishing emails. This will provide some insight into the decision-making process employed by participants in

responding to email stimuli. As suggested earlier, the nature of the 'line-up' task, with emails judged one after the other, means that there may be some comparison drawn across emails, which aid judgments. If this comparison is based on an expectancy to see an equal number of each stimuli type, then the impact of an expectation violation may affect accuracy. The majority of studies in which participants are required to categorise two stimuli types include equal proportions of each. Therefore, a scenario in which the expectation of this distribution is violated may impair performance. For this reason, accuracy on the email task is compared with data from Study 3, where participants saw fifty per cent phishing emails, in order to establish whether accuracy is impacted by varying the proportions of each email type.

Alternatively, performance may be affected by the patterns of detection that a user is accustomed to in day-to-day email management. Most users will receive more legitimate than phishing emails (The Radicati Group, 2015), and so the condition in which participants see a majority of legitimate emails is more representative of a genuine inbox. If this more realistic representation of email stimuli does affect accuracy, then a difference between the two varied proportion conditions would be expected. As stated above, the effect of email distribution is a concept that has not previously been considered in relation to its impact on decision-making. The motivation to develop this variation on the email task was to make the stimuli set more representative of a real inbox though, so based on this, the following hypothesis was generated:

H¹: Participants will perform with higher accuracy when the stimuli set for the email legitimacy task contains a majority of legitimate emails, akin to a real

inbox, compared to the majority phishing condition, and the equal proportions condition in Study 3.

Alongside the email task, this study will also assess the same cognitive measures as in Study 3, which outlined cognitive reflection, inhibition, and sensation seeking as predictive factors in performance on the email legitimacy task. Although these were the only factors that were found to be significant, all eight measures were administered again in this study to assess (i) whether the previous findings are replicable and (ii) whether any other apparent relationships emerge with a different variation of email stimuli. Of particular interest are the factors that are highlighted in previous literature as predicting susceptibility, but were not replicated in Study 3. For example, Modic and Lea's (2012) finding that agreeableness and extraversion, of the big five factors of personality, are associated with susceptibility, as well as Holtfreter et al.'s (2010) findings that self-control levels link to offline fraud susceptibility. The same hypotheses about the direction of these relationships are considered here as in Study 3, based on previous research.

4.2 Method

4.2.1 Participants. This study consisted of 98 participants, of which 83 were female and 15 male. All were first year psychology students at Lancaster University. The age of participants ranged from 18 to 48 years, with a mean age of 19.22 years (SD = 3.80).

4.2.2 Design. This study incorporates a between participants manipulation on the email legitimacy task (majority phishing vs. majority legitimate emails). In this study, the outcome measures for performance on the email legitimacy task act as dependent variables, whilst the proportion of emails seen and cognitive measures (Cognitive Reflection Test, Brief Need for Closure scale, International Personality Item Pool, Flanker Task, Brief Self-control scale, Brief Sensation Seeking scale, Stroop Task, and Reading Span Task) are considered as independent variables.

4.2.3 Materials.

Email usage questionnaire. The email usage questionnaire (shown in Appendix H) was the same as that used in Study 3. This provided data on participants' day-to-day usage and familiarity with emails, and experience of responding to phishing emails in the past.

Email legitimacy task. The design of the email task remained predominantly the same as in Study 3, with email stimuli displayed in a random order to avoid any order effects, and participants being asked to rate, on a 6-point Likert scale, how confident they were that each email was either phishing (1 – definitely phishing) or legitimate (6 – definitely legitimate). The stimuli in this version included the same 36 emails stimuli as used previously, but with an additional set of 18 emails to allow for the varied proportions of phishing and legitimate emails. However, each participant still only saw 36 emails. In the majority phishing condition, participants saw 27 phishing and 7 legitimate emails. In the majority legitimate condition, participants saw 27 legitimate and 7 phishing emails.

As in Study 3, participants were offered the chance to enter a prize draw for one of three cash prizes based on highest accuracy on the email task. The aim of this was to motivate participants to perform well in the email task.

Cognitive tasks. The same sets of tasks as previously outlined in Study 3 were included here, with each participant completing one set of tasks. However, data from the reading span task will not be included in the analysis due to a recording error. As previously, the measures were kept in the same two sets, outlined in Table 4.1.

Table 4.1 *Sets of cognitive measures*

| Set 1 | Set 2 |
|-------------------------------------|--|
| Need for closure scale | Brief self-control scale |
| Cognitive reflection test | Brief sensation seeking scale |
| International personality item pool | Stroop task |
| Flanker task | Reading span task (data not in analysis) |

4.2.4 Procedure. Participants completed this study online after signing up to one of the four conditions (without knowing the difference between these) via the University participant recruitment system. The conditions are outlined below in Table 4.2.

Table 4.2 *Outline of conditions*

| Cognitive task set | Email stimuli condition | N |
|--------------------|-------------------------|----|
| Set 1 | Majority phishing | 25 |
| Set 2 | Majority phishing | 24 |
| Set 1 | Majority legitimate | 24 |
| Set 2 | Majority legitimate | 25 |

The initial part of the study was presented through qualtrics.com, an online survey-building site. Participants gave consent for this study by clicked a statement

to acknowledge that they had read and understood information about their right to withdraw, anonymity, and confidentiality of data. Only once participants had acknowledged this statement were they able to proceed to the next screen (example of online consent form used can be seen in Appendix M). At the next screen, participants were asked to complete the email usage questionnaire. This asked them for some basic demographic information such as age and gender, as well as information about their internet and email usage. At the end of the questionnaire, participants had the opportunity to leave a contact email address for inclusion in the prize draw for performance on the email task. Participants were then automatically redirected to another web page, developed using PsyScript online software (Slavin, 2014), where they completed the remainder of the tasks. This software is only compatible with Google Chrome and Safari web browsers, which participants were informed of when signing up for the study. Participants first completed the email legitimacy task. All participants judged the legitimacy of 36 emails in total, with some participants shown more phishing emails, whilst others saw more legitimate emails. Full instructions were provided on how to complete the task, as outlined in Study 3. Emails were displayed one at a time, once a participant had selected their response on the scale; they were then shown the next email.

Upon completion of the email task, participants completed the set of cognitive tasks allocated to their condition. The order that the tasks were completed in was randomised to avoid any order effects that may bias results. Each task was preceded by a full set of instructions and in some cases, with the behavioural tasks, included practise trials to help participants understand the task. Upon completion of all three elements of the study, participants were shown a debrief sheet on screen.

Data from the various tasks was downloaded and collated, with the participant's identification number being the same across all in order to match the sets of responses. This was then stored on an encrypted hard drive in preparation for analysis.

4.2.5 Data collation. Performance on the email legitimacy task was measured using number correct and confidence score as outcome measures, as described in Study 3. Similarly, the email usage questionnaire and cognitive measures were scored in the same way as previously described. As previously, performance on the Flanker and Stroop tasks was measured by the mean difference between (correct) response time on congruent and incongruent trials. To screen for outliers, *z* scores were generated for trial response times and any values below -2.56 or above 2.56 were removed before the mean values were extracted. For the Flanker task, one further case was removed from the analysis, as this was an extreme outlier (*z* = -6.57) compared to response times recorded for other participants.

4.3 Results

4.3.1 Email legitimacy task. Data from all of the conditions were collated to monitor average accuracy on the email task. When taking number correct as the measure of susceptibility, the mean score was 22.69 (SD = 4.73; theoretical maximum of 36), whilst a mean confidence score of 23.03 (SD = 23.21) was generated. The two measures of susceptibility from the email task, number correct (measuring behavioural accuracy) and confidence score (measuring meta awareness), are highly correlated, $r(96) = .96, p < .001$.

In order to establish whether the proportions of email types affected accuracy, independent t-tests were conducted to measure the difference between the two conditions. No significant difference in performance was found when either the number of correctly identified emails ($t(96) = 0.81, p = .42, d = .16$) or the confidence score ($t(96) = 0.95, p = .34, d = .19$) was used as the measure of susceptibility. When compared to data from Study 3 on email task accuracy, where participants saw equal numbers of phishing and legitimate emails, a one-way ANOVA demonstrated that the number of emails correctly identified was found to be significantly different ($F(2, 319) = 8.51, p < .001, \eta^2 = 0.05$). Mean values for each of the conditions are displayed in Table 4.3. Tukey post-hoc analyses demonstrated lower accuracy in the majority phishing condition than in the equal mix, $p < .05, d = .38$. Similarly, participants performed with lower accuracy in the majority legitimate condition compared to the equal proportions condition, $p < .01, d = .52$. There was no significant difference between the more phishing and more legitimate conditions though, $p = .58, d = .16$.

The same analysis was run using the confidence score, which again demonstrated a significant difference between the three conditions, $F(2, 319) = 10.03, p < .001, \eta^2 = 0.06$. Mean values for each condition are displayed in Table 4.3. Tukey post-hoc analysis demonstrated that participants performed with lower confidence in the majority phishing condition than in the equal mix, $p < .05, d = .41$. Again, participants performed with higher confidence in the equal proportions condition than with more legitimate emails, $p < .001, d = .57$. No significant difference was found between the more phishing and more legitimate conditions, $p = .75, d = .19$.

Table 4.3 *Descriptive statistics for performance on the email task across conditions*

| Stimuli condition | Accuracy (number correct) | | Confidence score | |
|----------------------------|---------------------------|------|------------------|-------|
| | Mean | SD | Mean | SD |
| Majority phishing | 23.08 | 4.38 | 25.27 | 21.11 |
| Majority legitimate | 22.31 | 5.06 | 20.80 | 25.15 |
| Equal proportions (Exp. 3) | 24.57 | 3.43 | 33.08 | 17.14 |

4.3.2 Email usage questionnaire. Data was collated across the two email task conditions from this study to measure whether responses to the questionnaire were predictive of a number of factors: perceived risk, past response, and accuracy on the email legitimacy task. Backward multiple linear regression analysis demonstrated that for perceived risk, no significant predictors were found from the email usage questions, with the most predictive model containing only question 4 – ‘How many phishing emails do you receive in an average week?’ - ($\beta = .18, p = .08$), generating an R^2 value of .03 (adjusted $R^2 = .02, F(1, 96) = 3.11, p = .08$). When considering past response likelihood, again the best model contained only question 4, which was found to be a significant predictor ($\beta = -.24, p = .02$), producing an R^2 value of .06 (adjusted $R^2 = .05, F(1, 96) = 5.97, p = .02$). This finding suggests that participants who are used to receiving more phishing emails were less likely to have responded to a phishing email in the past. When assessing how predictive the email usage questionnaire was of accuracy on the email task using the number of emails correctly identified, no significant predictors of performance were found, with the best model containing only question 2 ($\beta = -.17, p = .10$), generating an R^2 value of .03 (adjusted $R^2 = .02, F(1, 96) = 2.80, p = .10$).

4.3.3 Cognitive measures – Set 1. Table 4.4 shows the means and standard deviations for performance on each of the cognitive tasks in this group. Correlational analyses between all of the cognitive measures and performance on the email task demonstrated no significant relationships, $r_s < .15$ and $> -.15$.

Table 4.4 *Descriptive statistics for cognitive measures in set 1*

| Cognitive measure | Mean | SD |
|-------------------------------------|--------|-------|
| Need for closure scale | 3.22 | 0.59 |
| Cognitive reflection test | 0.80 | 1.06 |
| International Personality Item Pool | 162.33 | 15.78 |
| Extraversion | 31.12 | 6.88 |
| Agreeableness | 38.20 | 7.12 |
| Conscientiousness | 32.22 | 5.93 |
| Neuroticism | 26.98 | 5.15 |
| Intellect | 33.80 | 5.70 |
| Flanker task ($N = 48$) | 39.53 | 65.65 |

Backward multiple regression analyses found none of the cognitive measures to be significant predictors of performance. Using the number of emails correctly identified, the best-fit model included the need for closure scale ($\beta = -2.07, p = .09$) and neuroticism ($\beta = -0.25, p = .09$), generating an R^2 value of .09 (adjusted $R^2 = .05, F(2, 45) = 2.12, p = .13$). When taking the confidence score as the dependent variable, the best model again included the need for closure scale ($\beta = -9.17, p = .13$) and neuroticism ($\beta = -0.99, p = .16$), with an R^2 value of .06 (adjusted $R^2 = .02, F(1, 45) = 1.54, p = .23$).

4.3.4 Set 2. Means and standard deviations for each of the measures included in this group are shown in Table 4.5. Again, correlational analysis

demonstrated no significant relationships between the cognitive measures and performance on the email task, $r_s < .07$.

Table 4.5 *Descriptive statistics for cognitive measures in set 2*

| Cognitive measure | Mean | SD |
|-------------------|-------|-------|
| Self-control | 37.02 | 7.92 |
| Sensation seeking | 3.32 | 0.80 |
| Stroop task | 80.72 | 70.80 |

Backwards multiple regression analyses found none of the cognitive measures to be significant predictors of performance. Using the number correct as the dependent variable, the best-fit model has only the self-control scale as a predictor ($\beta = .06$, $p = .66$), producing an R^2 value of $< .01$ (adjusted $R^2 = -.02$, $F(1, 47) = 0.20$, $p = .66$). When confidence score was taken as the dependent variable, self-control was again the only predictor variable in the best model ($\beta = .07$, $p = .65$), producing an R^2 value of $< .01$ (adjusted $R^2 = -.02$, $F(1, 47) = 0.21$, $p = .65$).

4.4 Discussion

The aim of this study was to understand whether accuracy in email judgments was affected by the stimuli viewed alongside each email, when these consist of either more phishing or more legitimate emails. Performance was compared between the two conditions – majority legitimate and majority phishing – as well as being compared with performance in Study 3, where participants saw equal numbers of phishing and legitimate emails. Hypothesis 1 predicted that participants would perform with higher accuracy in the majority legitimate condition, as this was more representative of a real inbox and the decision-making processes employed in day-to-day email management. In addition to assessment of the email legitimacy task,

this study also aimed to establish whether findings regarding cognitive measures predicting susceptibility in Study 3 were replicated with an alternative version of the email task.

Results demonstrated that overall performance on the email task did not differ depending on whether the participant saw either more phishing or more legitimate emails. This suggests that attempting to make the email task more realistic of how a user is used to managing their own inbox, with a higher proportion of legitimate than phishing emails, does not affect accuracy. However, when compared to data from Study 3, where the task included fifty per cent phishing emails, participants were found to perform with lower accuracy and confidence. This may be due to a violation of participant expectation in terms of the amount of each type of stimuli included in the set, rather than the nature of between stimuli comparisons. In an experimental setting, participants often witness an equal split of stimuli when asked to make categorical decisions. It is possible that participants' responses are affected by this expectation, and so, consciously or subconsciously, they are monitoring the amount of responses they give for each category. If a participant were to receive the same emails in real life to their own inbox, in which they are used to receiving more legitimate than phishing emails, they may be less influenced by such expectations, and thus manage the emails without such bias. This provides opportunity for further research in which participants are naive to the nature of the study and are acting more naturally in terms of how they manage an email inbox, without any expectancy of differentiating between two types of email.

An alternative explanation considered was that the additional stimuli included in this study, compared to Study 3, were more difficult for participants to

accurately judge. Further examination of responses to the emails that were included in both Studies would highlight whether participants performed with lower accuracy across all stimuli in the experiment, or just on the novel stimuli.

Findings from the email usage questionnaire demonstrated that there was no difference in performance on the email task based on the amount of time participants spent on the internet, or the amount of emails they were used to receiving. However, participants who received less phishing emails on average were more likely to have responded to a phishing email previously. It is possible that these participants have less exposure to phishing emails, or believe that they do, and so are less familiar with scams and how to avoid these. There may also be an element of underreporting, in that these more susceptible participants do actually receive the same amount of phishing emails as other participants, but are unaware of the fraudulent nature of these and do not consider them a threat, so ignore them. An interesting strand of future research could consider voluntary monitoring of participants' email inboxes to compare the self-reported amount of phishing emails received, compared to the true number. This would provide valuable insight into how many phishing emails go undetected, whether these are responded to or not.

Unlike Study 3, none of the hypotheses relating to cognitive measures included were supported, with no evidence of any significant predictors of email task performance. One interpretation of this finding is that the predictive variables detected in the previous study were false positives and not actually predicting susceptibility. However, an alternative, perhaps more likely explanation is that the manipulation of the email task stimuli in this experiment, which was found to reduce accuracy, also affects the predictive influence of the cognitive factors measured.

Overall performance was found to be lower across the two conditions in this experiment, in comparison to Study 3. If this were due to a violation of what participants expect, this would suggest that the email judgment task is not measuring the same thing on these two occasions, and thus the reported relationship with the cognitive variables between the two studies is not comparable.

One additional factor to be considered in explaining the difference in performance on the email task between this and the study reported in Chapter 3 is the fact that this study was conducted online. It is possible that when completing research studies online, in their own time and space, participants are less likely to engage with the tasks at hand, and thus there is likely to be some variation in their performance. This explanation applies to all aspects of the study – both the email task and the cognitive measures. Therefore, this difference in experimental environment may explain the lack of replication between the two experiments. However, there is literature to suggest that patterns in data analysis do not differ between lab-based and online studies in decision-making (e.g. Germine et al., 2012; Klein et al., 2014). Therefore, we cannot conclude that this demonstrable variation in our results is due to the online nature of the study. Further investigation would be required to establish whether this lack of replication would also be found if the current study were to be conducted in a lab environment.

4.4.1 Conclusion. In summary, although the findings from Study 3 were not exactly replicated in this study, there are some valuable insights to be taken from the patterns of response data reported. From a methodological perspective, findings demonstrate that varying the expected proportions of two categories of stimuli can

impair accuracy. It is possible that the lack of replication between this and Study 3 is due to the measure actually being a more realistic measure of susceptibility to email fraud, and the cognitive factors reported having no relationship to susceptibility. However, given the lack of replication with any other previous research, such as Modic and Lea's (2011) findings that agreeableness and extraversion were linked to susceptibility, or Holtfreter et al. (2010) who demonstrated a link between offline fraud susceptibility and self-control, it is more likely that this task is simply a less valid measure of susceptibility.

In order to investigate the relationship between cognitive variables and email fraud susceptibility, it is important to establish the most reliable measure possible of this, within the ethical and practical scope of experimental research. Further testing would be necessary to establish the reliability and validity of this email task as a measure of susceptibility, as this is not necessarily representative of how participants would manage their own inbox, given the artificial nature of the task in a lab environment. Ideally, this would involve comparison between susceptibility as calculated on this task, and susceptibility based on a more realistic measure – such as a simulated phishing attack or role-play scenario where the participant is unaware of the nature of the study.

Chapter 5

Assessing the reliability of the email legitimacy task

Chapter summary

The study presented in this chapter assessed the test-retest reliability of the email legitimacy task. Performance is compared within participants across two test sessions using the varied proportions version of the email task. An additional cognitive measure – the Moses Illusion task – was also administered, which is considered in relation to performance on the email task. Findings demonstrate a strong relationship in performance between the two test sessions, suggesting good test-retest reliability for the varied proportions version of the email task. There was no evidence of a relationship between susceptibility, as measured by the email task, and performance on the Moses Illusion task.

5 Study 4a

5.1 Introduction

As discussed in Chapter 1, the literature to date on understanding the psychology behind email fraud susceptibility has employed a range of methods for measuring susceptibility. These approaches vary from questionnaires (e.g. Modic & Anderson, 2014) to simulated phishing attacks (e.g. Jagatic et al., 2005; Guéguen & Jacob, 2002). However, given the paucity of literature in this area, none of these methods have been extensively tested, and there is little evidence corroborating their validity or reliability.

As outlined in Study 3, an email legitimacy task was developed as a measure of susceptibility to be used for the majority of the thesis. This task allows for controlled assessment of how well participants can differentiate between legitimate and phishing emails, giving an indication of their likelihood to recognise phishing emails in their own inbox. In a continuation of this, the current study will allow for test-retest reliability analysis between two test sessions, in order to establish whether the task consistently measures susceptibility in the same way. Strong reliability will support the notion that responses on each occasion are not simply due to chance (Carmines & Zeller, 1979). In relation to email management, this assessment also provides insight into the nature of susceptibility. Chapter 1 outlines alternative approaches to understanding individual differences in susceptibility, which bring to light the question of whether some users are consistently more susceptible, or whether situational variables are a prominent cause. Strong reliability between two separate test sessions would suggest that susceptibility is a consistent

construct. The following hypothesis was generated to assess reliability of the email task:

H₁: Reliability analysis will demonstrate a good level of consistency ($r > .8$; Cicchetti, 1994) in performance on the email legitimacy task across test sessions.

Findings from Study 3 demonstrated the importance of certain cognitive variables in predicting performance on the email task – namely cognitive reflection, inhibition, and sensation seeking. These findings were not replicated in Study 4 though, which may have been due to methodological differences in the email task. This study includes another cognitive measure that was not previously included, allowing an initial assessment of the relevance of this task to susceptibility. It will then be included in further studies later in the thesis, which use alternative versions of the email task to determine whether findings are reliant on the specific email task employed. The Moses illusion task was originally reported by Erickson and Mattson (1981), but then expanded upon by Reder and Kusbit (1991) and consists of a set of questions, some of which are distorted and so have no sensible answer, but induce an intuitive response. The key example, after which the task is named, is the question “how many animals of each kind did Moses take onto the ark?”. This question is distorted because it was Noah who reportedly took animals onto his ark in the biblical story. Therefore, there is no sensible answer to the question, but an intuitive answer (‘two’) is induced in relation to the story relating animals and an ark.

These questions assess a participant’s ability to recognise the distortion and suppress the intuitive response. Preliminary research from a student project at Lancaster University demonstrated that participants who were more susceptible to

the Moses illusion (i.e. could not recognise the distortions) also demonstrated higher susceptibility on an email legitimacy task (Harrison, 2015), similar to that used in this thesis. One explanation given for the occurrence of the Moses illusion is that participants do not properly encode the information in a question, assuming that they know an answer from reading only part of the question (Park & Reder, 2004). In terms of real world email fraud victimisation, it is possible that users overlook some crucial cues in an email that indicate its legitimacy because they have already come to a conclusion based on partially encoded information. Similar to the Moses illusion, this means that they rely on an instinctive impression that the email is either phishing or legitimate without taking the time to consider additional information available to them that may contradict this first impression. Based on this, the following hypothesis was generated in relation to this task:

H₂: Participants who demonstrate high susceptibility to the Moses illusion, through not recognising distorted questions, will demonstrate lower accuracy on the email legitimacy task.

5.2 Method

5.2.1 Participants. A sample of 60 participants, consisting of 6 males and 54 females, with an age range between 18 and 29 years ($M=18.87$, $SD=2.28$) was collected from participants who had previously completed Study 4. Participants signed up for the study using the University's Sona system, where it was advertised with the criteria that participants must have completed the previous study in order to take part.

5.2.2 Design. This study incorporates a within subjects design, with all participants completing the email legitimacy task at two test sessions, allowing for

test-retest analysis. The Moses illusion will also be considered as an independent variable in this study, in relation to performance on the email legitimacy task as a dependent variable.

5.2.3 Materials.

Email legitimacy task. The same design was used for this email task as was employed in Study 4, in order to measure the reliability of this task. Across the two conditions in Study 4, a total of 54 emails were included, with each participant seeing 36 emails. During the current experiment, participants were shown all 54 emails and asked to rate each on a 6-point scale indicating how confident they were that the email was either phishing or legitimate, as in the previous studies. Given the anonymous nature of data collected in Study 4, the condition a participant was in was not recorded with their contact email address. Therefore, all participants were invited to respond to the 54 emails, providing their participant number from Study 4 so that it could be established which email stimuli they had previously judged.

To assess reliability, only the responses for the 36 emails that the participant had seen in the previous study were included – depending on the condition they were allocated to – so as to provide a direct comparison. As in the previous experiments, an accuracy score based on the number of emails correctly identified was calculated for each participant, as well as a confidence score based on the scale responses given.

Moses illusion. The Moses illusion task consists of a set of questions that allow the researcher to assess a participant's ability to detect distortions in information presented to them. In this study, we included ten questions making up

the Moses illusion, five of which were distorted. These can be seen in Figure 5.1 below. The distorted questions provide a measure of how well a participant can recognise a discrepancy and suppress an intuitive response. The questions were presented in a random order, and were then followed up by a set of five knowledge questions intended to clarify that the participant does know, for example, that it was Noah, rather than Moses, who took the animals onto the ark.

1. *How many animals of each kind did Moses take on the ark? (d)*
2. *What country was Margaret Thatcher prime minister of?*
3. *What did Goldie-Locks eat at the Three Little Pigs' house? (d)*
4. *What is the name of the shape whose area is "pi-r-squared"?*
5. *What English rock group did the late John Lennon sing with?*
6. *Who found the glass slipper left at the ball by Snow White? (d)*
7. *What phrase followed "To be or not to be" in Macbeth's famous soliloquy? (d)*
8. *Who was the first man to walk on the moon?*
9. *What is the name of the long sleep some animals go through during the entire winter?*
10. *What is the name of the molten rock that runs down the side of a volcano during an earthquake? (d)*

(d) indicates distorted question

Figure 5.1 Questions included in the Moses illusion task

In line with previous research, participants were given the following instructions that clearly told participants to answer 'can't say' to any question that seemed distorted:

You will now be asked to answer a series of general knowledge questions. You should treat each question literally and not give an answer to any question that seems distorted. If a question seems distorted, you should answer "can't say", and if there is a question that you do not know the answer to, you should answer "don't know".

5.2.3 Procedure. Participants signed up to this study via the University's Sona system, approximately 4 months after Study 4 was conducted. The advert on the system was set up in a way that meant only those participants who had participated in Study 4 were able to sign up. Once they had signed up, participants were given access to a web address that took them to the study. All of the tasks were completed online, via qualtrics.com and so were accessible via any web browser, unlike Study 4, which was restricted due the software used. Once participants had read the information sheet, displayed onscreen, and given consent, they completed the email task. As in the previous two studies, a scale was displayed for each email, with participants being asked to rate each email as to how confident they were that it was either phishing or legitimate. Full instructions were provided for participants to remind them how to complete the task, and also to inform them about the prize draw that was included as an incentive to perform at their best in the task. An option to leave an email address for this draw was given at the beginning of the experiment, once consent had been gained.

Following the email task, participants were asked to complete the Moses illusion task. The target questions were displayed first in a random order, followed by the knowledge questions intended to establish whether participants were aware of the correct answers to distorted questions. Upon completing all of the questions, a debrief sheet was displayed onscreen. Data was downloaded from Qualtrics on a daily basis and this was then transferred to an encrypted hard drive in preparation for data analysis.

5.2.4 Data collation. The email legitimacy task will be scored in the same way as studies 3 and 4, generating two outcome measures – number correct and confidence score. In order to score the Moses illusion task, each distorted question response was coded as either 1 – illusion occurred, whereby the participant gave the intuitive answer, or 0 – where the illusion did not occur and the participant gave the correct answer. If the participant gave an incorrect answer different to the intuitive response, this was also scored 0. In instances where the illusion occurred, a participant's response to the knowledge question relating to this were checked. If the participant demonstrated that they did not know the correct answer to this question, they would be given a score of 0 rather than 1, as the occurrence of the illusion may have been due to lack of knowledge rather than a reliance on the intuitive response. These scores were then totalled for each participant, and the total indicated their susceptibility to the Moses illusion, with a higher score indicating higher susceptibility.

5.3 Results

5.3.1 Reliability analysis. Performance on the email task was analysed using both binary correct or incorrect responses to stimuli, as well as confidence scores based on the scale responses provided by participants (as described in Studies 3 and 4). In this experiment, the mean number of emails correctly identified was 22.48 (SD = 5.28), whilst the mean confidence score was 21.90 (SD = 27.24). A simple correlation demonstrated a strong positive relationship between these two measures of susceptibility on the email task, $r = .95$, $p < .01$.

In order to assess the re-test reliability of the email legitimacy task across two test sessions, correlation coefficients were calculated for the number of emails correctly identified and the confidence score. Figures 5.2 and 5.3 demonstrate the spread of scores for each measure of susceptibility. For the number of emails correct, there was a significant positive correlation, $r = .63$, $p < .01$. A paired sample t-test confirmed that there was no significant difference between performance at test 1 ($M = 22.78$, $SD = 5.03$) and test 2 ($M = 22.48$, $SD = 5.28$), $t(59) = 0.52$, $p = .61$. For the confidence score, a significant positive correlation was found between the test sessions, $r = .69$, $p < .01$. Again, a paired samples t-test was conducted, which confirmed that there was no significant between scores at test 1 ($M = 23.23$, $SD = 24.89$) and test 2 ($M = 21.90$, $SD = 27.24$), $t(59) = 0.50$, $p = .62$.




Figure 5.2 Number of emails correctly identified at test session 1 and test session 2




Figure 5.3 Confidence score at test session 1 and test session 2

5.3.2 Moses illusion. The Moses task was scored using the number of times the illusion occurred for each participant based on five distorted questions, meaning the maximum possible score was 5, with a higher score indicated higher susceptibility to the illusion. The mean score for this sample was 2.10 (SD = 1.35). A simple linear regression analysis was conducted to establish whether performance on the Moses illusion task was predictive of susceptibility to email fraud, as measured by performance on the email legitimacy task. For the number of emails correctly identified, the Moses illusion was not found to be predictive of task performance ($\beta = .36, p = .48$), producing an R^2 value of .01 (adjusted $R^2 = -.01, F(1, 58) = 0.51, p = .48$). Similarly, with the confidence score as the dependent variable, the Moses illusion was not found to predict performance ($\beta = 2.46, p = .35$), producing an R^2 value of .02 (adjusted $R^2 = -.00, F(1, 58) = 0.88, p = .35$).

5.4 Discussion

The main aim of this study was to establish the test-retest reliability of the email legitimacy task, developed for use in this thesis as a measure of susceptibility to email fraud. Hypothesis 1 predicted that the task would demonstrate good reliability, above .8. Consistency in performance on the task across two test sessions, approximately four months apart, demonstrated a significant positive correlation, although this was not quite strong enough to confirm good reliability (Cicchetti, 1994). The direction of this finding suggests that susceptibility is consistent across different test sessions, but that there may be additional factors affecting accuracy that can vary between situations. Given that this task is newly developed and has not previously been tested for reliability, the findings provide a positive indication of the usefulness of this task as a lab based measure of susceptibility. Further assessment of individual email stimuli and differences in accuracy on the novel stimuli seen in this task may provide additional insight into how well the task measures susceptibility, and how it might be improved.

It may be beneficial to also assess the reliability of the email task when equal proportions of phishing and legitimate emails are used. Study 4 found lower accuracy on the varied proportions version of the email task, compared to the equal proportions version used in Study 3. This was thought to demonstrate the effect of a violation of observer expectancy – whereby participants would normally expect to see equal numbers of each stimuli type in a discrimination task. In this experiment, participants did see equal numbers of phishing and legitimate email stimuli though, and accuracy was akin to that in Study 4. Therefore, it is possible that the lower accuracy is due to the additional stimuli included in Studies 4 and 4a – these may be

more difficult for participants to accurately judge. In terms of reliability though, the findings from this study still suggest a positive relationship in performance across test sessions, a finding which should apply regardless of any variation on the exact stimuli included in the task.

Further assessment of the task would provide a more robust assessment of the task and what it is measuring. Tests of validity would ensure that the task is measuring susceptibility, rather than simply ability to differentiate between a sample of email screenshots. This could be attained by comparing performance on the email legitimacy task and participant likelihood to respond to a simulated phishing attack. This could also be measured by comparing performance on the email task between past victims of email fraud and a control group, however careful sampling would be required to ensure that members of the control group are actually less susceptible, as discussed in Chapter 1.

The Moses illusion task was also administered in this experiment. Based on exploratory data from a student project (Harrison, 2015), which suggested that this paradigm might relate to email task performance, Hypothesis 2 stated that higher susceptibility to the Moses illusion would predict lower accuracy on the email task. However, no significant relationship was evident. This may be due to different decision-making processes employed for each of these tasks. Previously in the thesis, differences in accuracy on the email task have been explained by dual-system reasoning, with the suggestion that those participants who demonstrate higher susceptibility are more reliant on intuitive responses to the email stimuli. The Moses illusion on the other hand may be explained by reliance on prior semantic knowledge, making some participants think they know the answer based on partially

encoded information, without reading the question in full. This theory of encoding behaviour can also be applied to phishing emails, with the suggestion that participants who demonstrate higher susceptibility may be basing their judgments on partial information, without taking in all of the cues available. However, the lack of relationship between the two tasks suggests that there are different decision-making processes involved in each. In order to ensure that the lack of relationship is because there are different cognitive processes involved in the performance on each task and not due to the specific version of the email task used in this case, it would be useful to assess performance on the Moses illusion with alternative measures of susceptibility. For example, in the original email task used in Study 3, with equal numbers of phishing and legitimate emails, related cognitive factors, such as cognitive reflection and inhibition, were found to be predictive of susceptibility. It may therefore be beneficial to look at performance on the original email task and the Moses illusion alongside these other cognitive variables to assess their predictive validity.

5.4.1 Conclusion. This study demonstrates a strong positive relationship between performance on the email task across test sessions, although this does not quite reach the level of good reliability. Further reliability and validity testing of this and the original version of the email task would establish how well these measure real world susceptibility. The lack of relationship between the email task and the Moses illusion task suggests that these employ different decision-making processes. Therefore, there is no evidence at this stage that performance on the Moses illusion could be utilised in future efforts to predict susceptibility. However, this task will be

assessed alongside the original email legitimacy task to ensure that the null finding is consistent across the task variations.

Chapter 6

Thinking aloud about email legitimacy judgments

Chapter summary

The study reported in this chapter employs a think-aloud protocol along side the email legitimacy task, providing qualitative data about the cues and heuristics that participants rely upon when making decisions regarding email legitimacy. The cognitive reflection test and Moses illusion were also included in this study to provide concurrent evidence of effects reported earlier in the thesis. Findings from the think-aloud data demonstrate a reliance on basic, out-dated cues in judging email legitimacy, as well as highlighting the persuasive power of factors such as familiarity and email relevance. These are valuable insights that may help to inform efforts to educate users and reduce susceptibility. The Moses illusion was found to be close to significance as a predictor of accuracy on the email task in this experiment, contrary to earlier null findings in Study 4a.

6 Study five

6.1 Introduction

In previous experiments, significant cognitive predictors of susceptibility – namely, cognitive reflection, inhibition, and sensation seeking - have been identified using an email legitimacy task as a measure of email behaviour. Whilst these findings provide valuable data about how individual differences can influence accuracy in email decision-making, the direct insight into *how* participants reach the decisions that they do about each email is fairly limited. Assumptions can be made about the cues utilised, based on information commonly portrayed to users about the generic traits of a classic phishing email, and from self-reports of previous email management behaviour, as demonstrated in Study 1. However, there is little past research reporting empirical evidence to support these assumptions. Therefore, we employ a think-aloud protocol in this study to provide a more in-depth account of how participants make judgments during the email legitimacy task.

The most recently published advice from the Citizen's Advice Bureau (2016) on how users can recognise a phishing attack is as follows:

- *The sender's email or web address is different to the genuine organisation's addresses*
- *The email is sent from a completely different address or a free web mail address*
- *The email does not use your proper name, but uses a non-specific greeting such as 'dear customer'*
- *The email threatens that unless you act immediately your account may be closed*
- *You're asked for personal information, such as your username, password or bank details*
- *The email contains spelling and grammatical errors*
- *You weren't expecting to get an email from the company that appears to have sent it*
- *The entire text of the email is contained within an image rather than text format*

- *The image contains a link to a bogus website*

This information addresses only the most basic phishing emails, relying on out-dated information that has long since been surpassed by the technological advance on the part of the fraudsters. Whilst alternative advice services, namely Action Fraud (2016), target some issues associated with more sophisticated phishing attacks, such as the use of spoofed email addresses, this still seems fairly limited. In addition, there appears to be little psychological evidence to demonstrate whether advice such as this is successful in helping users to detect phishing emails. A key aim of this study is to establish the extent to which users are reliant on out-dated cues such as these, which are no longer sufficient to protect them against more sophisticated phishing attacks.

One approach to this would be to adopt the design of the preceding experiments, but then follow up on the response to each email decision by asking participants why they judged an email as either phishing or legitimate. However, as suggested by van Someren, Barnard, and Sandberg (1994), it can be hard for subjects to give an explicit reason for a decision they have made. Giving the subject the opportunity to contemplate such a question may result in an inaccurate account of their decision-making process, as this is reliant on a memory of how the decision was reached. An alternative approach, which avoids the potential bias from participants interpreting their own cognitive processes, is to ask them to think aloud concurrently whilst completing a task. Originally emerging in the 1940s (e.g. Duncker, 1945; de Groot, 1946), the think-aloud protocol has become a widely accepted method for gathering data about information processing techniques employed by participants in reaching a decision, to enhance quantitative data regarding the decision itself. With

extensive guidelines about the utilisation of this method published by Ericsson and Simon (1993), the think-aloud protocol is used in a range of fields from medicine to human computer interaction, providing insight into various different task processes.

In relation to the field of cyber security, the think-aloud protocol has been employed to assess how participants make judgments about the legitimacy of various stimuli, as well as the effectiveness of training systems. Jakobsson, Tsow, Shah, Blevins, and Lim (2007) report findings from a think-aloud study in which participants rated the 'phishiness' of a set of email and webpage stimuli. They identified a number of cues utilised by participants in their decision-making process, including spelling, emphasis on security, poorly formatted URLs, relevance, personal approach, and security symbols. This work highlighted a number of cues that might now seem fairly out-dated with the relative sophistication of phishing emails (bearing in mind the study was published in 2007). Therefore, it is expected that the think-aloud protocol in this study will provide a contemporary assessment of the (self-reported) cues and heuristics that users rely on when judging the legitimacy of emails in our stimuli set, and whether this has changed and advanced with the simultaneous advancement of techniques employed by the fraudsters.

In situations where participants are uncertain of how to respond, they often rely on heuristics – utilising prior knowledge and experience to make a decision more rapidly, but not always more accurately (Tversky & Kahneman, 1973; Gilovich, Griffin, & Kahneman, 2002). Based on stereotypical cues to phishing, that are highlighted in many educational materials, it is possible that in the email legitimacy task participants will rely on a set of heuristics that do not fit with the more sophisticated techniques employed in current phishing attacks. This is of particular relevance with

the methodology used in this study, where participants may not always be familiar with the expected format of emails from a given company. In these cases, participants may fall back on heuristics that they associate with phishing to make a decision about legitimacy. An understanding of heuristics employed by users in this task will help to inform future educational efforts by outlining the cues that are often missed by participants, as well as the common misconceptions about reliable cues.

In addition to insight into the cues used by participants, the inclusion of a think-aloud protocol to the email legitimacy task will allow observation of the cognitive processes involved in the decision-making process. Study 3 outlined cognitive reflection, inhibition, and sensation seeking as predictors of accuracy on the email task. This was interpreted as a reliance on intuitive, system 1 decision-making processes, whereby these participants were unable to engage more rational System 2 reasoning. It is therefore predicted that a lack of contemplation, and consideration of solitary cues will be seen in those participants who are less accurate in their legitimacy responses.

The use of think-aloud protocols to assess cognitive processes has received some criticism. Whilst reporting of cues used in recognising phishing emails can be fairly objective, the interpretation of more subliminal cognitive processes involved in decision-making may be more subjective (van Someren et al., 1994). It has therefore been argued that think-aloud protocols are only effective when there is a solid theoretical underpinning to the research (Boren & Ramey, 2000). Within the current research domain – phishing emails - previous quantitative research from Study 3 has outlined cognitive predictors of accuracy on the email task. The think-aloud protocol will therefore offer a convergent approach to demonstrate that these cognitive

variables are evident in the decision-making process. It is hoped that, in combination with quantitative data reported throughout the thesis, this data will contribute to more evidence-based approaches to educating and training users in an effort to reduce susceptibility and future victimisation.

In addition to the qualitative data gathered in this study, performance on the email legitimacy task will be compared with that exhibited in previous studies in the thesis. It has been argued by some that concurrent thinking aloud whilst completing can result in changes of consciousness, which may impact performance by making participants more aware of the decision-making process (Smagorinsky, 1998). By comparing performance across studies in the thesis, the effect this has on accuracy can be assessed. Given the increased amount of attention that is likely to be drawn to certain elements of the emails due to speaking out loud about them, the following hypothesis was generated:

H¹: Participants will demonstrate higher accuracy on the email legitimacy task in this study compared to Study 3, where the same task was administered.

Alongside the email task, participants in this study will also complete two cognitive tasks. First - the cognitive reflection test (CRT; Frederick, 2005), which measures ability to reflect upon a problem, suppressing an impulsive response in order to attain the correct one. In Study 3 this was shown to be a significant predictor of susceptibility on the email task. However, this finding was not replicated in Study 4 when the proportions of phishing and legitimate emails in the legitimacy task were varied. Therefore, it has been included in this study to provide further data

on the relevance of this test as a predictor of susceptibility, as measured by the email legitimacy task:

H²: Higher levels of cognitive reflection will predict increased accuracy on the email legitimacy task.

Second - the Moses illusion (Erickson & Mattson, 1981; Reder & Kusbit, 1991) was included as part of this experiment. This task consists of a set of questions, some of which are distorted, designed to measure a participant's ability to recognise the distorted questions and suppress the intuitive response. Although there was no evidence of a relationship between this and email task performance in Study 4a, it is of interest to establish whether this null finding is replicated when the proportion of phishing and legitimate emails are equal. Based on previous findings from a student project (outlined in Study 4a), which demonstrated a link between this task and susceptibility, the following hypothesis was generated:

H³: Participants who demonstrate higher susceptibility to the Moses illusion will perform with lower accuracy on the email legitimacy task when there are an equal number of phishing and legitimate stimuli.

6.2 Method

6.2.1 Participants. A sample of 51 participants, all students and staff of Lancaster University, took part in this study, with 14 males and 37 females. There were 28 native English speakers in the sample, and 23 non-native speakers.

Participants ranged in age from 18 to 36 years, with a mean age of 23.86 (SD = 3.76).

6.2.3 Materials.

Email legitimacy task. This study will replicate the same email legitimacy task outlined in Study 3, as a measure of how accurately and confidently participants are able to differentiate between phishing and legitimate emails. The novel aspect in this study is that whilst completing the email task, participants were asked to think aloud in order to elicit verbal reports that potentially provide information about *how* people make decisions surrounding emails, as well as what decisions they make. In addition to the instructions provided for Study 3 (section 3.2.3.2) outlining the email task, the following written instructions were included to explain the think-aloud element of the task:

During this task, the voice recorder will remain on the desk to record your vocalisations. Whilst judging each email please say out loud everything that goes through your mind with regard to whether the email is phishing or legitimate.

This task was also explained to the participant verbally by the researcher. The emails were displayed in a random order during the task, with display order being automatically recorded so that the vocalisations could be linked to the appropriate email.

Cognitive tasks. In this experiment, participants were asked to complete the cognitive reflection test (Frederick, 2005; outlined in more detail in section 2.2.1.3) and the Moses illusion task (Reder & Kusbit, 1991; outlined in more detail in section 5.2.3). The cognitive reflection test provides a measure of a participant's ability to suppress an intuitive response in order to reach a correct answer. The Moses illusion recognises whether a participant is able to spot distortions in a set of questions and

acknowledge these, rather than providing an intuitive response that overlooks the distortion.

6.2.4 Procedure. Participants completed this study in a lab within the psychology department as they were required to think aloud and have their vocalisations recorded, meaning this could not be completed online. As in Study 4a, materials were presented via Qualtrics (www.qualtrics.com), and all participants completed the task on an Apple iMac. Once consent had been given, participants were asked to provide some basic demographic information, and were then asked to complete the cognitive reflection test. During this task, participants were asked to vocalise their thoughts – this was important to help participants get used to thinking aloud and feel more comfortable with the process (Ericsson & Simon, 1998). The recordings of this task were not including in later analysis. Following this, participants completed the email legitimacy task, again whilst being recorded with a voice recorder placed on the desk in front of them. The researcher was sat in the next room whilst participants completed the study so they could feel more comfortable thinking aloud without any experimenter bias. Upon completing the email task, the voice recorder was switched off before participants completed the final task - the Moses illusion questions. Once all of the tasks were completed, participants were shown an onscreen debrief sheet and given the opportunity to ask any questions that they may have had.

Audio recordings were transferred to an encrypted hard drive before being transcribed. All files were transcribed by a third party agency (www.uktranscription.com) due to the volume of data collected and time constraints

on the project. Once transcribed, the original audio files were destroyed (transcripts can be seen in Appendix N). Data recorded via Qualtrics was downloaded, and again stored on an encrypted hard drive in preparation for analysis.

6.2.5 Data collation. Outcome measures for the email legitimacy task were generated in the same way as previous studies, with a number correct and confidence score calculated for each participant. Similarly, the Moses illusion was scored as described in Study 4a.

Think-aloud data was transcribed and then coded for each participant to establish which section of text referred to which email, dependent on the random order in which these were presented. Once coded by email, each section of data was also coded as either phishing or legitimate, and as either correctly or incorrectly judged on legitimacy by the participant. Cues were identified across the following categories: correctly identified phishing emails, incorrectly identified phishing emails, correctly identified legitimate emails, and incorrectly identified legitimate emails. Further to the cues used to identify phishing emails, the transcripts were also examined for evidence on the cognitive processes involved in decision-making. Based on findings from Study 3, which showed cognitive reflection, inhibition and sensation seeking to be significant predictors of accuracy, differences in the presence of impulsive and risk-taking behaviour between participants were sought, to support these findings.

6.3 Results

6.3.1 Email legitimacy task. As in previous chapters, performance on the email legitimacy task is assessed through the number of emails accurately identified when scale scores were converted into binary responses, as well as a confidence score based on scale responses. The mean number of emails correctly identified was 26.84 (SD = 3.18; theoretical maximum score of 36), and the mean confidence score was 47.41 (SD = 16.94). Correlations between these, and accuracy in detecting phishing and legitimate emails individually are shown in Table 6.1. These demonstrate similar patterns to those shown in Study 3, with strong relationships between measures of accuracy and confidence, and negative relationships between measures of phishing and legitimate emails. Using the same criteria as outlined in Study 3, each participant's mean response score was examined to look for any biases in the use of the legitimacy scale. Any participant who had a mean score below 2.5 or above 4.5 was considered to show bias towards one end of the scale, suggesting they were either being under or over cautious in their responses. Using these criteria, 7.84% of participants were found to demonstrate bias, with all of these being towards the 'definitely phishing' end of the scale, demonstrating over caution in their responses.

A series of independent samples t-tests demonstrated that there was no difference in performance based on gender for either the number correct ($t(49) = 0.41, p = .68, d = .13$), or the confidence score ($t(49) = 1.06, p = .29, d = .33$). In addition, participants were asked in this study whether they were native speakers of English to establish whether this would affect performance. No significant differences in performance were found from either the number correct ($t(49) = 0.92,$

Table 6.1 *Correlations between performance measures on the email legitimacy task*

| | 1 | 2 | 3 | 4 | 5 | 6 |
|--------------------------|-------|--------|-------|-------|------|------|
| 1. Number correct | 1.00 | | | | | |
| 2. Phishing correct | .23 | 1.00 | | | | |
| 3. Legitimate correct | .77** | -.44** | 1.00 | | | |
| 4. Confidence score | .91** | .24 | .68** | 1.00 | | |
| 5. Phishing confidence | .20 | .82** | -.35* | .40** | 1.00 | |
| 6. Legitimate confidence | .59** | -.48** | .86** | .70** | -.19 | 1.00 |

Note. ** $p < .01$, * $p < .05$ (two-tailed)

$p = .36$, $d = .26$), or the confidence score ($t(49) = 0.09$, $p = .93$, $d = .03$).

As in Study 3, the mean score for each email was examined to look for emails that were scored on the incorrect side of the scale by a majority of participants. The same two phishing emails were found to fit these criteria as in Study 3 – one from Amazon that was correctly identified by 25 participants ($M = 3.68$) and one from PayPal that was correctly identified by only 9 participants ($M = 4.82$). Both followed the exact layout that would be expected from an email for the company, with correct formatting and logos. However, both came from incorrect email addresses, which was the principal cue that participants could have used to identify the fraudulent nature of these emails. The qualitative data associated with these emails are outlined below in section 6.3.3.3.

In order to establish whether the think-aloud protocol affected performance on the email task, an independent samples t-test was conducted between scores from this study and those from Study 3. A significant difference was found, with participants correctly identifying more emails in this study ($M = 26.84$, $SD = 3.18$) than in Study 3 ($M = 24.57$, $SD = 3.43$), $t(273) = 4.32$, $p < .001$, $d = .69$. In addition,

participants demonstrated higher confidence in this study ($M = 47.41$, $SD = 16.94$) than in Study 3 ($M = 33.08$, $SD = 17.14$), $t(273) = 5.40$, $p < .001$, $d = .84$.

6.3.2 Cognitive tasks. Participants in this study completed the cognitive reflection test ($M = 1.31$, $SD = 1.07$) and the Moses illusion ($M = 1.73$, $SD = 1.28$) in addition to the email legitimacy task. Backward multiple linear regression analysis demonstrated that when taking the number of emails correctly identified neither of these tasks was predictive of accuracy on the email task, generating a best-fit model that included only the Moses illusion ($\beta = 0.67$, $p = .06$). This produced a non-significant R^2 value of .07 (adjusted $R^2 = .05$, $F(1, 49) = 3.83$, $p = .06$). With the confidence score as the measure of susceptibility though, the best model was that containing just the Moses illusion, which was found to be a significant predictor ($\beta = 4.01$, $p < .05$), producing an R^2 value of .09 (adjusted $R^2 = .07$, $F(1, 49) = 4.98$, $p < .05$). This demonstrates that a higher score on the Moses illusion (indicating higher reliance on intuitive responses) predicted higher confidence on the email task.

6.3.3 Think-aloud analysis.

Construction of cue-type categories. Five cue types were consistently identified across all email responses, and so an informal content analysis was conducted for each cue type in each of the phishing/legitimate, correct/incorrect categories. Figure 6.1 shows the proportional occurrence of each cue type reported by participants whilst they were making decisions about email legitimacy, by category.

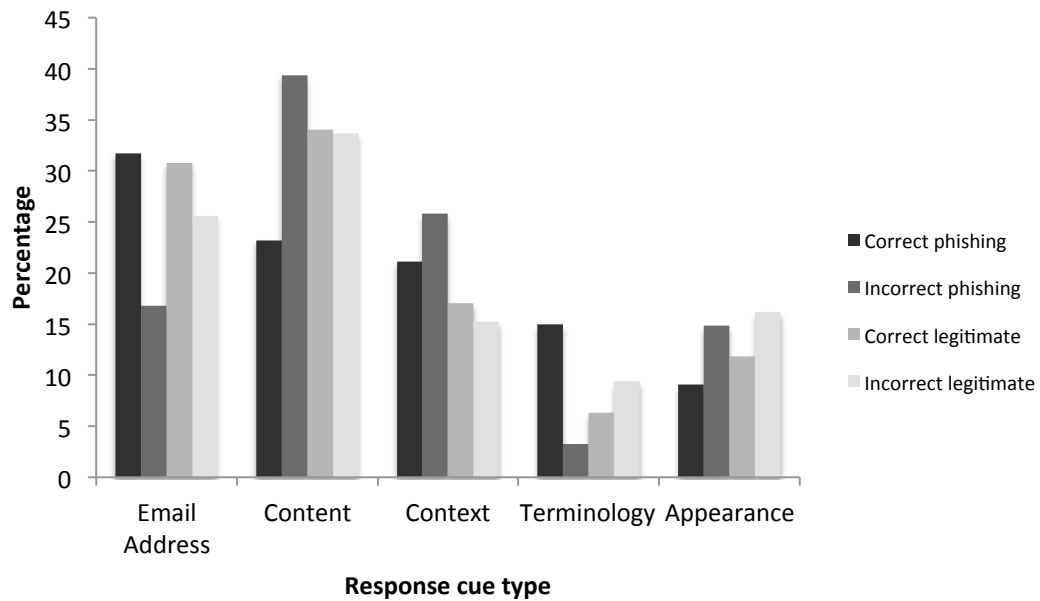


Figure 6.1 Frequency of reference to cue types in email decision-making

When participants erred, they believed legitimate emails to be phishing emails, and vice versa. Therefore, in the detailed descriptions of our cues to responses below, these are grouped together for correctly identified phishing emails with incorrectly identified legitimate emails, and correctly identified legitimate emails with incorrectly identified phishing emails.

Emails identified as phishing.

Email address. One of the most common cues used to identify phishing emails correctly, as shown in Figure 6.1, was the email address that the message was sent from. When identifying phishing emails, or misidentifying legitimate emails, participants often noted that the email address of the sender “doesn’t sound very official” (P302, Line 1), and “doesn’t look very trustworthy” (P308, Line 106). More specifically, when comparing the email address and any URL links within the email, participants would often notice the address “is not really related to the link that’s in

the email” (P305, Line 8), due to a discrepancy in the domain names. It is evident from the frequency of this cue being reported that it is seen as important to a lot of participants, with one noting, “if it wasn’t for the email address, I’d have probably said it was legitimate” (P312, Line 202).

Content of the email. Many participants considered the believability of what the email was asking of the recipient when judging the legitimacy of the emails. Based on their own knowledge of phishing emails, some participants noted for emails with URL links that “they would say ‘please log on’ or ‘please go to the website to change your password’ rather than following a link” (P336, Line 192). In addition, participants were wary of emails asking for money or personal information – “it asks you to give him the money so I would say it’s definitely phishing” (P307, Line 125); “it’s also asking [you] to fill out your information” (P304, Line 8). This kind of request, alongside a sense of urgency made participants particularly wary:

It is asking you to click there to put in banking details, which I’ve said I’m sketchy about, and the fact it says you have to do it within 72 hours or you’ll be suspended permanently. (P313, Line 131)

Those emails that contained URLs or attachments to download provided participants with further cues to educate judgements of legitimacy. Participants often reported that an attachment “doesn’t look right” (P310, Line 136) or a link “looks absolutely ridiculous. It just seems far too long” (P310, Line 73). As well as recognising there was something suspicious about such links and attachments, some participants went on to demonstrate an understanding of the consequences of such

risks, with comments such as “I’m certain [the links] will give me all kinds of viruses” (P306, Line 36).

A final interesting, and noteworthy, point made by some participants regarding the content of certain emails made reference to advice to the recipient about phishing emails and scams. Whilst this will be discussed in more detail as a cue for legitimate emails, it should be noted that some participants actually reported such warnings as a cue to phishing:

The fact that it says ‘be wary of scams. Do not send money via Western Union’ seems a bit strange to be at the bottom of the email, so it might be like a double bluff. (P312, Line 132)

Context of the email. This category describes cues in which participants consider the likelihood of a company sending out a specific email. For example, many participants recognised that “big companies wouldn’t ask you to do that” (P312, Line 67), when an email was asking for information, or requesting that the recipient follow a link. There were numerous references to contact from companies regarding issues such as fraudulent activity, particularly banks, suggesting that “they would phone you or send you a letter” (P351, Line 8) in such a situation.

Alternative scenarios also raised suspicion in terms of the likelihood of such an event occurring, for example noting that “they wouldn’t disable an account just for a security issue” (P306, Line 57), when the recipient was informed that their account would be deactivated if they did not respond. In addition, regardless of the specific scenario, some participants recognised the ‘too good to be true’ nature of

some emails, commenting that “nobody is ever giving anything to anyone, especially through the internet” (P308, Line 50).

Terminology. The way in which the recipient was addressed in the email was a commonly used cue to phishing, with many participants noting that emails not addressed to a specific person were “not personal” (P302, Line 151), “kind of informal” (P310, Line 62), and “a little bit too casual” (P307, Line 8). Further to this, the spelling and grammar within the main text of the email was also relied upon for some participant’s decisions, with “poor grammar” (P350, Line 178) and “all the typos make it just really obvious” (P351, Line 207).

Appearance of the email. Many emails were judged as phishing because there were “no logos or company styling” (P302, Line 73) or in cases where the “logo is not correct” (P303, Line 39). The overall appearance and format of an email was evidently important in the decision-making process, with participants comparing between the emails and picking out those that were “not formatted like the other emails” (P305, Line 49). Poor formatting, such as “the use of exclamation points in the subject makes it seem unprofessional, as does the grammatical error of capitalising all the words in the subject” (P324, Line 3). This type of cue also extends to the links embedded within an email, with one participant observing “that doesn’t look like a legitimate link because it’s got loads of random letters and numbers and words” (P320, Line 11).

Emails identified as legitimate.

Email address. Participants demonstrated confidence in the legitimacy of an email when it was received from “a fairly sensible sounding email address” (P302, Line 85). As seen in Figure 6.1, this was the second most common cue used to

correctly identify legitimate and incorrectly identify phishing emails. Further to the email address itself, participants also relied upon a comparison of the address and any URLs included in the email, with the conclusion that an email could be “legitimate, because the domain is consistent” (P317, Line 32).

Content of the email. The content of an email was the most commonly used cue when participants were judging an email as legitimate. The believability of what an email was asking of the recipient was a common response, with participants reassured by content that “doesn’t ask you to do anything” (P353, Line 3), for example “not asking for a PIN” (P302, Line 26) or “doesn’t require me to enter or give any personal ID, credit cards or something like that” (P326, Line 20). In addition to the information requested, participants were more trusting of an email with “no particular rush for you clicking on the links” (P302, Line 142). Unlike with phishing emails, participants felt that an email was more likely to be legitimate if there was no sense of urgency to the matter at hand.

Whilst some participants took warnings about phishing emails as a sign of a double bluff by the sender, the majority were reassured by these and reported them as an indicator of a legitimate email, “there’s a lot of stuff like ‘learn more to protect yourself from spoof emails’, a lot of links there which you would think are legitimate” (P310, Line 97).

Emails containing URLs were often assessed based on whether the “links in it look like they could be real links” (P305, Line 89). Links that were associated with the company the email purported to come from, and that were consistent with the domain name in the email address of the sender were considered to be “a secure link”

(P313, Line 107), thus allowing the participant to conclude that an email was legitimate.

Context of the email. Results demonstrate that a participant's prior knowledge about emails sent from the company sending an email provided an important cue to correctly judging its legitimacy, as they are able to recognise that "it's just conceivable that the bank will send an email like this" (P307, Line 18). On the other hand, some participants noted that they were unfamiliar with certain companies and their usual email communications, and so were making judgements despite the fact that they "don't" really know what it would look like" (P302, Line 117).

In addition to cues used to recognise legitimate emails though, this category also highlights a crucial issue with the email legitimacy task – that participants cannot make a fully informed decision about the emails without knowing the specific context. Across both the legitimate and phishing emails, although more apparent in the legitimate category, participants made comments that their decisions were based on an assumption that the email they were judging was relevant to the recipient, noting that "you would know whether or not you'd just ordered that specific thing" (P306, Line 99).

Terminology. As noted when discussing phishing emails, participants often relied upon the quality of spelling and grammar within emails when judging their legitimacy, with emails being viewed as more trustworthy when they "can't see any spelling mistakes" (P310, Line 69). The overall tone of the email also provided a cue to legitimacy, with one participant noting, "the wording on this email seems to be too professional for it to be a phishing scam" (P327, Line 10).

Further to this, a personalised approach including the name of the recipient was often noted as a reason for judging an email as legitimate, even in the case of some phishing emails, where “the fact that it gives the actual name of the person, that seems, that make me think it’s more legitimate” (P312, Line 25).

Appearance. Many participants commented on the presence of logos and recognisable branding as a cue to the legitimacy of an email, with some acknowledging “letters which contain logos make me confident that they are not phishing” (P346, Line 284). The format of the email was also discussed; with those emails judged as legitimate generally being perceived as having a more “professional” (P310, Line 152), “well laid out” (P315, Line 80) appearance, and “how you’d expect it to be” (P305, Line 19). Finally, the sophistication of emails was mentioned by some, with participants making legitimate judgments because they didn’t “think phishing emails look that technical, so legitimate” (P321, Line 271).

Phishing emails mistakenly judged as legitimate by the majority. As highlighted earlier, there were two phishing emails that were marked as legitimate by a majority of participants. For the Amazon email, participants who incorrectly judged this as legitimate were reassured by the “official and legitimate” (P346, Line 24) appearance of the email, as well as the fact that the email “doesn’t ask you to do anything” (P307, Line 42). Similarly, participants who judged the PayPal email as legitimate were reassured that the email “doesn’t ask me to click on any address and [is] just the information of the payment that I have made” (P303, Line 69).

Evidence of cognitive processes in email decision-making. In Study 3, three cognitive variables were found to be significant predictors of accuracy on the email legitimacy task – cognitive reflection, inhibition, and sensation seeking. From

the think-aloud data, it is clear that some participants rely more on impulsive decisions, rather than contemplating all aspects of an email. This was evidenced through participants either not mentioning any cues at all, “I’ll say this is sketchy and give it a three” (P337, Line 36), or relying on a single clue. This is possibly the first thing they noticed about an email and made a decision based on that, for example, “the address it’s sent from looks alright, so I’ll rate that one quite high” (P305, Line 36).

Although there are differences within participant responses to show that they think more rationally about some emails than others, this difference is clearer between participants. Some participants contemplated all aspects of the email stimuli, giving detailed accounts of their contemplative process and progression from an initial intuitive response, to contemplate all cues, such as:

Okay, the first email seems okay on first sight. It’s from Adobe; it’s got the Adobe signature. Well, you know, logo. Okay, it’s not asking from them to give a password; it says it’s reset. It’s an external link; the external link has Adobe in it. It’s to a specific user, rather than to several. It also gives a specific date, so you could presumably Google that to make sure that there was any evidence online. I’m going to say that’s very probably legitimate. (P306, Line 1)

Whilst other participants made immediate decisions with little time spent pondering the cues available to them to aid their decisions: “this looks [like] phishing” (P308, Line 112). This suggests, as demonstrated in Study 3, that there are individual

differences in ability to reflect upon a problem, and to suppress an impulsive response, in relation to the decision-making processes involved in email management. There is little evidence that participant decisions are influenced by their inclination towards sensation seeking behaviour though. Participants do not consciously report making decisions because the rewards outlined in the email stimuli entice them. However, some participants spent time contemplating the cues in an email that signify phishing, and then change their response at the last minute to legitimate, for example:

The information is quite weird. HR recruitment, but this is not an official account, by accounting or something like that, it is a private account. I don't think the HR department will actually give a private account. 'Attached the application form, along with directions'. You know what, I think this one is a little confusing. I will choose four. (P354, Line 90)

This may be a demonstration of an impulse to take the riskier option, despite the acknowledged warning signs that the email is not legitimate.

6.4 Discussion

The main aim of this study was to develop a more in-depth understanding of the cues and cognitive processes used by participants to aid decision-making whilst completing the email legitimacy task. Through the collection of qualitative, think-aloud data, a set of cue categories was compiled. In addition to this, the study also asked participants to complete the Cognitive Reflection Test (CRT) and the Moses

illusion. In Study 3, the CRT was a significant predictor of individual differences on the email task, although this was not replicated in Study 4. Study 4a showed no difference in scores on the Moses illusion based on email task performance. These tasks were included in this study to establish whether the previously reported findings are replicated with a different participant sample and an alternative variation of the email legitimacy task.

6.4.1 Email legitimacy task. Hypothesis 1 predicted that performance on the email task in this study would be better than in Study 3, as the think-aloud protocol would induce more rational decision-making in participants. This was supported, with higher performance in this experiment, based on both the number of emails correctly identified and the confidence demonstrated by participants in their decisions. It is likely that requiring participants to talk out loud whilst contemplating their judgments affected the decision-making process by bringing otherwise subconscious observations into consciousness (Smagorinsky, 1998). This may have resulted in participants being more engaged with the emails and confident in recognising cues that they would otherwise have missed. In relation to consideration in other chapters of the relevance of dual-system processing systems, it is possible that the think-aloud element in this task, which requires contemplation of reasoning, resulted in participants engaging in system 2 thinking styles. Participants in previous studies may have been making more rapid decisions without considering why they were doing so and what cues may aid in their decisions.

As well as participant gender, for which the null findings from previous chapters were replicated, this study also considered the influence of native language on email task performance. Results showed no evidence of a difference in

performance between native and non-native English speakers. This is a slightly surprising result, especially in an study where the participants were also asked to speak aloud for the duration in English, thus increasing the cognitive load required to complete the task. The finding provides a reassurance though that non-native speakers are not necessarily more susceptible, and that there is no evidence to suggest the search for predictors of susceptibility should be limited by language comprehension in non-native speakers.

6.4.2 Cognitive measures. In this experiment, participants were asked to complete the CRT and the Moses illusion task. In the Moses illusion task, participants are measured on their ability to spot discrepancies in the questions asked which made them distorted – for example, asking how many of each animal Moses took on the ark with him, when in fact it was Noah who took animals on the ark. Based on initial findings from a student project (Harrison, 2015), Hypothesis 3 proposed that higher susceptibility to the Moses illusion would predict lower accuracy in the email task. In contrast to findings from study 4, the Moses illusion task was found to be a significant predictor of confidence score, and was close to significance in predicting the number of emails participants were able to correctly identify. However, our results demonstrate the opposite, suggesting that score on the Moses illusion (based on the number of times the illusion occurred) increased with accuracy and confidence on the email task.

Based on findings from Study 3, which demonstrated that the CRT was a significant predictor of performance on the email task, Hypothesis 2 stated that higher levels of cognitive reflection would predict higher accuracy in the email task.

However, previous findings were not replicated in this experiment, despite there being little difference in mean scores on the CRT for each experiment. For both of the cognitive tasks included in this experiment, the higher mean accuracy on the email task compared to previous studies may mean that the predictive capacity of these variables cannot be assessed in the same way. The think-aloud protocol encouraged participants to verbalise their decision-making, which in turn meant increased contemplation of the stimuli. Therefore, participants who would otherwise have relied on inaccurate, impulsive responses may have been discouraged from doing this given the nature of the task in this experiment.

6.4.3 Think-aloud responses. Five major cue types were extracted from the think-aloud response data, which allowed for an understanding of the types of rationale participants generate when judging the legitimacy of emails. However, the most notable finding from this study is the overreliance participants had on cue heuristics that are ‘stereotypical’ and traditional to educational information about phishing emails, such as the email address of the sender in relation to embedded URLs, and grammatical errors. These findings seem to partially replicate those from Jakobsson et al. (2007), despite substantial technological advances in the generation of emails since 2007 that mean such heuristics cannot necessarily be relied upon any more. For example, it is easy for a fraudster to forge the email address of the sender so it appears to come from a legitimate address associated with the purported sender. This data emphasises that many users still rely on these slightly out-dated heuristics about what a phishing email looks like and how it can be easily detected, especially when they are unfamiliar with the sender. The poor reliability of such cues

is evidenced by the inaccuracy of some judgments where participants relied upon these – such as using a trustworthy looking email address as a cue to legitimacy in what is actually a convincing phishing email.

A further common response from participants was to judge emails based on their familiarity with what an email from a given company *should* look like, from personal experience. Participants seemed to dismiss some emails as phishing if they were unfamiliar with the company, which may be indicative of how they would react to an unsolicited email in their own inbox. This suggests that attempts to reduce susceptibility should possibly focus more on spear phishing emails that purport to come from companies known to the recipient, than generic phishing emails, which are more easily recognised. The relevance and familiarity of emails seems to play a major role in the decision-making process, and so phishing attempts which resemble familiar companies and seem relevant to transactions or behaviours that the user normally engages in are likely to be more successful. In fact, the overwhelming likeness of some phishing emails to what a genuine email would look like can sometimes lead to participants overlooking the above mentioned ‘typical’ cues to phishing. As demonstrated by the two phishing emails that were judged incorrectly by the majority, the incorrect sender email address was overlooked due to the accuracy in appearance of the rest of the email. This is a finding that is worth considering in future research and efforts to address susceptibility within users.

One interesting finding that came from the think-aloud data relates to the presence of security warnings within an email. Whilst a lot of users saw this as a cue to the legitimacy of an email, there were some participants who felt that the inclusion of information explaining how to recognise a phishing email may have been

a form of 'double-bluff', whereby the sender was trying to make the email more believable by including information about phishing scams. Although the majority of users took these warnings at face value, the fact that even a few participants did not raises scepticism over the effectiveness of such warnings.

In addition to the cues used to recognise phishing emails, this study also considered whether the qualitative data collected could be used to support previous findings from Study 3, which found cognitive reflection, inhibition, and sensation-seeking to be significant predictors of accuracy in the email task. This was demonstrated to some extent, with some participants seeming to make intuitive decisions with little consideration for cues present in the email, whilst others contemplated multiple cues before making a decision. A difference between participants in the aspects considered in decision-making supports prior findings that individual differences in cognitive reflection and inhibition may influence decision-making style. Without more substantial analysis though, it would not be possible to assume how this influences accuracy based on the qualitative data. There was also a difference within participants though, with decisions being made more impulsively for some emails than others. This is to be expected, given the range of obvious and less obvious email examples in the stimuli set, and does not take away from the evidence on a difference between participants. There is less evidence of the impact of sensation-seeking behaviour from the qualitative data, although some participants did contemplate the cues to indicate one type of email, and then changed their mind despite of these. This was not seen consistently for specific participants though, and seemed to be more email dependent.

Although there are some indicators of cognitive variables influencing the decision-making process, these cannot be used to definitively conclude that some participants consistently engage in intuitive decisions, and some demonstrate sensation-seeking traits. This is partly due to the subjective nature of qualitative data, but may also be due to the differences between participants in how well they engage with the think-aloud process. Whilst some participants comfortably verbalised every decision made, others were more hesitant in doing this, and were silent for periods of time whilst completing the email task. A familiarisation period was included, with participants asked to think aloud whilst completing the CRT in order to make them comfortable with vocalising their thoughts. However, it is possible that this was not enough for some participants, and so some were still reserved in their verbalisations, thus not providing a thorough account of their decision-making. There is little that could be changed in this methodology to improve this, as it is a case of some participants being more reserved than others, regardless of familiarity, and must be accepted as a limitation of the think-aloud protocol.

One further limitation, highlighted by many participants, points to a recurring issue with the methodology of using an email legitimacy task - the relevance of email stimuli to the recipient. At the beginning of the task, participants were instructed to assume that unless otherwise obvious they should assume that the email was relevant to the recipient (i.e. they did have an eBay account, and a bank account with NatWest etc.). It is evident though that there is some ambiguity in these instructions, and this assumption is not clear-cut for all of the emails. When receiving an email, a user will know whether or not that email is relevant to them, and the context surrounding the email. However, in this task participants have little context to base

their judgments on, thus making it difficult to come to an accurate conclusion in some cases. One solution to this is to include contextual information with each email. For example, with an Amazon order confirmation, there might be a short statement to inform the participant that ‘the recipient bought an item exactly like this 3 days ago from Amazon and then received this email’. Alternatively, studies in which participants are put in a situation where they receive emails related to a task they are completing may be a more reliable way of measuring susceptibility. This way, the participants are receiving the email themselves rather than looking at a screenshot, they are able to interact with the email inbox and make judgments about the relevance to the task, without any ambiguity about a fictional recipient.

6.4.4 Conclusion. The previous studies in this thesis have highlighted individual differences in accuracy on the email legitimacy task as a measure of susceptibility to email fraud. The data collected in this study contributes complimentary insight into the cues used by participants in making decisions about legitimacy. The security awareness information that is currently available (e.g. Citizens Advice Bureau, 2016; Action Fraud, 2016) seems to make assumptions about the cues that users are aware of, and those that are most beneficial to reducing susceptibility, with little empirical evidence. This study highlights the concerning reliance by our participants on out-dated cues to phishing which are easily masked in more sophisticated phishing attacks. In an attempt to address reliance on inaccurate cue heuristics, future efforts to educate users should focus on more sophisticated spear phishing attacks that are able to closely replicate the emails sent from genuine companies, and include personalised information. Given the ever changing nature of

technology and the techniques available to the fraudsters, it would be beneficial to educate users about how phishing attacks are conducted, how information is accessed, and how these are likely to advance in time, rather than focusing on a limited set of cues that will no longer be relevant in the near future.

This study has also highlighted a limitation in the use of an email legitimacy task as a measure of susceptibility, in that interpretation of the emails is subjective on a participant-by-participant basis. Without more detailed information regarding the context in which an email was received, there is a level of ambiguity about the believability of an email. This is an issue that needs to be considered in using this task as a method for measuring susceptibility. The task provides an ethically sound, controlled measure of email decision-making, but this is not necessarily representative of how participants would manage their own email inbox. Therefore, it might be useful to consider alternative approaches to assessing susceptibility that consider the context of emails being judged. It should also be noted that whilst efforts were made to include a range of email types of ranging subtlety in the email stimuli, these are limited to 36 emails. Therefore, the cues highlighted from this sample do not necessarily represent decision-making behaviour across all phishing emails, but instead provide an overview of the cues most commonly relied upon.

Chapter 7

The impact of cognitive load on email fraud susceptibility

Chapter summary

The study reported in this chapter assesses the impact of a secondary task on performance in the email legitimacy task. Participants were asked to complete one of four secondary tasks simultaneously to the email task. The secondary tasks included verbal (counting backwards out loud) and motor tasks (finger tapping), of differing complexity in order to assess the impact of varying levels of increased cognitive load. The cognitive reflection test was also administered here. Findings demonstrated that when compared to a control group, accuracy on the email task whilst completing a simultaneous verbal task was significantly lower. No effect was found for the motor tasks, regardless of complexity. The cognitive reflection test was found to be a significant predictor of accuracy on the email task here, replicating findings from Study 3.

7 Study six

7.1 Introduction

In the previous chapters, individual differences in performance on the email legitimacy task have been demonstrated through quantitative analyses, as well as qualitative analyses based on the various cues that participants report using to judge email legitimacy. A number of cognitive and contextual variables have been found to predict accuracy on the email task, which has been used as a lab-based measure of susceptibility to email fraud. The current study will focus on contextual influences on susceptibility, considered the impact of a secondary task on email judgment accuracy.

In an office environment, it is likely that internet users will be managing a number of tasks simultaneously, and will not be able to give undivided attention to one task. Whilst concentrating on other tasks, an email might come in telling the user that they need to respond immediately to change their password and avoid a suspension on their email account. Still concerned with the other tasks that are higher on their priority list, the user may be more likely to click on this email and enter their personal information in order to return to the more pressing task at hand.

Study 3 demonstrated how manipulating time pressure affected accuracy on the email task, with participants found to show higher susceptibility when under a time constraint to complete the task. One explanation for this considers dual-system decision-making processes (e.g. Stanovich, 1999; Kahneman, 2000; and Evans, 2003), which suggest that there are two behavioural responses in decision-making – system 1, which relies on more intuitive, rapid responses, and system 2, which involves more rational and analytic reasoning. The deployment of these systems is thought to

depend upon both situational factors, as well as individual differences between subjects that may influence the process. In the increased time pressure condition, email decisions were likely to be rushed, therefore inducing a reliance on system 1 type reasoning in decision-making. This is consistent with findings reported by Yan and Gozu (2012), which demonstrate that participants who were asked to make decisions as quickly as possible performed worse on a similar email task, compared to participants who were told to take their time with each decision.

Although Study 3 found no evidence to suggest that working memory span influenced decision-making in the email legitimacy task, this does not mean that working memory has no role in this process. The assessment of individual differences in working memory span differs from experimental manipulations of working memory capacity. Through the use of secondary cognitive tasks, this study aims to manipulate working memory capacity by increasing cognitive load, in order to emulate a scenario in which a user is completing multiple tasks at once. Past research has demonstrated a link between limitations on working memory capacity and an increased reliance on intuitive responses in decision-making (Kyllonen & Christal, 1990; Markovits, Doyon, & Simoneau, 2002). Therefore, in a situation where the user is required to divide attention whilst managing emails, they may be reliant upon intuitive responses that increase decision-making errors. This provides a convergent assessment to time pressure as a situational variable affecting accuracy in legitimacy judgments.

Whilst dual-system decision-making models provide one possible explanation for individual differences in email management, this is not the only explanation. Theories of divided attention suggest that situational factors can influence how well

a person is able to perform multiple tasks at once (Kane et al. 2001), due to the amount of attention they are able to allocate to each. Beede and Kass (2006) demonstrated the relevance of divided attention in a driving scenario, assessing driving performance whilst engaging in conversation on a hands free mobile device. Performance was negatively affected by simultaneous mobile phone use – suggesting that participants were not able to effectively divide attention between these two tasks. This scenario can also be applied to email management – in a situation where a user might be completing multiple tasks simultaneously.

In a more experimental setting, Jameson, Hinson, and Whitney (2004) demonstrated that participants were more likely to make impulsive decisions during a gambling task when they completed a digit recall task that interfered with working memory load. According to their argument, depriving the primary task of full attentional control meant that participants were less able to inhibit impulsive responses. This effect was not found with secondary tasks involving auditory suppression or a keypad task though, in which a number came on screen after each trial of the gambling task and the participant was required to press this on the keypad. Neither of these tasks have the same level of cognitive load as the digit recall task, implying that the nature of the secondary task is a key component in how it modulates processing and decision-making in the primary task.

This study considers two different secondary task types, one verbal and one motor, each with a simple and a complex condition, in order to compare performance on the email task, and thus establish the level at which cognitive load impacts decision-making ability. The verbal task involves counting backwards in ones (simple) or sevens (complex), a commonly used articulatory task. The motor task

involves finger tapping, with either a sequential (simple) or a non-sequential (complex) pattern to be constantly repeated on a second keyboard whilst completing the email task (outlined in Kane & Engle, 2000). Based on our previous finding that time pressure negatively impacted performance, the following hypothesis was generated regarding performance between the secondary task and control conditions:

H₁: Participants who complete a secondary task whilst simultaneously completing the email legitimacy task will perform with lower accuracy than a control group with no secondary task.

Within the secondary task types, it would be expected that the more complex version of each task – counting back in sevens and tapping a non-sequential pattern - would induce higher levels of cognitive load, thus allowing the following hypothesis to be generated:

H₂: Participants who complete the complex condition of each secondary task type will perform with lower accuracy on the email legitimacy task than those in the simple and control conditions.

In addition to the response accuracy outlined in previous studies using the email task, this study will examine participants' response times to each email stimulus, as well as the time spent selecting a response for each. These data will be used to evaluate whether time spent assessing an email and considering its legitimacy is important to the accuracy of a response. As previously demonstrated in Study 3, inducing a time pressure element on participants led to reduced accuracy. Therefore, the response time data will allow investigation of whether self-induced

time limitations on decision-making have the same effect on performance. Further to this, it will indicate whether increased cognitive load influences the decision-making process in terms of the amount of time participants spend contemplating stimuli before making a legitimacy assessment across secondary task conditions. Based on the previous finding from Study 3, of impaired accuracy when cognitive load was increased through induced time pressure, the following hypotheses were generated regarding the response time data:

H₃: Participants who spend longer looking at the email stimuli and contemplating scale responses will perform with higher accuracy on the email legitimacy task.

H₄: Participants in the control condition will spend longer looking at the email stimuli and contemplating scale responses than those simultaneously completing a secondary task.

Following up research from Study 5, which suggested that many participants confidently rely on cue heuristics (that are not always accurate) to detect phishing, this study investigated the relationship between self-reported confidence in knowledge of phishing and ability to detect phishing emails, and performance on the email task. This will demonstrate whether participants have confidence in cue heuristics that are sometimes inaccurate, thus biasing the accuracy of their responses, or if confidence is grounded in an improved ability to recognise phishing. Previous research has provided contradictory findings about the effect of knowledge and experience on susceptibility. Friedman et al. (2002) found that participants who were highly experienced with technology were no better at defining what is meant by a secure connection than other professional participants with no specific

technological experience. However, Downs et al. (2007) demonstrated that ability to correctly recognise fraudulent URLs and recognise security symbols was associated with lower susceptibility. In order to establish how knowledge and confidence affect performance on the email legitimacy task to be used in this study, the following hypotheses were generated:

H₅: Participants who report feeling more confident in understanding what 'phishing' is will perform with higher accuracy on the email legitimacy task.

H₆: Participants who report feeling more confident in their ability to differentiate between phishing and legitimate emails will perform with higher accuracy on the email legitimacy task.

Results from Studies 3, 4, and 5 have not provided entirely consistent findings with respect to the relevance of the Cognitive Reflection Test (Frederick, 2005) as a predictor of accuracy on the email legitimacy task. Study 3 found the CRT to be a significant predictor, but this finding was not replicated in Studies 4 and 5. It is possible that this was due to variations in methodology, or it may be that the finding from Study 3 was a false positive. In order to clarify this, the CRT was included in this study, with the following hypothesis to be tested:

H₇: Higher scores on the Cognitive Reflection Test will predict higher accuracy on the email legitimacy task.

7.2 Method

7.2.1 Participants. A sample of 102 participants took part in this study, with 76 females and 26 males. Participant ages ranged from 18 to 68 years, with a mean

age of 20.51 (SD = 6.14). When asked whether they were native English speakers, 77 participants reported that they were, whilst 25 reported they were not.

7.2.2 Design. This study incorporated a 2 (secondary task: verbal vs. motor) x 2 (task difficulty: simple vs. complex) between subjects design, with participants completing a secondary task simultaneously to the email legitimacy task. Outcomes measures from the email legitimacy task were considered as dependent variables, whilst secondary task condition and performance on the cognitive reflection test were analysed as independent variables.

7.2.3 Materials.

Email legitimacy task. Prior to the email task, this study asked participants to complete a few questions about their confidence in understanding and detecting phishing to establish how this relates to accuracy. Following this, participants completed the email task as outlined for Study 3. In addition, participants were asked to complete one of the secondary tasks listed below in Table 7.1, whilst working through the email stimuli. Instructions for the secondary tasks were given both verbally and on screen to ensure that participants understood and completed them simultaneously. For the verbal tasks, participants were instructed to start at 350 and count backwards either in ones (simple) or in sevens (complex). The instructions also outlined that if participants got confused, made a mistake, or reached zero, they should start counting again from 350. For the motor tasks, a second computer was placed on the desk, and participants were asked to tap a given sequence on the keyboard of this computer. For the simple task, the sequence was V B N M on a

Qwerty keyboard, so the sequence was consecutive along the keyboard. In the complex finger tapping condition, the sequence was M B N V (shown in Figure 7.1) – so the same letter keys were used, but these were in a different order to how they appear on the keyboard, making it more difficult to tap as a secondary activity (Kane & Engle, 2000).

Table 7.1 *Secondary tasks whilst completing the email task*

| Task | Description | N |
|----------------|-------------------------------------|----|
| Simple verbal | Count backwards in ones | 21 |
| Complex verbal | Count backwards in sevens | 21 |
| Simple motor | Consecutive finger-tapping sequence | 19 |
| Complex motor | Non-consecutive tapping sequence | 18 |
| Control | No secondary task | 23 |

Emails were displayed in a random order on screen, with participants being instructed to continue to the next page as soon as they had made a decision about each email. The response scale was then displayed on the next page, where participants inputted their chosen response. For every email, the time spent on the screen displaying the email image, and the time spent selecting a scale response was recorded in order to establish how long participants took to make decisions about legitimacy.

Cognitive reflection test. As described in previous chapters, the cognitive reflection test, a measure of ability to suppress an intuitive response in order to reach an accurate one, was administered in this experiment.

7.2.4 Procedure. This study was completed in a lab within the psychology department to allow the researcher to observe that participants continued to complete the secondary tasks simultaneously whilst judging the email images. As in




Figure 7.1 Finger tapping sequence for complex condition of secondary task

Studies 4, 4a, and 5, all materials for this study were hosted on Qualtrics (www.qualtrics.com), and these were completed on an Apple iMac.

Participants were first asked to provide some basic demographic information and answer questions about their confidence in understanding and detecting phishing emails. Following this, the email legitimacy task was administered, whilst the participant also completed a secondary task – either counting backwards, tapping a finger sequence, or no task in the control condition. Once they had completed the email legitimacy task, participants were reminded that they could stop the secondary task and then completed the cognitive reflection test.

Debrief information was shown on screen once participants had completed all of the tasks. Data recorded via Qualtrics was downloaded, and stored on an encrypted hard drive in preparation for analysis.

7.2.5 Data collation. As in the previous studies, binary responses from the email legitimacy task were taken to generate a number correct score, whilst scale responses were used to calculate a confidence score. Similarly, the Cognitive Reflection Test was scored in the same way as described in Study 3.

7.3 Results

7.3.1 Email legitimacy task performance. The mean number of emails correctly identified was 23.44 ($SD = 4.55$), and the mean confidence score was 28.75 ($SD = 22.11$). A one-way ANOVA demonstrated a significant difference in accuracy between studies 3, 5, and the current study, $F(2, 351) = 16.58, p < .001, \eta^2 = .09$. Tukey post-hoc analyses demonstrated that accuracy on the email task in this study ($M = 23.04, SD = 4.56$, with control condition excluded) was significantly lower than both study 3 ($M = 24.57, SD = 3.43, p < .01, d = 0.38$) and study 5 ($M = 26.84, SD = 3.18, p < .001, d = 0.97$).

An independent t-test demonstrated no significant difference in the number of emails correctly identified ($t(100) = 1.28, p = .20, d = 0.30$) or confidence score ($t(100) = 1.86, p = .07, d = 0.44$), dependent on participant gender. A significant difference was shown between native and non-native speakers though, with native speakers ($M = 24.01, SD = 4.23$) able to accurately judge more emails than non-native speakers ($M = 21.68, SD = 5.12$), $t(100) = 2.27, p < .05, d = 0.50$. Similarly, native

speakers ($M = 31.81$, $SD = 19.84$) were found to have higher confidence scores than non-native speakers ($M = 19.32$, $SD = 26.23$), $t(100) = 2.52$, $p < .05$, $d = 0.54$.

The mean values and standard deviations for performance in each of the secondary task conditions are shown in Table 7.2. For the number of emails correctly identified, a one-way ANOVA demonstrated that there was a significant difference in performance between the secondary task conditions, $F(4, 97) = 7.97$, $p < .001$, $\eta^2 = .25$. Based on Tukey post-hoc analyses, performance in the simple verbal condition was found to be significantly lower than the simple motor ($p < .01$, $d = 1.03$), complex motor ($p < .01$, $d = 1.26$), and control conditions ($p < .05$, $d = 0.87$). Similarly, performance in the complex verbal condition was significantly lower than the simple motor ($p < .01$, $d = 1.21$), complex motor ($p < .01$, $d = 1.48$), and control conditions ($p < .01$, $d = 1.02$). No significant difference in performance was found between the two verbal conditions, or between the two motor and control conditions, with all $ps > .05$.

When taking the confidence score as the measure of susceptibility, again there was a significant difference in performance across the conditions, $F(4, 97) = 6.91$, $p < .001$, $\eta^2 = .22$. Further Tukey post-hoc analyses demonstrated the same pattern, with confidence in the simple verbal condition significantly lower than the simple motor ($p < .05$, $d = 0.89$), complex motor ($p < .01$, $d = 1.23$), and control conditions ($p < .05$, $d = 0.89$). Performance was also lower in the complex verbal condition than the simple motor ($p < .05$, $d = 0.96$), complex motor ($p < .01$, $d = 1.36$), and control conditions ($p < .05$, $d = 0.96$). No significant differences were found between the two verbal conditions, or between the two motor and control conditions, with all $ps > .05$. For both measures of susceptibility, results demonstrate

that accuracy is lowest in the two verbal conditions, whilst performance in the two motor conditions does not differ from the control condition where participants had no secondary task.

Table 7.2 Means and standard deviations for performance on the email task

| Condition | Number correct | | Confidence score | |
|----------------|----------------|------|------------------|-------|
| | Mean | SD | Mean | SD |
| Control | 24.83 | 4.33 | 35.83 | 19.56 |
| Simple verbal | 20.95 | 4.60 | 16.67 | 23.55 |
| Complex verbal | 20.62 | 3.94 | 16.43 | 20.76 |
| Simple motor | 25.26 | 3.72 | 35.68 | 19.13 |
| Complex motor | 25.94 | 3.22 | 40.83 | 14.74 |

7.3.2 Response times on email task. Mean response times for looking at emails and selecting scale responses are indicated by secondary task condition in Figure 7.2. The mean time spent looking at emails was 19.89 seconds ($SD = 8.75$), whilst the mean time spent selecting scale responses was 3.77 seconds ($SD = 1.30$). Paired samples t-tests compared response times between phishing and legitimate emails. For the time spent looking at each email, there was a significant difference ($t(101) = 2.13, p < .05, d = .21$), with participants spending more time looking at legitimate ($M = 20.43, SD = 9.50$) than phishing ($M = 19.36, SD = 8.70$) emails. There was no significant difference for the time spent selecting scale responses though, $t(101) = -.76, p = .45, d = 0.08$.

A one-way ANOVA demonstrated that there was no significant difference in how long participants spent looking at the emails between the five conditions, $F(4, 97) = 0.54, p = .71, \eta^2 = .02$. There was a significant difference between the conditions on the time spent in selecting scale responses though, $F(4, 97) = 8.19, p < .001, \eta^2 = .25$. Tukey post-hoc analyses demonstrate that participants in the simple

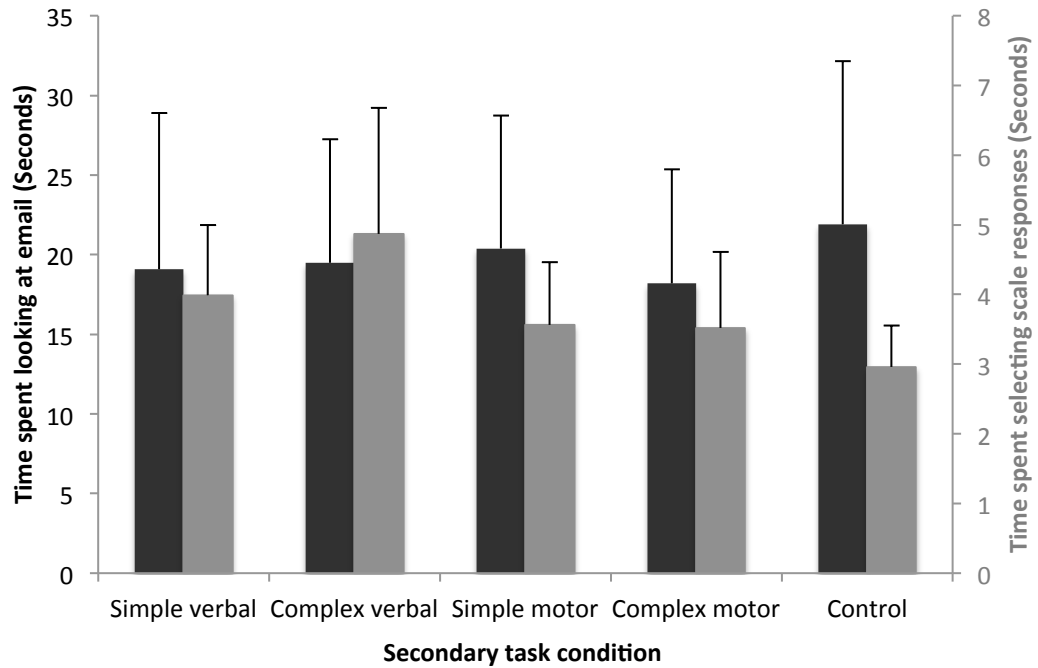


Figure 7.2 Mean response times on email legitimacy task by secondary task condition

verbal condition spent significantly more time selecting scale responses than those in the control condition ($p < .05$, $d = 1.25$), but not than any of the other conditions, $ps > .05$. Participants in the complex verbal condition spent longer selecting scale responses than those in the simple motor ($p < .01$, $d = 0.92$), complex motor ($p < .01$, $d = 0.90$), and control conditions ($p < .001$, $d = 1.41$). There was no significant difference between the two verbal conditions, or between the two motor and control conditions, with all $ps > .05$.

Correlational analysis demonstrated that there was a positive relationship between time spent looking at each email and the number of emails correctly identified ($r(101) = .36$, $p < .01$). No significant relationship was found between email correctly identified and time spent selecting scale responses though ($r(101) = -.12$, $p = .23$).

7.3.3 Confidence in recognising phishing emails. Following backwards multiple regression analysis, the best-fit model for the number of emails correctly identified included only the question about understanding of the term 'phishing' ($\beta = 1.70, p < .01$), producing an R^2 value of .09 (adjusted $R^2 = .08, F(1, 100) = 10.14, p < .01$). Similarly, when the confidence score is taken as the dependent variable, the best-fit model includes understanding of the term 'phishing' only ($\beta = 8.25, p < .01$), producing an R^2 value of .09 (adjusted $R^2 = .08, F(1, 100) = 10.19, p < .01$). Both of these findings demonstrate that increased confidence in understanding of the term 'phishing' predicted higher accuracy on the email legitimacy task.

7.3.4 Cognitive reflection test. Correlational analysis demonstrated a significant positive relationship between performance on the cognitive reflection test and the number of emails accurately identified ($r = .31, p < .01$), as well as confidence score ($r = .32, p < .01$). There was no significant relationship between performance on the cognitive reflection test and either the time spent looking at emails ($r = .19, p = .05$) or the time spent selecting scale responses ($r = .18, p = .07$).

Simple linear regression analyses demonstrated that the cognitive reflection test was a significant predictor of the number of emails correctly identified ($\beta = 1.30, p < .01$), producing a model with an R^2 value of .10 (adjusted $R^2 = .09, F(1, 100) = 10.68, p < .01$). For the confidence score, the cognitive reflection test was again found to be a significant predictor ($\beta = 6.51, p < .01$), producing an R^2 value of .10 (adjusted $R^2 = .09, F(1, 100) = 11.36, p < .01$). These findings indicate that a higher score on the CRT, indicative of ability to inhibit an intuitive response, was predictive of higher accuracy in differentiating between phishing and legitimate emails.

7.4 Discussion

The main aim of this study was to explore the impact of increased cognitive load on performance in the email legitimacy task. Based on previous findings from Yan and Gozu (2012), and Study 3 of this thesis, we expected to find that performance would be impaired when participants were under increased cognitive load, through inclusion of a secondary task. This study deploys different manipulations of cognitive load, and varied load complexity, thus developing upon previous findings. In addition, response time data was analysed for the first time, to establish whether the time spent looking at and contemplating the legitimacy of emails improved ability to recognise phishing. Finally, the cognitive reflection test was administered to participants in order to clarify contradictory findings in earlier studies in the thesis.

7.4.1 Email task performance and cognitive load. Initial analysis

between performance on this and the previous studies in the thesis that used the same email task demonstrated lower accuracy rates. This suggests that the inclusion of a secondary task whilst completing the email legitimacy task impairs performance. A difference was also highlighted between native and non-native speakers, suggesting that the inclusion of a secondary task unrelated to email management might make it more difficult for non-native speakers to comprehend and make accurate legitimacy judgments about the email task.

Hypothesis 1 predicted that participants in the control condition would perform with higher accuracy and confidence than those in the secondary task conditions. Results demonstrated that this was only true in relation to conditions

with a verbal task, with accuracy and confidence in these being lower than for the motor task and control conditions. Further to this, Hypothesis 2 stated that performance on the email task would be lower for participants in the complex secondary task conditions. There was no evidence to support this, with performance being lower in both of the verbal task conditions, regardless of complexity.

These findings suggest that the nature of the secondary task is relevant to its influence on susceptibility. Previous chapters in the thesis have highlighted dual-systems theories of reasoning as an explanation for findings that demonstrate lower inhibition and increased time pressure predict susceptibility. However, in this study a dual-system approach would suggest that the complexity of the secondary task should predict susceptibility, as this increased level of cognitive load would lead participants to become reliant on intuitive responses. As there was no evidence of task complexity affecting performance, this encourages consideration of alternative explanations.

Hiscock (1986) outlines how different types of secondary tasks can result in varying levels of disruption to the primary task based on the cerebral hemispheres that they are engaging. Past research has demonstrated that requiring participants to count aloud can interfere with reading comprehension more than alternative non-verbal tasks (Slowiaczek & Clifton Jr., 1980; Daneman & Newson, 1992). Therefore, the requirement to verbalise the backwards counting in this study may have interfered with participants' sub-vocalisation whilst reading and comprehending the email stimuli. If this were the case, then the verbal tasks, in both the simple and complex conditions, would have caused more interference and impaired ability to

make fully informed decisions about email legitimacy with only partially encoded information, compared to the motor tasks.

Alternatively, the vocalisation of the verbal tasks may have added experimenter bias. In the motor tasks, participants were completing the task on a second keyboard with no verbalisations, meaning that any errors participants made were not apparent to the researcher in real-time. This may have led participants to feel less pressure to perform well on the secondary task, meaning they were less concerned about it and it did not affect performance in the same way as the verbal tasks.

7.4.2 Response times and email task performance. Hypothesis 3 predicted that participants who spent longer looking at the email stimuli and contemplating their scale responses, would perform with higher accuracy on the email task. In addition, Hypothesis 4 predicted that in the control condition, participants would spend longer looking at and responding to the emails, as the lack of additional cognitive load would allow more consideration of the stimuli.

Analysis of the response time data in relation to accuracy demonstrated that participants who spent longer looking at the email stimuli did perform better on the task. In line with dual-system theories of reasoning discussed in previous chapters, this may be due to participants taking the time to engage in more rational decision-making, with all cues within the email taken into account. When the analysis took into account the secondary task conditions, participants in the verbal task conditions were found to spend longer selecting their scale responses than participants in the other conditions, contrary to our hypothesis. As no difference was found between

the conditions on time spent looking at the email stimuli, this would suggest that although participants were clicking through the email stimuli at the same rate, participants in the verbal condition needed longer to contemplate a scale response. This may be because they encoded less information whilst looking at the email itself, due to the increased cognitive load, and so were less prepared to make a decision when faced with the scale response screen. Given that accuracy was lower for the verbal task conditions, it seems that the extra time spent selecting scale responses is more likely due to attention deficits than to taking more time deliberating and reaching the correct response.

In addition to the hypotheses outlined, further analysis of response times demonstrated that participants spent more time looking at the legitimate, than the phishing, emails. Given the nature of this email task – with participants being asked to differentiate between phishing and legitimate emails – this might be explained by participants looking for cues to phishing in the legitimate emails, which were not present. In a task where participants were not knowingly searching for phishing emails, this finding may be less evident. It is unlikely that in real life email management, users spend longer contemplating legitimate emails, and so this finding is thought to reflect the nature of the methodology, rather than that of cognitive processing in relation to email management behaviour.

7.4.3 Understanding of the term ‘phishing’. Hypothesis 5 stated that participants who self-reported higher confidence in their understanding of the term ‘phishing’ would perform with higher accuracy on the email task, which was supported in the findings from this experiment. When considering participants’

confidence in their ability to recognise phishing emails, Hypothesis 6 predicted that, again, an increased level of confidence would predict higher accuracy on the email task. However, no differences in performance on the email task were shown. This suggests that knowledge, although only self-reported here, is more important to performance than confidence in ability. Previous research has demonstrated that users who have a better knowledge of the internet and cyber security are less susceptible to phishing attacks (Downs et al., 2007). Our finding supports this, suggesting that a good understanding of what phishing is may influence ability to detect phishing emails.

This finding may be an important consideration in future educational efforts to reduce susceptibility. Rather than a focus on the individual cues typically present, which can vary massively between emails, information about what phishing is, what information fraudsters are trying to access, and the tactics used to access this may be more beneficial. In addition, the lack of evidence that confidence in ability to detect phishing emails improves accuracy emphasises the issue that some users demonstrate a level of overconfidence, self-reporting higher confidence than their accuracy warranted.

7.4.4 Cognitive reflection and email task performance. Based on findings from Study 3, which found evidence of a relationship between performance on the CRT and on the email task, Hypothesis 7 proposed that this would be replicated in the current experiment. One interpretation of the previous finding was that this reflected individual differences in the deployment of system 1 / system 2 decision-making processes. This would suggest that participants who were better

able to suppress an intuitive response to the cognitive reflection questions were also better at rationally analysing email stimuli, engaging system 2 reasoning to make more accurate decisions. This finding was replicated in the current experiment, with participants who demonstrated higher cognitive reflection performing with higher accuracy on the email task.

However, there was no evidence of such a relationship in Studies 4 and 5. This may be attributed in part to the methodological differences between the experiments. In Study 4, the varied proportions of phishing and legitimate emails in the email task means that susceptibility was being measured slightly differently to those versions of the task that include equal amounts of each email type. Whilst in Study 5, the inclusion of the think aloud protocol was found to increase accuracy on the email task. The inclusion of a secondary task in the current study also influenced performance on the email task, yet the relationship with the CRT was still evident. In each case, the predictive power of the cognitive reflection test is relatively small, and so any noise from methodological or environmental differences may be responsible for the deviations from significance between experiments. Further examination of performance on the CRT and email susceptibility, as measured by alternatives to the forced-choice email legitimacy task, may provide a more robust outcome in terms of its ability to predict susceptibility.

7.4.5 Conclusions and future directions. From this experiment, it has been shown that increasing cognitive load through the use of a secondary task can impair performance on the email legitimacy task. However, this impairment is task dependent with performance only impaired for participants asked to complete a

verbal secondary task. Those participants who completed a simultaneous motor task showed no evidence of decreased accuracy in performance. In relation to real world email management behaviour, future research might focus on establishing how these different secondary tasks relate to day-to-day activities that users might be engaging in whilst also managing their email inbox. The articulatory nature of the verbal task may result in higher cognitive load and impaired sub-vocalisation, compared to the motor task, making it more difficult to focus attention on the email task. In an office setting, it might be the case that talking on the phone to a colleague induces cognitive load in the same way due to the articulation, and might temporarily increase susceptibility to any phishing emails coming in during that time, where a non-articulatory task such as filing paperwork might not. Further research of email behaviour in an actual office environment could investigate this further, highlighting the situations in which users are most at risk of victimisation.

When contemplating the improvement of educational efforts to reduce susceptibility to phishing emails, there are some important considerations from the findings in this experiment. The impact of a verbal secondary task on ability to divide attention, and the replication of findings that cognitive reflection predicts accuracy in email management, suggest the importance of rational, contemplative decision-making. This is supported further by findings regarding response times, with evidence that participant accuracy improved when they took more time to contemplate the email stimuli, before going on to make a decision about legitimacy. This need for contemplation and rational decision-making is something that could be encouraged through a training programme, by emphasising the cues that can be missed and consequences of these when users make more impulsive decisions.

Chapter 8

Cognitive and situational predictors of email fraud susceptibility in a simulated office environment

Chapter summary

The study reported in this chapter introduces an alternative methodology for assessing susceptibility to phishing emails. An office simulation task was developed, in which participants completed a number of day-to-day office tasks, including email management. This allowed for assessment of response behaviour, whilst participants remained naïve to the purpose of the experiment. The cognitive variables found to be predictive in Study 3 were included here to assess whether findings were replicated with a potentially more ecologically valid measure of susceptibility. Additional situational factors, including time pressure, email relevance, and priming about phishing emails, were also manipulated to assess how these influenced response behaviour. Findings somewhat replicated Study 3 with regard to the cognitive variables found to predict susceptibility, including cognitive reflection and sensation seeking. In addition, email relevance provided a further partial explanation of response behaviour.

8 Study seven

8.1 Introduction

In Studies 3-6 an email judgment task was employed as a measure of susceptibility, which explicitly asked participants to make legitimacy judgments about a set of email screenshots. Whilst this methodology allows for experimenter control over independent variables, and does not raise any particular ethical concerns, it is limited in the sense that participants are actively seeking to differentiate between phishing and legitimate emails in a forced-choice task, potentially reducing the ecological validity of the task. Therefore, this study introduces an alternative methodology, an office simulation, which allows participants to interact with an email inbox in an office environment, whilst remaining naïve to the true purpose of the research.

Previous studies in the thesis produced some mixed results regarding the relevance of specific cognitive variables in predicting individual differences in susceptibility to phishing emails. Whilst Study 3 identified cognitive reflection, inhibition, and sensation seeking as predictors of susceptibility, findings were not consistent in the later studies. Studies 4 and 5 found no significant predictors of susceptibility, but Study 6 did replicate findings of a significant relationship between susceptibility and the cognitive reflection test. Given that susceptibility was measured using a task that was inspired by previous work (e.g. Yan & Gozu, 2012) but adapted for the research in this thesis, and in addition used a novel set of stimuli, one possibility is that the lack of consistency in the findings from these studies is due to the nature of the task. There is no evidence as yet to show whether judgment tasks in a lab-based environment capture the exact behaviour that participants

would engage in with their own inbox. No relevant literature has been published to demonstrate this. However, reflection upon this matter suggests that the participant in a laboratory environment does not have anything to lose in this situation. Whilst participants were given a monetary incentive to perform well, this may not emulate the fear and urgency that is potentially experienced when you receive an email threatening to suspend access to your bank account, for example. Instead, participants are actively seeking out the phishing and legitimate emails in order to perform well and obtain the monetary reward. In addition, the email task described in previous studies in the thesis incorporated a 6-point response scale, meaning that participants were forced to make a choice between phishing and legitimate, where in real life the decision-making process would not have the same clear cut, binary outcomes. This response process allowed for a richer and more sensitive account of email judgment than a simple binary measure, whilst also allowing for the calculation of binary scores. However, it does not provide an insight into the range of possible reactions that users may have to receiving phishing emails – for example, one user may open and read an email, whilst another may also decide to click an embedded link, or delete it all together, with different responses demonstrating varying levels of susceptibility.

An alternative methodological approach to measuring susceptibility, whilst still retaining the experimental control of a lab-based setting, involves role-play scenarios in which the participant is asked to manage an email inbox under the premise that they are in the role of a fictional character. Previous examples of this type of methodology are outlined in more detail in Chapter 1 (e.g. Downs et al., 2007; Hong et al., 2013; Parsons et al., 2013), but all follow the same basic premise.

Participants are given a set of emails and asked how they would manage these. The examples in previous literature follow a similar structure to the email legitimacy tasks though, with forced-choice decisions, meaning that this methodology is subject to the same limitations in terms of how representative it is of real world email behaviour.

In this experiment, the concept of role-play scenarios is developed as a methodology for understanding more about participants' email management behaviour without making apparent the purpose of the research. An office simulation scenario was developed, in which participants take on the role of an employee and complete a number of office-based tasks. Simultaneously, they are tasked with managing a personal email inbox. Although the research examples above which used a role play task did not explicitly tell participants about the nature of the study, the sole focus on email management in the task and response options provided may have given a strong clue, or at least focused the participants' attention more to studying the emails than they would be likely to do in a day-to-day work environment. Therefore, by focusing the study on the completion of the office tasks, this study aims to ensure that participants remain naïve to the nature of the study.

Study 3 revealed that time pressure led to lower accuracy on the email judgment task (replicating Yan & Gozu, 2012). In addition, convergent evidence from Study 6 demonstrated that a secondary verbal task impaired accuracy on the email task. From this, one possible explanation may be the influence of dual system theories of reasoning and decision-making (Stanovich, 1998; Kahneman, 2000; and Evans, 2003). Increased cognitive load from an induced time pressure or secondary task might lead to participants relying on system 1, intuitive decision-making

processes more heavily. To offer convergent evidence on these findings, the current study includes a time pressure manipulation on the office tasks being completed, generating the following hypothesis:

H₁: When time constraints are in place on the office tasks, participants will demonstrate higher susceptibility to phishing emails received.

Parsons et al. (2013) demonstrated that knowledge of the purpose of a study on response to phishing emails, in which participants were asked how they would respond to a set of email stimuli, led to better performance. This is a major limitation of the email legitimacy task, as participants are asked to make explicit judgments between phishing and legitimate email stimuli. Whilst the office simulation methodology aims to provide a measure of susceptibility where participants remain naïve to the nature of the study, an additional manipulation involved priming participants about the presence of phishing emails in their inbox. That is, half of the participants were given written and verbal information to suggest that some of the emails they received may not be legitimate. Whilst participants are still not being explicitly told that the study is about email behaviour, as those from the previous chapters have done, and as Parsons et al. did, it would still be expected that when participants are primed about an influx of phishing emails they would act more cautiously when managing their email inbox. In a real world organisation, warnings are often circulated to alert employees to recent phishing emails or spates of attacks. In including the priming condition, the aim is to emulate such a scenario, highlighting that phishing emails may be present in their inbox during the study, and that they should be wary of these. This will allow assessment of the effectiveness of

such warnings when participants are also engaged with other tasks, which could potentially make them forget about these. Previous research has suggested that prior knowledge about security cues used to recognise phishing (Downs et al., 2006) and domain-specific experience (Vishwanath et al., 2011) does not reduce response likelihood to phishing emails. However, in a study that implemented an intervention to raise awareness of telephone-based social engineering attacks, an immediate improvement one week following the intervention was reported. This effect was lost two weeks after the intervention though, and response rates to a simulated attack actually demonstrated increased susceptibility in comparison to a control group who received no intervention (Bullée, Montoya Morales, Junger, & Hartel, 2016). It is therefore of interest in this study to assess whether a warning about phishing emails immediately before the study has the same success as the short term intervention on telephone based attacks, with the following hypothesis being tested:

H₂: Participants who are primed about an influx of phishing emails will demonstrate lower susceptibility to phishing emails received.

One limitation highlighted in the previous studies in the thesis, in particular through the qualitative data collected for Study 5, is that participants struggle to judge the email stimuli without information about the relevance of each email to the recipient. Participants were told in this task to assume that the recipient held accounts at the companies the emails purported to come from, and to assume that the email was relevant unless it was clear otherwise. However, there were a number of discrepancies between participants as to whether each email might be relevant, which affected judgment of these. This emphasises a point though that is present in

day-to-day email management - the relevance of an email to the recipient is potentially a very influential factor on their likelihood to respond to it. If a user banks with NatWest and receives an email from them warning that their account will be suspended, they may be more likely to reply than if they received the same email from Barclays, with whom they do not hold an account. Hadlington (2015) outlines areas of key research investigating how the relevance of technological interruptions, such as incoming emails, to a primary task being completed can affect the user. Interruptions that are irrelevant to the task at hand are shown to disrupt concentration and increase stress (Li, Edwards, & Lee, 2002), whereas interruptions that contain information related to the primary task can benefit the user and encourage completion of the primary task. However, such research has not considered whether the effect of relevance is a symmetric or an asymmetric tool – relevant information may inform the primary task, but can the relevance of the primary task also help to inform decisions relating to the interruption itself. That is, when an incoming email interrupts a primary task, is a user more likely to pay attention to, or demonstrate susceptibility to, an email that is related to a primary task they are working on? In order to test this, some of the email stimuli in this study were designed to relate to the office tasks being completed by participants, whilst others were related to the job role and office in general, or were completely irrelevant. The following hypothesis was generated based on previous research focused on primary task interruptions:

H₃: Participants will demonstrate higher susceptibility to the phishing emails that are relevant to the office tasks they are completing.

This study aims to establish an experimental scenario in which assessment of email management behaviour is representative of real world behaviour, based on average response rates. As in previous studies in the thesis, cognitive factors are considered in their ability to predict response behaviour. This provides a valuable comparison to assess whether findings where an explicit judgment task was used as a measure of susceptibility are replicated with an alternative, potentially more ecologically valid measure. It is hoped that the experimental masking of the subject expectancy effect, and the implementation of a substantial role-play scenario to encourage participants to get into the character of the role appointed to them, will provide an alternative and complementary measure of susceptibility to phishing emails. In Study 3, cognitive reflection, inhibition, and sensation seeking were found to be significant predictors of performance on the email judgment task. Study 6 provided further evidence of cognitive reflection as a predictor of susceptibility. In addition to these predictors, a measure of self-control is also included in this experiment. Based on previous literature from Holtfreter et al. (2010), which demonstrated a link between susceptibility to telemarketing fraud and self-control, this is included to establish whether the alternative methodological approach in this study will lead to a replication of this finding. Therefore, assessment of the findings from previous studies in this thesis, as well as the wider literature, will demonstrate how a more realistic email management environment affects predictive power, with the following hypothesis:

H₄: Lower levels of cognitive reflection, inhibition, and self-control, and higher levels of sensation seeking behaviour will significantly predict increased susceptibility to phishing emails.

Finally, Studies 5 and 6 provide contradictory evidence regarding the relationship between native and non-native participants and susceptibility. In these studies the participant was focused solely on the completion of the email judgment task. In Study 5, no difference in accuracy on the email task was found, despite participants being required to vocalise their decision-making. However, in Study 6, where the participant completed a secondary task at the same time as the email task, accuracy was found to differ between native and non-native English speakers, with native speakers demonstrating higher accuracy. Therefore, it is of interest in this study to establish whether native language affects performance when participants are completing multiple tasks involving reading and writing in a second language simultaneously. As the email management is a secondary task in this scenario, with the main focus being on the completion of the office tasks, it is possible that participants' will pay less attention to the language used in the emails and thus miss otherwise useful cues, which may have an amplified negative effect in non-native speakers. Based on this concept, and on findings from Study 6, the following hypothesis was generated:

H₅: Non-native speakers will demonstrate higher susceptibility to the simulated phishing emails included in the office scenario.

8.2 Method

8.2.1 Participants. A sample of 51 participants, consisting of Lancaster University students, comprised 29 females, and 22 males. Participants ranged in age from 18 to 40 years, with a mean age of 22.08 (SD = 4.32). From this sample, 25

participants reported that they were native English speakers, whilst 26 reported they were non-native.

8.2.2 Design. This study involves a 2 (time pressure: time pressure on office tasks vs. no time pressure) x 2 (priming: primed about phishing emails vs. not primed) x 2 (relevance: phishing email relevant to office task vs. not relevant) mixed design, with the last factor being within participants. Participants were systematically allocated to the between factors conditions (time pressure and priming) depending on the office simulation session they attended. Response behaviours to the simulated phishing emails are taken as dependent variables, whilst time pressure, priming, relevance, cognitive reflection, inhibition, sensation seeking, and self-control are considered as dependent variables.

8.2.3 Materials.

Cognitive tasks. Based on findings from Study 3, the Flanker task, brief sensation-seeking scale, and cognitive reflection test were administered to offer a conceptual replication, with a more realistic measure of susceptibility. In addition, the brief self-control scale was included, based on prior research from Holtfreter et al. (2010), to establish whether this is replicated with a novel email management task. These tasks are all outlined in detail in Chapter 3.

The Flanker task (Eriksen & Eriksen, 1974) provides a behaviour measure of inhibition, asking participants to suppress an initial, intuitive response to the stimuli presented in order to reach the correct response. The brief sensation-seeking scale (Hoyle et al., 2002) takes a mean score of an 8-item scale to provide a measure of

risk-taking behaviour. The CRT (Frederick, 2005) provides a 3-item measure of ability to reflect upon a problem, suppress an initial response in order to re-think the question and reach the correct answer. Based on the number of correct responses, participants receive a score between 0 and 3 on this task. Finally, the brief self-control scale (Tangney et al., 2004) provides a 13-item measure of participants' self-control levels. Once specified items are reversed, a total sum of responses is calculated for each participant to provide their self-control score.

Office tasks. During the office simulation section of the study, participants were given a set of tasks to complete, during which time they would also receive a number of emails. The office tasks were developed based on the usual day-to-day tasks of someone in an administrative role within an office environment. A total of 12 tasks were developed, with each participant completing six of these. Examples of the tasks included: data entry, letter writing, transcription, and office supply pricing. Step by step instructions for the tasks was provided in an information pack, and any additional materials required were provided on the computer desktop.

Phishing and legitimate emails. During the office simulation part of the study, participants were logged into an email account on a Lancaster based email server, developed specifically for research studies such as this one. An individual email address was created for each of the four character roles that participants could be assigned to. As part of the information pack, and verbal instructions from the researcher, participants were told that they were required to manage the email inbox as part of their role, and that they should treat this as if it were their own inbox. Over the duration of the office simulation, participants received 15 emails, five of which were designed to emulate phishing emails. These emails were sent from

@ex2010.lancs.ac.uk (the Lancaster based email server, also used for the participant inboxes), @gmail.com, or @yahoo.com addresses, with a number of different accounts having been set up for the purpose of the study. Emails were designed based on genuine phishing and legitimate email examples, and were varied in terms of how subtle they were as examples of phishing or legitimate emails. In addition, the relevance of the emails to the tasks that the participants were completing was varied. Table 8.1 outlines the make-up of the phishing emails received during the study, including information about relevance, as these are the emails of interest in the analysis that follows (see Appendix O for screenshots of these phishing emails).

Table 8.1 *Summary of phishing email stimuli*

| Email | Email address | Subject line | Relevance to tasks |
|-------|-----------------------------------|----------------------------------|---------------------------|
| 1 | acountmanagar@gmail.com | Account update – action required | None |
| 2 | h.smithson@ex2010test.lancs.ac.uk | Check payroll details | None |
| 3 | travel@ex2010test.lancs.ac.uk | Update to travel system | For character roles 2 & 4 |
| 4 | cheapofficesuplies@gmail.com | Cheap office supplies website! | For character roles 1 & 3 |
| 5 | loteryexpert@gmail.com | Guarenteed lottery WIN!! | None |

Post-study questionnaires. Following the office simulation exercise, each participant completed a questionnaire about the difficulty of each of their office tasks, including the email management, and how motivated they were to complete these. In addition, they were asked an open-ended question about what they felt the study purpose was. This was included to assess whether any participants had realised that the study was focusing on email management and phishing emails

(questionnaire can be seen in Appendix P). Following a verbal debrief from the experimenter that identified the true purpose of the study, participants were then administered a survey containing a screenshot of each phishing email sent during the study and asked how they responded to this (Did not open it; Deleted it; Opened it; Downloaded attachment/Clicked link; Responded; I do not remember receiving this email) and why. As well as providing statistical information on response behaviour, open-ended questions provided some qualitative information about why participants chose to respond or delete each email. This self-report data was used, in combination with information from each inbox and browser history gathered by the researcher at the end of each session, to assess decision-making.

8.2.4 Procedure. Participants signed up to a study advertised as being about task management and office behaviour, in order to ensure they remained naïve about the true purpose. After signing up, participants were asked to complete the cognitive measures assessment online at home, prior to the office simulation section of the experiment. This was administered using Qualtrics software (except for the Flanker task, which was programmed using a web-based version of PsyScript; Slavin, 2015). The online materials also solicited basic demographic information.

The main phase of the study comprised the lab-based office simulation. This lasted for approximately three hours. Upon arrival at the lab, participants sat at a desk with a laptop and an information pack explaining the tasks that they were required to complete during the study. Each session contained up to four participants (some were fewer due to non-attendance), and so there were four individual character roles, email accounts, and sets of tasks for the participants.

Participants were told that they would take on the role of an administrative assistant for a printing company called 'Lancaster Printhouse'. Verbal instructions were given by the researcher outlining that participants were required to complete six tasks over the course of the three hours, that in addition to this they were required to manage the email inbox that they were logged into on their laptop, and that any problems should be directed to the email address provided for the 'office manager' (an account that was monitored by the researcher during the office sessions).

Participants were also asked at this point to forward completed task documents to the 'office manager', allowing the researcher to ensure that they were actually making an effort to complete these. For those participants in the primed condition, verbal and written instructions were also given to inform participants that the company had recently seen an influx of phishing emails and so they should be wary of the legitimacy of emails received. In the time pressure condition, additional instructions emphasised that participants had only 20 minutes to complete each of the tasks. These additional instructions were simply omitted for those conditions without priming or time pressure.

Tea and coffee making facilities were available in the room, and there was a video camera recording the session, all of which participants were informed of before the commencement of the study. The video camera was set up to monitor whether there was any interaction between the participants, in particular whether they became aware of the purpose of the study through discussion of the emails.

Participants were then left to commence the tasks, with the researcher in the next room. The researcher then sent emails to each of the participant's email accounts approximately every 10 minutes for the duration of the experiment. A

schedule for the delivery of the emails was set out before the commencement of the study to ensure that the relevance of the emails to the task that the participant should have been completing at that time was accounted for. In the time pressure condition, additional emails were sent to participants from the 'office manager' email account 20 minutes after they had begun each task asking for the associated documents, in order to encourage them to complete this as soon as possible and move onto the next as a way of inducing time pressure. Some participants worked through the tasks quicker than others across the two time pressure conditions, and in the cases where a participant completed all of the tasks allocated to them, an additional task not included in their original task set was sent to them.

After all 15 simulated emails had been sent to participants (normally after around two hours and 40 minutes), the researcher re-entered the lab and told participants that they could stop working on the office tasks. The post-study questionnaire was administered, after which the researcher debriefed participants - explaining that the study was focused on email management and that some of the emails that they had received were designed as phishing emails. Finally, participants were asked to complete the second post-study questionnaire, about their responses to the phishing emails. Participants were then given a written debrief sheet and thanked for their participation. In addition to the responses given in the post-study questionnaire, after participants had left, the researcher looked at each of the testing laptops and recorded from the email account, browser history, and download history, how the participant had managed each of the emails. This accounted for any mis-reporting by the participants on how they responded to the emails.

All data collected both online and on the paper questionnaire about task motivation was recorded and transferred to an encrypted hard drive in preparation for analysis.

8.2.5 Data collation. Participants were asked what they thought the true purpose of the study was, having completed the office simulation phase. In addition, video recordings of the office simulation sessions, and verbal discussion with participants following the study were used to inform judgements of awareness. Only two participants indicated that they were aware of the focus of the study being phishing emails and response likelihood. It is possible that they had heard this from friends who had already completed the study (despite instructions to participants not to tell friends who may take part). These two participants were excluded from the analysis as their understanding of the purpose of the study may have affected response likelihood, and would not provide the naturalistic assessment intended.

For each of the phishing emails, there was a record of whether each participant opened, clicked the link or downloaded the attachment, replied to, and deleted each phishing email. Participants' actions in relation to each of the phishing emails were collated to provide a score for the number of emails opened, clicked, replied to, and deleted. In addition, based on the post-simulation comments participants gave when asked why they deleted or did not open an email, a score was generated for the number of emails recognised as phishing.

As previously outlined in Study 3, performance on the Flanker task was assessed using a mean difference score taken between the congruent and incongruent trials, where the target item was correctly identified. For all participants,

z scores were generated in order to identify any outliers in the response time data. Any scores below -2.56 or above 2.56 were removed before mean values were calculated. Data failure and partial completion of the task by some participants meant that there were 11 cases missing for the Flanker task. One further case was removed, as this was an outlier ($z = -3.59$) when compared to mean response times from other participants. Scores for all other cognitive measures were calculated as outlined in Study 3.

8.3 Results

8.3.1 Response likelihood to phishing emails. Self-report data asking participants about why they deleted or did not open an email emphasised that the action of deleting or ignoring an email did not necessarily mean that the participant was aware of its suspicious nature. Many participants reported their reason for not opening or deleting an email being its irrelevance to the study or the task they were completing at the time. Table 8.2 shows the number of participants who either did not open or deleted each email, and the proportion who did so because they recognised the suspicious nature of the email. Given that we are interested in ability to detect phishing emails, further reference to the number of email 'recognised' will only include those that were identified as phishing emails.

Table 8.2 *Frequency for reasons given by participants who deleted or did not open each email*

| | N | Recognised as phishing | Irrelevant to study/ tasks |
|---------|----|------------------------|----------------------------|
| Email 1 | 6 | 83.3 | 16.7 |
| Email 2 | 3 | 33.3 | 66.7 |
| Email 3 | 4 | 50.0 | 50.0 |
| Email 4 | 3 | 33.3 | 66.7 |
| Email 5 | 13 | 61.5 | 38.5 |

Frequencies of response types for each email are shown below in Table 8.3.

From this, it can be seen that the rates at which participants replied to the emails are not dissimilar to the five per cent average response rate outlined by the a recent cybercrime report (Norton, 2014). However, victimisation can occur from simply downloading a file or clicking a URL, the rates for which are shown to be much higher. Correlations between the different response types recorded are shown in Table 8.4. These demonstrate significant relationships between the number of emails opened and the number in which the URL was clicked or file downloaded, as well as between the number deleted and the number recognised as phishing. In addition, a negative relationship can be seen between the number of emails opened and the number deleted, as well as the number recognised as phishing, as might be expected, as some participants were more wary and deleted emails on first suspicion of their illegitimacy, without even opening them.

Table 8.3 *Frequency for each action on each phishing email*

| | Opened | Clicked/ downloaded | Replied | Deleted | Recognised |
|---------|--------|------------------------|---------|---------|------------|
| Email 1 | 92.2 | 64.7 | 5.9 | 9.8 | 10.2 |
| Email 2 | 92.2 | 45.1 | 9.8 | 5.9 | 2.0 |
| Email 3 | 92.2 | 27.5 | 0.0 | 11.8 | 4.1 |
| Email 4 | 94.1 | 47.1 | 3.9 | 7.8 | 2.0 |
| Email 5 | 68.6 | 5.9 | 0.0 | 19.6 | 16.3 |

Table 8.4 *Correlations between response types to phishing emails*

| | 1 | 2 | 3 | 4 | 5 |
|---------------|--------|-------|------|-------|------|
| 1. Opened | 1.00 | | | | |
| 2. Clicked | .45** | 1.00 | | | |
| 3. Replied | .07 | .18 | 1.00 | | |
| 4. Deleted | -.39** | .00 | -.12 | 1.00 | |
| 5. Recognised | -.71** | -.35* | -.18 | .57** | 1.00 |

Note. ** $p < .01$ * $p < .05$ (two-tailed)

Based on self-report responses from the post-study questionnaire, the mean score for motivation to manage emails was 3.96 ($SD = 1.43$; rated on a 6-point scale), whilst the mean difficulty rating was 1.88 ($SD = 0.97$; rated on a 6-point scale). A series of simple linear regressions demonstrated that motivation to perform well in the email management task did not predict response behaviour based on any of the measures above, $R^2s < .05$, $Fs < 2.45$, $ps > .12$. Further simple linear regression analysis demonstrated that perceived difficulty of the email management task was also not predictive of response behaviour, $R^2s < .02$, $Fs < 0.28$, $ps > .60$.

Analyses of Studies 5 and 6 showed contradictory results regarding performance on the email judgment task between native and non-native English speakers. However, an independent samples t-test demonstrated that in this study there was a significant difference in the number of emails that native and non-native participants clicked the link or downloaded the attachment for. Native speakers ($M = 1.54$, $SD = 1.32$, $N = 24$) were less likely to react in this way to the emails than non-native speakers ($M = 2.28$, $SD = 1.14$, $N = 25$), $t(47) = 2.10$, $p < .05$, $d = 0.60$. In addition, a significant difference was found with the number of emails recognised as phishing, with native speakers recognised more emails ($M = .63$, $SD = 1.10$) than non-native speakers ($M = .08$, $SD = .28$), $t(47) = 2.41$, $p < .05$, $d = 0.63$. Consistent with previous studies though, there was no evidence of a difference in response likelihood on any of the measures based on gender, $ts < 1.69$, $ps > .10$, $ds < 0.49$.

Time pressure. A manipulation of the office task meant that some participants were under increased time pressure whilst completing these tasks. Response reactions were compared between those in the time pressure condition and those in the no time pressure condition. Figure 8.1 demonstrates the mean number of phishing

emails participants reacted in this way to for each of the time pressure conditions, which suggests an inverse relationship than that seen in Study 3. Participants in the time pressure condition demonstrating slightly *lower* susceptibility. A series of independent t-tests found no significant differences between the two time pressure conditions on any of the measures of response behaviour though, $t_s < 1.79$, $p_s > .08$, $d_s < 0.52$.

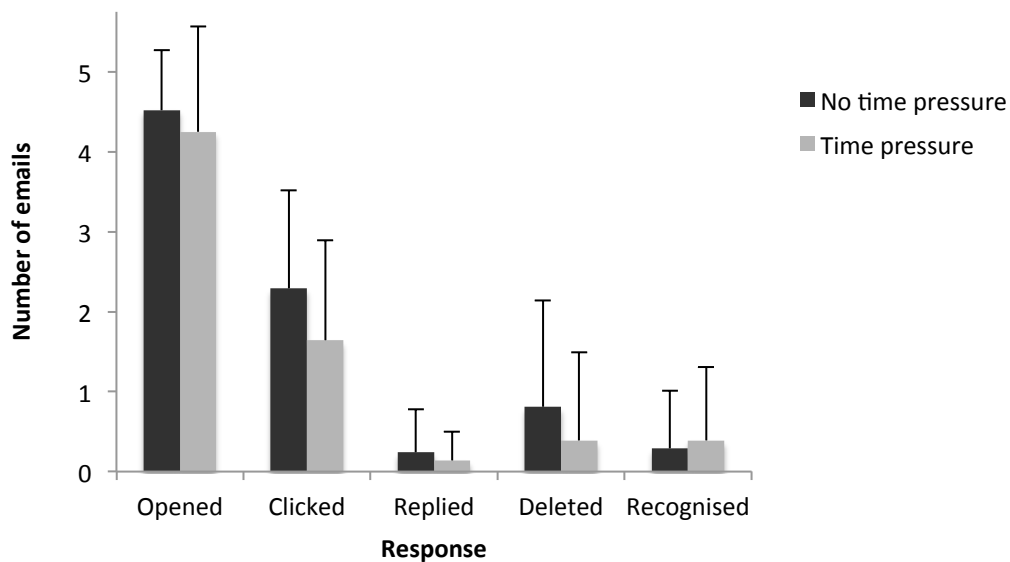


Figure 8.1 Graph to show mean number of emails and standard deviation for each response type by time pressure condition

Priming. An additional manipulation was priming of participants to an influx of phishing emails into the company, and a warning to be aware of the legitimacy of the emails coming in. Again, participants' reactions were compared between the primed and not primed conditions, with results shown in Figure 8.2. A significant difference was found in the likelihood to delete the emails, with those in the primed condition ($M = 1.04$, $SD = 1.54$) more likely to delete the phishing emails than those not primed ($M = 0.08$, $SD = 0.28$), $t(47) = 2.99$, $p < .01$, $d = 0.87$. No significant

differences between the two priming conditions were found in the other response actions, $ts < 1.03$, $ps > .31$, $ds < 0.30$.




Figure 8.2 Graph to show mean number of emails and standard deviation for each response time by priming condition

Email relevance. Two of the phishing emails (Email 3 and email 4) were designed to be relevant to the office tasks which some of the participants were completing. Chi-square analyses were conducted to establish whether the email being relevant to the participant's tasks affected their response type. The relevance of email 3 was found to affect whether the email was opened, $\chi^2 (1, N = 49) = 4.18$, $p < .05$, $V = .29$, with more participants in the relevant condition ($N = 24$) than the not relevant condition ($N = 21$) opening the email. There was also a difference in whether the URL was clicked or attachment downloaded for this email, $\chi^2 (1, N = 49) = 3.95$, $p < .05$, $V = .28$, with those in the relevant condition more likely to respond in this way ($N = 10$) than those in the not relevant condition ($N = 4$). Likelihood to reply, delete

the email, or recognise it as phishing was not affected by relevance though, χ^2 s < 2.00, $ps > .16$, $Vs < .20$. For email 4, there was no effect of the relevance of the email to the participant's tasks on their response type, χ^2 s < 2.17, $ps > .14$, $Vs < .21$.

8.3.2 Cognitive measures. Descriptive statistics for each of the cognitive tasks are shown in Table 8.5. Correlations for these measures and the measures of response behaviour for the phishing emails are shown in Table 8.6. These demonstrate a significant relationship between the number of phishing emails opened and performance on the cognitive reflection test, with more emails opened indicated lower levels of cognitive reflection. In addition, a significant relationship between the cognitive reflection test and number of emails deleted is also demonstrated, this time with more emails deleted indicating higher levels of cognitive reflection.

Table 8.5 *Descriptive statistics for cognitive measures*

| Cognitive measure | N | Mean | SD |
|-------------------------------|----|-------|-------|
| Cognitive reflection test | 49 | 1.20 | 1.17 |
| Flanker task | 37 | 11.35 | 38.08 |
| Brief self-control scale | 49 | 37.45 | 6.57 |
| Brief sensation-seeking scale | 49 | 3.06 | 0.84 |

Backward multiple regression analysis was conducted to establish the relationship between the cognitive measures completed by participants and their responses to the phishing emails received during the office task. The best fitting models for each of the response types to the emails are shown in Table 8.7, with those measures not included indicated with a dash. Notable significant findings demonstrate that for the number of emails opened, the best model included the cognitive reflection test ($\beta = -.33$, $p < .05$), and the brief self-control scale ($\beta = .03$, $p =$

Table 8.6 *Correlations between cognitive measures and responses to phishing emails*

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|---------------|--------|-------|------|-------|------|------|------|------|
| 1. Opened | 1.00 | | | | | | | |
| 2. Clicked | .45** | 1.00 | | | | | | |
| 3. Replied | .07 | .18 | 1.00 | | | | | |
| 4. Deleted | -.39** | .00 | -.12 | 1.00 | | | | |
| 5. Recognised | -.71** | -.35* | -.18 | .57** | 1.00 | | | |
| 6. CRT Score | -.36* | -.02 | .01 | .34* | .29* | 1.00 | | |
| 7. Flanker | -.02 | -.02 | .07 | .07 | -.08 | .01 | 1.00 | |
| 8. BSCS Scale | .19 | .01 | -.01 | -.24 | -.21 | -.11 | -.25 | 1.00 |
| 9. BSS Scale | .08 | .13 | .00 | .21 | -.01 | -.12 | -.04 | -.14 |

Note. * $p < .05$, ** $p < .01$ (two-tailed)

.26), with the cognitive reflection test being a significant predictor. This model produced an R^2 value of .17 (adjusted $R^2 = .12$, $F(2, 34) = 3.36$, $p < .05$). When considering the number of emails deleted, the best model included the cognitive reflection test ($\beta = .34$, $p < .05$), the brief self-control scale ($\beta = -.04$, $p = .18$), and the brief sensation seeking scale ($\beta = .44$, $p < .05$), with cognitive reflection and sensation seeking being significant predictors. This model produced an R^2 value of .27 (adjusted $R^2 = .20$, $F(3, 33) = 4.03$, $p < .05$).

Table 8.7 *Summary of best-fit regression models for cognitive measures predicting response to phishing emails*

| | Opened | Clicked | Replied | Deleted | Recognised |
|----------------|--------|---------|---------|---------|------------|
| CRT Score | -.34* | -.23 | - | .34* | .28 |
| Flanker | - | - | .07 | - | - |
| BSCS Scale | .18 | - | - | -.21 | -.20 |
| BSS Scale | - | .19 | - | .32* | - |
| Adjusted R^2 | .12 | .04 | -.02 | .20 | .08 |

Note. Parameter estimates are standardised coefficients.

* $p < .05$ (one-tailed)

8.4 Discussion

The primary aim of this study was to demonstrate an alternative method for assessing susceptibility, which offered a more naturalistic measure of real life email management behaviour, given the ethical constraints of research involving cyber attacks. Using an office simulation scenario that involved email management, susceptibility was measured by various response reactions to a number of email stimuli sent to participants during the experiment. In addition, manipulations of time pressure, priming about phishing emails, and relevance to the office tasks being completed, allowed for a more substantial exploration of the potential situational factors influencing susceptibility. Factors reported as significant predictors of susceptibility in Study 3, were also included here to assess whether findings would be replicated with an alternative measure of email behaviour.

As part of the nature of this study, it was important to identify whether participants were naïve to the purpose, in order to emulate a real office scenario where they would not actively be seeking to differentiate phishing from legitimate emails. Parsons et al. (2013) reported that participants were better at detecting phishing emails when they had been explicitly told that this was the purpose of their task, and so it was important that this remained masked for the integrity of the behavioural assessment. Our findings demonstrate that only two participants were aware of the nature of the study. In order to maintain a realistic measure of susceptibility, these data were removed from further analysis. However, that only two participants were aware that their response likelihood to phishing emails was being measured suggests that the scenario designed for this study was successful in distracting participants from the true purpose of the study. It is hoped that this

provides a more ecologically valid measure of email behaviour, as it is clear that for the majority of participants, attention was focused on the office tasks assigned to them.

The alternative approach in the study design here means that a simple accuracy score as a sum of emails correctly identified is not feasible given the different response reactions that participants could have when managing an actual inbox. Therefore, a set of different response types was recorded to measure participants' susceptibility. The mean number of emails opened by each participant is unsurprisingly high, given that participants were told that email management was one the tasks they were required to complete as part of the study. Whilst opening an email can prove dangerous in some cases, with malware infected images being automatically downloaded, this does not demonstrate the highest level of susceptibility. The number of emails in which the participant clicked the URL or downloaded the attached file on the other hand, demonstrates more susceptibility. For the attachments, downloading the file could install malware on a user's computer, allowing the cyber attacker access to their system, including any confidential information stored there. Although clicking an embedded URL does not necessarily mean that the participant would have gone on to input personal information to a phishing website, it shows the initial susceptibility to follow this link, which can in itself sometimes prove harmful to the person's computer. On the other hand, the number of emails deleted and recognised demonstrate a heightened level of self-defence against phishing. Although deleting an email does not necessarily mean that the participant did not first demonstrate susceptibility to it, this does indicate that they may have become suspicious of the email or are uninterested in

any future response to it. The number of emails recognised as phishing provides a more specific account of why the emails were deleted or not opened, and can be used of an indicator of whether a participant is aware of phishing emails during email management. This does not mean that all participants who did not recognise the illegitimate nature of the email would become victims of the phishing attack, many would ignore it because it wasn't relevant to them or they didn't deem it important enough, but recognition does demonstrate a heightened awareness of such attacks, which is likely to reduce susceptibility in other scenarios as well.

The disparity between the two types of responses – opening or clicking an email vs. deleting or demonstrating awareness of phishing – is emphasised by the negative correlations between these response types. Collectively, the set of phishing emails designed for this study seem to be measuring the same thing. This suggests that as a measure of overall susceptibility, the email management task is able to distinguish between participants who are good at identifying phishing emails and those who are poor at this, regardless of the type of phishing email. However, it should be noted that the levels of susceptibility are based on five phishing emails, compared to eighteen in the original email legitimacy task. In order to maintain the integrity of the study as an office simulation, where employees would be unlikely to receive an extensive amount of phishing emails in a three-hour time frame, the data density in this measure is sacrificed. The extent to which this affects the validity of the task as a measure of susceptibility cannot be quantified from the current data set. Further development of the office simulation could consider a direct comparison between participant responses on a number of different measures of susceptibility.

8.4.1 Situational manipulations. Hypothesis 1 predicted, based on findings from Study 3, that an induced time pressure for completion of the office tasks would lead to increased susceptibility to the phishing emails received during the experiment. However, unlike previous findings, no effect of time pressure was evident in this study. It was thought that increasing time pressure would reduce attention to the details in the phishing emails which identify them as such, as was apparent previously in the thesis, and in research from Yan and Gozu (2012). However, in this study the time pressure was part of the office tasks, not the email judgement - participants were told that they only had 20 minutes to complete each task. This was necessary to ensure that participants remained naïve to the purpose of the study – as an emphasis on rapid response to the email stimuli may have drawn added attention to these. In Study 3, as in Yan and Gozu’s study, time pressure was added to the email task itself. The different nature of these time pressures may explain the differing results observed. When focusing solely on an email task, participants are under pressure to make a decision about each email more rapidly, meaning that they are more likely to make mistakes and ignore cues available to them. In this study though, the participants were given multiple tasks to complete simultaneously and so *they* decided how to manage their time between the office tasks and email management, when under a limited time pressure. It is possible that when time pressure was added to the office tasks, participants’ performance on these would have been jeopardised as opposed to their performance on the emails. In relation to a real-world office environment, this suggests that the impact of time pressure to simultaneous tasks may not be detrimental to email management behaviour, as it seemed that participants simply reacted to the emails when they had

chance around completing the office tasks. There was no evidence to suggest that participants allocated any less attention to these emails than they otherwise would have.

Hypothesis 2 proposed that participants in the priming condition would demonstrate lower susceptibility, as they were alerted to the possibility of receiving phishing emails during the experiment. However, this was only demonstrated when the number of emails deleted was taken as a measure of susceptibility. This suggests that although participants were more vigilant in managing the emails in their inbox, deleting those that they felt were irrelevant or illegitimate, their response likelihood was not affected. This result leads to consideration of the most effective way to reduce susceptibility in users. It is apparent that priming about phishing, in the most basic sense, does not have any systematic impact on behaviour. This supports previous research reported by Downs et al. (2006), which demonstrated that knowledge of security cues to recognise phishing emails did not reduce response likelihood in a role-play email judgment task. Further research has suggested that increased knowledge of security issues is only beneficial in combination with rational cognitive processing (Vishwanath et al., 2011). Even with extensive knowledge about phishing emails, a lack of information processing - whereby users are reliant on intuitive responses without contemplation of the cues available to them – means that response likelihood will remain the same as users without this knowledge. Educational efforts should therefore consider moving away from simply informing users about the issue, and attempt to provide more in depth explanations of how a successful phishing attack is constructed, as well as considering psychological influencers.

An additional factor in the design of this study was the relevance of some of the phishing emails to the office tasks that the participant was asked to complete. Two of the phishing emails were designed to relate to the office tasks, but each was only relevant to two of the four character roles in each session. Hypothesis 3 stated that participants would demonstrate higher susceptibility to an email that was directly relevant to a task they were completing. However, this difference was only demonstrated in one of the phishing emails – Email 3, with participants for whom this email was relevant demonstrating higher likelihood of opening and downloading the attachment. This may be related to the email content, in addition to the relevance of the email. Email 3 was sent from a Lancaster email address, regarding an update to the travel system. Email 4 on the other hand was sent from a Gmail account, advertising a website for cheap office supplies. It is possible that email 4 was more easily identified as a phishing email, given its unprofessional appearance and the email address of the sender. These are factors that have previously been highlighted by participants as cues to phishing in Studies 1 and 5. The lower overall response rates for email 4 support this notion - it seems that participants were less likely to respond to this email, regardless of whether it was relevant to the tasks they were completing or not. Therefore, in relation to email 3, which showed higher overall believability, the influence of task relevance demonstrates that previous findings relating to the effect of this on primary task performance (as discussed by Hadlington, 2015), are somewhat replicated when considering performance on a secondary task (email management, in this case). Relevance is evidently not the only factor determining response likelihood in this case though - as the effect was only

shown on one of two task relevant emails - and should be considered alongside alternative explanations, such as the overall believability of the email.

8.4.2 Cognitive variables. Cognitive signatures of email judgment behaviour from Study 3 were administered here to establish whether these would be relevant in a different, more ecologically valid environment, where participants were treating the emails like they would their own inbox. Based on previous research, outlined in detail in Study 3, Hypothesis 4 proposed that lower levels of cognitive reflection, inhibition, and self-control, and higher levels of sensation seeking behaviour, would predict higher susceptibility to the phishing emails during this experiment.

As discussed above, there are a number of different ways of measuring susceptibility from how participants responded to the emails designed for this study. This variation in how susceptibility is measured may provide some explanation for the lack of consistency in the cognitive factors highlighted as predictive of each of these response types. All of the factors found to be predictive in study 3 were included in the best-fit model for at least one response type in this study, but only cognitive reflection and sensation seeking were found to be significant predictors. Self-control was also included in this study, despite no evidence that this related to susceptibility in Study 3, to establish whether results from Holtfreter et al.'s (2010) study would be replicated in this alternative scenario. However, this was not found to be predictive for any of the response types. This at least shows consistency of null effects across all of the thesis studies in which it was included.

The cognitive reflection test and the Flanker task are two alternative measures of inhibition and impulsive behaviour that have been linked in earlier studies to email decisions. Previous findings have implied that a person who is less able to inhibit an impulsive response or reflect upon this to reach an accurate decision, is also less able to accurately judge the legitimacy of an email. Based on response to the cognitive reflection test, findings suggest that inhibition has a role in the initial opening of the email, which is done with little systematic thought, and on likelihood to delete the phishing emails, which may be due to a lack of more in depth contemplation of how best to deal with the email, and rather a contentment to disregard the email and ignore it rather than take action to get rid of it. Although the Flanker task is also included as a measure of inhibition and impulsivity, the cognitive reflection test requires the participant to engage in contemplation of the problem in order to reach the accurate conclusion, whereas the Flanker task requires only an initial suppression of the intuitive response. This may explain why findings are not consistent across these two measures.

These findings suggest an increased susceptibility, across various methods of analysing this concept, in those participants with reduced impulsivity control during decision-making. This may be explained, as discussed in Study 3, by an inability to engage in more rational, system 2 processes, which may allow the consideration of all aspects of the email, including the consequences of responding to this based on prior knowledge of phishing emails. As discussed in the previous studies though, the amount of variance explained by performance on the cognitive reflection test was limited, and so a number of other factors, some of which are discussed through

manipulations in this experiment, are also influencing the response decisions made by participants.

Sensation seeking was also found to be a significant predictor of the number of phishing emails deleted, with those demonstrating higher levels of sensation seeking also more likely to delete more phishing emails. This is contradictory to what would be expected, as those who are more likely to engage in risk-taking behaviour were expected to engage more actively in behaviour without considering the negative consequences, as demonstrated in Study 3. It is possible that this is related to a methodological difference, rather than actual susceptibility. Participants were not given instructions to delete emails that they did not trust, and therefore the decision to delete emails may have been higher in those participants demonstrating higher sensation seeking as they are doing something that was not explicitly outlined to them, and may not be how the researcher wants them to respond.

In Study 5, there was no significant evidence of a difference in performance between native and non-native English speakers. However, Study 6 provided contradictory evidence, suggesting that there was a difference, with non-native speakers demonstrating higher susceptibility. Hypothesis 5 predicted that native speakers would demonstrate lower susceptibility, which was supported in this experiment, as in Study 6. Both Study 6 and the current study involved completion of unrelated simultaneous tasks, and so it is possible that the increased cognitive load involved in these takes away attention from the email management aspect of the experiments. Non-native participants may have overlooked some of the cues available to them to distinguish the phishing from legitimate emails, instead relying on intuitive responses that do not require extensive translation of the email content.

8.4.1 Limitations and future research. Whilst the phishing emails sent to participants in this study were designed for the purpose of the study only, these were based on the design of genuine phishing emails, and so are fairly representative of those that a user would receive in their own inbox. Further to this, the fact that participants were unaware of the nature of the study makes the responses more genuine, in the sense that they are not deliberately focusing attention on distinguishing between phishing and legitimate emails. However, there is still the issue of this being a laboratory-based study, in which there are no negative consequences for participants if they chose not to respond to an email. The mean self-reported motivation score was found to be just below 4 on a 6-point scale, demonstrating that participants were somewhat motivated to perform well on the email legitimacy task, but there is no evidence to suggest how well this reflects real world email management behaviour. With a topic as sensitive as vulnerability to fraud though, this methodology provides an ethical compromise, in which participants are not put at risk or embarrassed by their response decisions, as this is a simulation and not their own inbox, but at the same time they are able to interact with an actual inbox in the role of a character whose job it is to do so.

The partial replication of findings from Study 3, which demonstrate the predictive nature of cognitive reflection and sensation-seeking behaviour in susceptibility to email fraud, provides a more solid grounding for future considerations in efforts to educate users in order to reduce this susceptibility. The relationship between task relevance in an email from a trustworthy source, which could have been hacked, emphasises previous findings that educational efforts in tackling fraud susceptibility should focus on spear phishing attacks which are

specifically targeted at a user, with information relevant to them. The finding that time pressure and priming were not influential on likelihood to click the URL or download the attachment in an email, dismisses some previous assumptions in the field that may have influenced training efforts. Merely warning participants that phishing emails exist, and that they should take time to consider the email before deciding to respond is not sufficient to reduce their susceptibility.

8.4.2 Conclusions. This study introduces a novel methodology for measuring susceptibility, which ensures that participants are naïve to the nature of the study and allows interaction with an actual inbox, thus providing a more realistic measure of email management behaviour. Findings from Study 3 regarding cognitive factors influencing susceptibility were somewhat replicated in this study, across the different response types to the email stimuli. However, as previously reported in Study 3, the amount of variance explained by these cognitive factors is small, and so there are still a number of additional factors to be considered in explaining individual differences in susceptibility. Some of these were addressed in this experiment, with variables assessing the importance of situational factors such as time pressure, email relevance, and priming about an influx of phishing emails.

Although it would be useful to replicate findings from this and earlier studies in a scenario where participants demonstrate susceptibility to their own personal details, and personal inboxes, this study begins to bridge the gap by providing a methodology which measures email behaviour in a simulated office environment, whereby participants are asked to manage an email inbox in the same way they would if it were their own. The findings build upon the cognitive profile outlined in

previous studies in the thesis with findings regarding the relevance of an email to a concurrent task being completed, and null effects of priming participants to be aware of phishing emails. Together, these can be utilised to develop understanding of the decision-making process and to aid efforts to reduce susceptibility.

Chapter 9

General discussion

Chapter summary

This chapter provides an overview of the key findings from the thesis, highlighting aspects of email content, cognitive processing, and cognitive make-up that were found to affect performance in measures of susceptibility to email fraud. The practical applications of these findings are outlined, with particular emphasis on how findings might be used to enhance efforts to educate and protect users from victimisation. Limitations to the research conducted in the thesis are also outlined. Finally, conclusions regarding the implications and insights to be taken from the thesis are discussed.

9 General discussion

9.1 Background and main aims of the thesis

The main aim of this thesis was to offer a psychological perspective on email decision-making, and individual differences in ability to recognise authenticity and plausibility. A review of the relevant literature, reported in Chapter 1, established that knowledge about how decisions regarding email legitimacy are made is fairly limited, especially from a psychological perspective. Although there are a small number of studies that look at individual psychological constructs and their contribution to decision-making, no comprehensive theoretical explanation exists for this. Bringing together the little research that does exist to suggest that psychology can contribute to understanding susceptibility, three broad theoretical approaches were outlined at the outset of this thesis.

Firstly, the persuasive techniques employed by the sender in the email content are considered, with insights from the psychology of persuasion to suggest that social conformity/familiarity (Jagatic et al., 2005), authority (Guéguen & Jacob, 2002; Workman, 2007), and scarcity (Langenderfer & Shimp, 2001) may impact response likelihood. Secondly, situational influences that affect cognitive processing are thought to influence the decision-making process. Users who are reliant on intuitive responses when judging email legitimacy, which may be induced by scenarios such as multi-tasking or limited time constraints, have been shown to make more errors (Yan & Gozu, 2012; Harrison et al., 2016). Finally, the cognitive make-up of the individual user is considered, with previous research demonstrating links between fraud susceptibility and personality (Modic & Lea, 2011), as well as self-control (Holtfreter et al., 2010; van Wilsem, 2013). Further to this, additional

variables are highlighted as potential predictors of email behaviour based on links to impulsive behaviour or risky decision-making (e.g. sensation seeking, Hoyle et al. (2002); cognitive reflection, Frederick (2005)).

Through a number of empirical studies in the thesis, several important effects and phenomena have been established, whilst others have been disputed or found to be relatively unimportant. These findings have advanced a theoretical account of email decision-making, as well as offering insight for applied work and methodological approaches. In this chapter, these key findings will be reviewed, highlighting the psychological constructs that were found to be relevant to email decision making, the extent of these findings, the psychological constructs that were deemed to be less relevant, situational constraints that impair accuracy, and the differences in experimental approaches that can affect measurements of susceptibility. This will be followed by a discussion of the theoretical and methodological implications of these factors for email decision-making, limitations of the experiments, and prospects for further research.

9.2 Understanding the issue of email fraud victimisation

To begin with, a number of studies were conducted to explore how users viewed email fraud – their familiarity with the issue, personal experience of it, and approaches to protecting themselves against it. In Study 1, qualitative data from focus group discussions demonstrated that although participants had a clear understanding of email fraud, they did not necessarily feel protected against it. There were also certain aspects where participants demonstrated a slight naivety, for

example mentioning that they would be likely to trust any email that came from a lecturer's email account, as they are not likely to get hacked.

Further to this, Study 1 highlighted a reliance on generic cues to phishing, such as grammatical errors and unprofessional appearance. Based on limited media coverage and discussion with other internet users, participants' knowledge on how to recognise anything other than the most basic phishing emails seemed fairly minimal, emphasising the need for improved educational efforts to reduce victimisation. This was also emphasised by an assumption from participants that either older users or very young users who have no internet experience are more susceptible to fraud. There is no consistent evidence regarding age and susceptibility though, with some research demonstrating that older users are more susceptible (e.g. Shadel & Pak, 2007; Pak & Shadel, 2011), whilst other research implies younger users are more so (e.g. Sheng et al., 2010; Pratt et al., 2010). Therefore, such stereotypes about who is most likely to become a victim of fraud are unfounded.

Participants demonstrated a lack of knowledge regarding more sophisticated phishing attacks, in particular spear phishing, which may lead to overconfidence in their ability to recognise phishing attacks. In turn, this could leave them more susceptible to a sophisticated attack targeted towards them. This initial exploratory study provided support for the need to research explanations of susceptibility to email fraud, in order to improve efforts to reduce this.

The cues employed in email decision-making were explored further in Study 5, where a think aloud protocol was instigated whilst participants completed the email legitimacy task, providing additional qualitative data. This highlighted a number of commonly used cues to phishing, some of which again relied on out-dated

assumptions about elements of an email that are often spoofed in more sophisticated phishing attacks – such as the email address of the sender, and professional appearance of the email.

Findings from each of these studies suggest that although users have an awareness of phishing as a serious cyber security issue, their knowledge about how to protect themselves, especially from the more sophisticated attacks, may not be optimal. This highlights a level of susceptibility in itself, and encourages the need for research that works towards reducing this.

9.3 Measuring susceptibility

As discussed in Chapter 1, the paucity of literature assessing psychological influences to fraud susceptibility means that there is no consistently used and widely validated measure for assessing susceptibility. Research conducted to date has used a variety of methods, including a questionnaire approach (Modic & Anderson, 2014); forced-choice legitimacy judgment tasks similar to the one developed for this thesis (e.g. Holtfreter et al., 2010; Yan & Gozu, 2012); and simulated phishing attacks (e.g. Guéguen & Jacob, 2002; Jagatic et al., 2005). Each of these methodologies has advantages and disadvantages. For example, the judgment tasks allow more control in an experimental environment, whilst the simulated attacks provide a potentially more valid measure of real world susceptibility.

For the purpose of this thesis, an email legitimacy task was developed, providing an ethically sound and controlled measure of how participants respond to email stimuli. The emails included were genuine examples of phishing and legitimate emails, which participants were asked to judge on their confidence that each was

either phishing or legitimate. Study 2a provided some initial pilot data on a few of the email stimuli to be used later on in the thesis, demonstrating that some participants were more accurate in their judgments than others, and also that some of the emails were more difficult to judge than others.

The full version of this legitimacy task, with 36 email stimuli, was introduced in Study 3. The data from this study found few participants who demonstrated a response bias based on their use of the scales, by being either overcautious and marking most emails legitimate, or being under cautious and marking most emails phishing. In addition, when using the response bias calculation from signal detection analysis, few participants were found to show a bias in their response behaviour. These findings suggest that participants engaged with the task and were making an effort to differentiate between phishing and legitimate emails.

Variations on this email task were used through the thesis. The mean accuracy and confidence scores for each of these variations can be seen in Tables 9.1 and 9.2, with indication of where there were significant differences in performance between the experiments. These demonstrate generally higher accuracy and confidence in Studies 3 and 5 compared to the other experiments. Possible explanations for these differences, relating to the nature of manipulations included in the tasks, will be discussed in more detail in section 9.5.

Study 3 incorporated an equal mix of phishing and legitimate email stimuli, which is not necessarily representative of the mix that users are used to managing in their own inbox. In Study 4, the proportion of phishing and legitimate emails included in the stimuli set was manipulated to assess how this affected legitimacy judgments, and how participants manage and compare stimuli in a categorical

Table 9.1 *Mean accuracy and confidence score for the email legitimacy task across studies 3 to 6*

| Exp. | N | Mean accuracy | SD | Mean confidence | SD |
|------|-----|---------------|------|-----------------|-------|
| 3 | 224 | 24.57 | 3.43 | 33.08 | 17.14 |
| 4 | 98 | 22.69 | 4.73 | 23.03 | 23.21 |
| 4a | 60 | 22.48 | 5.28 | 21.90 | 27.24 |
| 5 | 51 | 26.84 | 3.18 | 47.41 | 16.94 |
| 6 | 102 | 23.04 | 4.56 | 26.68 | 22.49 |

Table 9.2 *Mean differences in accuracy on the email legitimacy task for each experiment*

| Exp. | 3 | 4 | 4a | 5 |
|------|---------|---------|---------|----------|
| 4 | -1.88** | | | |
| 4a | -2.09** | -0.21 | | |
| 5 | 2.27** | 4.15*** | 4.36*** | |
| 6 | -1.53* | 0.34 | 0.55 | -3.81*** |

Note. Scores for study 6 exclude participants in the control condition;

* $p < .05$, ** $p < .01$, *** $p < .001$

judgment task. The majority legitimate condition was designed to be more representative of a genuine inbox. By comparing performance between conditions with more legitimate, more phishing emails, or equal numbers of each, accuracy was found to be lower in both of the varied proportion conditions. This may be due to the violation of expectation that participants have of seeing equal amounts of each type of stimuli in a categorical judgment task impairs performance. Based on an assumption of a fifty-fifty chance likelihood of each email being phishing, participants may be monitoring how many times judgments of each type are made and comparing between stimulus, to keep these as close to equal as possible. An alternative explanation may be that participants found the additional stimulus included in this version of the email legitimacy task were more difficult to accurately judge.

In order to assess the reliability of the email legitimacy task as a consistent measure of susceptibility, test-retest data was collected in Study 4a. This demonstrated a significant positive relationship between response accuracy at test 1 and test 2, although this did not quite reach a good level of reliability, as outlined by Cicchetti (1994). This analysis comes from data using the varied proportions version of the email task though, meaning that it may not be representative of reliability in the task when equal proportions of email stimuli are used, especially given the differences in performance between these variations. Further assessment using the original email legitimacy task would be necessary to establish reliability for this.

For all of the studies using the email task, an accuracy score was generated based on the number of emails correctly identified, including both legitimate and phishing emails, as well as a confidence score based on the rating scale used to make judgments. For Study 3, signal detection theory was also reported as a measure of how well participants were able to differentiate phishing from legitimate emails. In doing this, the focus is on phishing emails, whereas the accuracy score accounts for how well a participant can recognise both phishing and legitimate. In relation to a real inbox, although the ability to recognise phishing emails seems a more pressing issue with more severe consequences, disregarding a legitimate email as phishing may also have adverse effects. For example, a legitimate email with information about updates to online banking, or about updating security software may be crucial to maintaining a secure online presence. Although not reported in order to keep results concise, the use of signal detection theory across all studies produced very similar results in terms of how cognitive and situational factors predicted ability to differentiate between the two stimulus types. Therefore, it was decided that the

accuracy score would be used throughout as the binary measure of susceptibility, as this also took into account ability to recognise legitimate emails.

Given some methodological limitations of the email legitimacy task, discussed in more detail in section 9.8, this thesis also introduced an alternative, potentially more ecologically valid measure of susceptibility. Developing upon role-play tasks used by some researchers to measure susceptibility (e.g. Downs et al., 2007; Parsons et al., 2015), an office simulation scenario was developed, during which participants were naïve to the true purpose of the study, having been told that this was a study of office behaviour and task management. This allowed decisions about emails received to be recorded without participants actively seeking to differentiate between phishing and legitimate emails, as they did previously in the email legitimacy task. As participants were actively engaging with an email inbox, this meant that a number of measures of response likelihood could be assessed, all measuring susceptibility to a different extent. Across these different measures of susceptibility, cognitive reflection and sensation seeking were found to be predictive, replicating findings from Study 3. The importance of these cognitive factors will be discussed more in section 9.4. However, what this demonstrates is that the two alternative methods for measuring susceptibility produce similar results. This suggests that in situations where there are time constraints or it is impractical, the email legitimacy task provides a more easily administered, web-based alternative that is measuring the same as methods that enhance ecological validity by allowing participants to interact with an email inbox similarly to how they would in real life.

9.4 Assessing the influence of cognitive make-up on susceptibility

Previous research has highlighted the influence of agreeableness and extraversion (Modic & Lea, 2011), as well as self-control (Holtfreter et al., 2010; van Wilsem, 2013) on fraud victimisation. However, there is a wealth of literature from other areas of psychology that could also be considered in relation to email fraud susceptibility. For example, risky decision-making and working memory capacity literature provide some potential explanations for why some internet users demonstrate higher susceptibility to email fraud than others. One possible explanation comes from a dual-systems approach to reasoning and decision-making (e.g. Stanovich, 1999; Kahneman, 2000; Evans, 2003), which would suggest that there are two systems available for generating a behavioural response. System 1 relies on intuitive, impulsive inclinations about the given stimuli, whilst system 2 implies a more rational, and analytic consideration. In Chapter 1, it was proposed that cognitive influences such as lower working memory capacity, or lower levels of inhibition, may indicate a reliance on system 1 decision-making processes due to a difficulty in engaging system 2 responses (Barrett & Tugade, 2004). Such factors may make some users more susceptible to email fraud as they are less likely to consider the cues to legitimacy and the consequences of responding to an email.

Based on previous research and predictions about dual-system decision-making, a battery of nine cognitive tasks was used throughout the thesis, measuring a range of variables including inhibition, self-control, and working memory. Each study assessed all or some of these variables, with findings across the studies summarised in Table 9.3. In Study 3, eight tasks were administered, with three of the variables found to be predictive of susceptibility – cognitive reflection, inhibition, and

sensation seeking. These support the notion of a dual-system approach to email decision-making, with participants who demonstrated an inability to suppress an intuitive response in the cognitive reflection and Flanker tasks performing with lower accuracy on the email task. Further to this, higher scores on the sensation seeking scale predicted lower accuracy, suggesting that these participants do not contemplate the consequences of their decisions, leading to a reliance on system 1, intuitive responses.

Again, the same eight tasks were administered in Study 4, but this time no significant predictors of susceptibility were found. As discussed above, accuracy on the email task in Study 4 was lower due to the varied proportions of phishing and legitimate emails, so this may explain a lack of replication for the cognitive predictors assessed across these two experiments. Similarly, when an additional task, the Moses illusion, was assessed in Study 4a there was no evidence to suggest that this was predictive of accuracy on the email task. However, this was later found to be close to significance in Study 5, suggesting that the alternative approach to measuring susceptibility, with varied proportions of email stimuli, may have affected the outcome in Study 4a. Despite the small predictive power of the Moses illusion in Study 5, this suggests that an ability to recognise distortion and suppress an intuitive response may contribute somewhat to email fraud susceptibility.

Only the Cognitive Reflection Test was included in Study 6, which was found to be predictive of accuracy on the email legitimacy task. This replicates the finding from Study 3, providing further support for the relevance of this task to email fraud susceptibility. Further support comes from Study 7, which found higher cognitive

Table 9.3 Summary of best-fit regression models for cognitive measures predicting response accuracy across experiments

| | Experiment | | | | | | | | | |
|---|-------------|------------|-----------|------------|-----------|-------------|------------|------------|------------|------------|
| | 3 | 4 | 4a | 5 | 6 | 7 | Deleted | Recognised | | |
| Cognitive Reflection Test | .31** | - | | - | 1.30** | -.34* | -.23 | - | .34* | .28 |
| Flanker Task | -.21* | - | | | | - | - | .07 | - | - |
| Brief Self-control Scale | | | .06 | | | .18 | - | - | -.21 | -.20 |
| Brief Sensation-seeking Scale | | -.23* | - | | | - | .19 | - | .32* | - |
| Agreeableness (IPIP) | -.13 | - | | | | | | | | |
| Neuroticism (IPIP) | - | -.26 | | | | | | | | |
| Brief Need for Closure Scale | - | -.26 | | | | | | | | |
| Moses Illusion | | | .36 | .67 | | | | | | |
| R ² (Adjusted R ²) | .15** (.12) | .05* (.04) | .08 (.04) | .01 (-.01) | .07 (.05) | .10** (.09) | .17* (.12) | .09 (.04) | .00 (-.02) | .27* (.20) |
| | | | | | | | | | | .14 (.08) |

Note. Parameter estimates are standardised coefficients.

** $p < .01$, * $p < .05$

- indicates a measure not included in model of best-fit.

Measures included in the thesis but not included in any best-fit regression models: Stroop task, Openness (IPIP), Conscientiousness (IPIP), Intellect (IPIP), Reading span task.

reflection to be predictive of less phishing emails being opened and more being deleted during the office simulation. In addition, Study 7 demonstrated that higher levels of sensation seeking predicted the number of phishing emails deleted by participants during the office simulation, suggesting more security conscious behaviour. This is contrary to findings from Study 3, which suggested that higher sensation seeking was associated with higher susceptibility.

Although there are some inconsistencies in the cognitive variables found to predict susceptibility to email fraud across the studies, it is possible that these are due to methodological differences between the experiments. The most consistent finding throughout the thesis is the cognitive reflection test as a predictor of response likelihood. This measure assess elements of inhibition, with participants required to suppress an initial impulsive response in order to reach the correct one through more rational, analytic thinking. Dual-system approaches to reasoning, as discussed above, may provide an explanation of these findings. Those participants who were better able to repress the intuitive responses in this task also demonstrated an increased ability to accurately differentiate between phishing and legitimate emails. This success may be explained by an ability to engage system 2 decision-making processes in a more efficient manner than other participants, allowing consideration of all information and cues available in an email, which may aid the decision. One basic difficulty should be noted in the application of dual-system theories though, which is that there is no objective measure of how these are affecting decision-making. Assumptions can be made about the relevance of these to factors such as cognitive reflection and inhibition, but there is no definitive evidence that intuitive and rational decision-making processes are influencing performance on

cognitive measures, or on the email legitimacy task, and so at this stage the relationship remains theoretical.

In addition, null findings demonstrate a number of cognitive variables that do not contribute to predicting susceptibility, thus narrowing down the search for explanatory variables. Some previous research findings relating to fraud victimisation were not replicated in the thesis. Modic and Lea (2011) demonstrated that agreeableness and extraversion were predictive of past response behaviour, and Holtfreter et al. (2010) reported a relationship between self-control and reported response likelihood to a number telemarketing fraud scenarios, but neither of these findings were replicated in Studies 3, 4, or 7 (included self-control only). This may be explained by methodological differences in assessing susceptibility – with Modic and Lea employing a measure of self-reported past victimisation, rather than hypothetical response likelihood, and Holtfreter et al. focusing on telemarketing fraud. It is possible that these measures are not assessing susceptibility in the same way as those employed in this thesis.

Additional variables that showed no relationship with susceptibility included cognitive closure, and working memory span. Although these findings are not conclusive given the nature of how susceptibility was assessed, there is no evidence to suggest that these variables should be considered in future efforts to understand response behaviour.

Despite consistent support from the different measures of susceptibility employed in this thesis to suggest that some cognitive variables reported play a role in the decision-making process, as seen in Table 9.3, the variation in performance explained by these variables is fairly low. Therefore, it is important to consider how

other factors may contribute to the decision-making process, in addition to differences in cognitive make-up between participants, in order to develop a more comprehensive understanding of susceptibility.

9.5 Assessing situational influences on susceptibility

As discussed above, performance on a number of cognitive tasks goes some way to explaining individual differences in susceptibility to email fraud, as measured by the email legitimacy task or the office simulation. However, it is apparent that this is only a partial explanation, given the small amount of variance associated with these factors. Therefore, it is crucial that other influences are considered in the effort to understand why some internet users demonstrate higher susceptibility than others. Across the studies in this thesis, a number of situational factors were considered, in addition to the cognitive elements, as potential predictors of email behaviour.

In Study 3, a time pressure manipulation was included in the email legitimacy task, which demonstrated lower accuracy when participants were given a limited time frame in which to complete the task, compared to a control group who received no time pressure. This manipulation can be considered in two real life scenarios – either when users are managing emails under increased cognitive load from pressure to complete other tasks, or when an email itself portrays a sense of urgency and need for immediate response. The finding from Study 3 replicates a similar task reported by Yan and Gozu (2012), where participants were instructed to either make decisions rapidly for each email presented to them, or to take time and rationally think through their decisions. However, in Study 7 a time pressure manipulation was

also included, whereby participants were told that they had a limited amount of time to complete each of the office tasks during a simulation scenario, but there was no evidence to suggest that this affected susceptibility to the phishing emails received. In this experiment, the time pressure related to the office tasks, whereas in Study 3 and in Yan and Gozu's research, the time pressure related specifically to the email task. It was predicted that, similar to a real world office setting, time pressure on the office tasks may lead to participants making more impulsive, rapid decisions when emails came in, so that they could return to the office tasks as promptly as possible. It seems that email decision-making is only affected when the increased time pressure related directly to that task though. In Study 7, it is possible that the time pressure on the office tasks distracted participants away from the email management element of the study, so findings may relate more to participants not paying attention to the emails and thus not responding at all, rather than accurately recognising them as phishing.

In order to further assess how alternative situational factors that may increase cognitive load affect susceptibility, participants in Study 6 were asked to complete a secondary task, alongside the email task. With both verbal and motor type secondary tasks included, results demonstrated that the verbal task conditions (both simple and complex) resulted in lower accuracy compared to the motor task conditions and a control group with no secondary task. This somewhat supports the previous findings regarding the impact of increased cognitive load on the decision-making process. However, it is evident that this effect is dependent on the type of cognitive load induced. The verbal tasks may have required more effort to complete simultaneously with the email task, than the motor tasks did, or may have interfered

with sub vocalisation to a greater extent (Slowiaczek & Clifton Jr., 1980; Daneman & Newson, 1992), making it more difficult for participants to comprehend information in the email stimuli.

One additional explanation for the findings relating to secondary verbal tasks may come from a body of literature suggesting that secondary tasks relating to the primary task make it more difficult to divide attention between the two (Kane et al., 2001). In Study 6, participants in the verbal secondary task conditions were required to simultaneously count backwards out loud – a task that may be more closely related to reading on the email task, than the motor secondary task was. If this is the case, this may explain why verbal tasks affected accuracy, whilst the motor tasks did not. However, this explanation may be undermined by the increased accuracy demonstrated by participants in Study 5, compared to the other studies using the email legitimacy task. In this experiment, participants were required to think aloud whilst completing the task – again, engaging verbal skills in a task that was this time directly related to the email task. Therefore, it seems that while a secondary task requiring reflection on the email stimuli can improve accuracy, a secondary task that involves the same processing skills as the email task, but without requiring direct contemplation of the stimuli may have adverse effects on accuracy. This may be because there is less need to divide attention when the secondary task relates directly to the primary as this aids information processing, and so the notion of divided attention regarding unrelated secondary tasks should not be dismissed without further investigation.

These considerations can be taken into account alongside dual-system decision-making theories. In this sense, it might be that the time pressure and verbal

secondary tasks result in increased cognitive load due to a need to divide attention effectively, leading participants to rely on intuitive, system 1 responses to email stimuli as they are unable to engage system-2 reasoning. Whereas, the think aloud protocol encourages them to engage more rational decision-making processes, by contemplating why they are making each decision.

In relation to real world email management, the impact of increased cognitive load is an important issue to understand. Although findings suggest that urgency to complete unrelated tasks does not have adverse effects on email management, as demonstrated by time pressure on the office tasks in Study 7, increasing cognitive load whilst completing the task of email management itself does. Often, phishing emails are designed to emulate a sense of urgency, encouraging the recipient to respond quickly without contemplating the consequences. The findings from this thesis would suggest that techniques like this, which induce an added cognitive load to the processing of an email, might improve the success rate of the attack. In order to investigate this possibility further, it would be beneficial to design email stimuli in which urgency is manipulated to establish whether effects found from an induced cognitive load are replicated in this scenario.

9.6 Assessing persuasive techniques employed in phishing emails

Although not a main focus, this thesis also considers some elements of phishing emails themselves that might impact their persuasive power on the recipient. There are two main factors that were highlighted across studies as being important to the decision-making process – familiarity and relevance to the recipient. In Study 1, data from focus group discussions demonstrated that

participants would be more trusting of an email that purported to come from someone familiar to them, especially if this was a person of higher authority, such as a lecturer. Similarly, in Study 5 participants reported feeling more confident in assessing emails from companies they recognised, as they felt that they could identify anomalies in the appearance or the content of these emails. Although this worked in some cases, it was not always an accurate method for detecting phishing, as emails can easily be modified to mirror a genuine email from a given company. This is demonstrated by the consistent misidentification of a sophisticated phishing email purporting to come from Amazon, which mirrors the exact format of an Amazon email, but is sent from an un-associated email address. In more sophisticated phishing attacks, even the email address of the sender can be spoofed, leaving fewer cues to assist the user in recognising the fraudulent nature of the email.

The relevance of the email to the recipient was also a key cue for participants in Study 5 during the decision-making process. Although the availability of information about context was highlighted as a methodological issue (addressed more in section 9.8) given the nature of the email task, this is still an important consideration for educating users about spear phishing attacks. Whilst an irrelevant email may be detectable to a user as phishing, relevance cannot necessarily be relied upon as a cue to judge an email as legitimate. Spear phishing attacks target users based on personal information collected about them, meaning that the emails are often very relevant to them. These techniques are likely to make an email attack more successful if users believe that emails that are relevant to them are always legitimate. Therefore, suggesting that although users are often able to engage

decision-making rationale that might help them detect the more obvious phishing emails, these do not always translate to spear phishing attacks.

9.7 Major contributions from the thesis

9.7.1 Theoretical contributions. This thesis aimed to contribute findings from a psychological perspective to expand upon an understanding of individual differences in susceptibility to email fraud. Whilst previous research demonstrated the potential influence of some psychological variables, such as self-control (Holtfreter et al., 2010) and agreeableness (Modic & Lea, 2011), no comprehensive profile of the cognitive and situational factors influencing susceptibility has been reported. In Chapter 1, a diagram was presented to demonstrate proposed theoretical explanations of susceptibility, including potential examples. The approaches outlined in this diagram were addressed across the thesis, and findings have allowed for the development of an evidence-based version of this, seen in Figure 9.1. This updated version also includes examples of methodological approaches employed to assess each of the concepts.

The findings from this thesis provide support, across multiple experiments, of cognitive factors that seem to differ between users, influencing susceptibility. This in itself provides support for the consideration of psychological perspectives in addressing the issue of email fraud victimisation. Given the paucity of previous literature in this area, a range of cognitive factors were included at the beginning of the thesis in order to narrow down the relevant constructs. As well as highlighting predictors of email behaviour, this also provided insight into the cognitive factors that are less important in protecting users against email fraud.

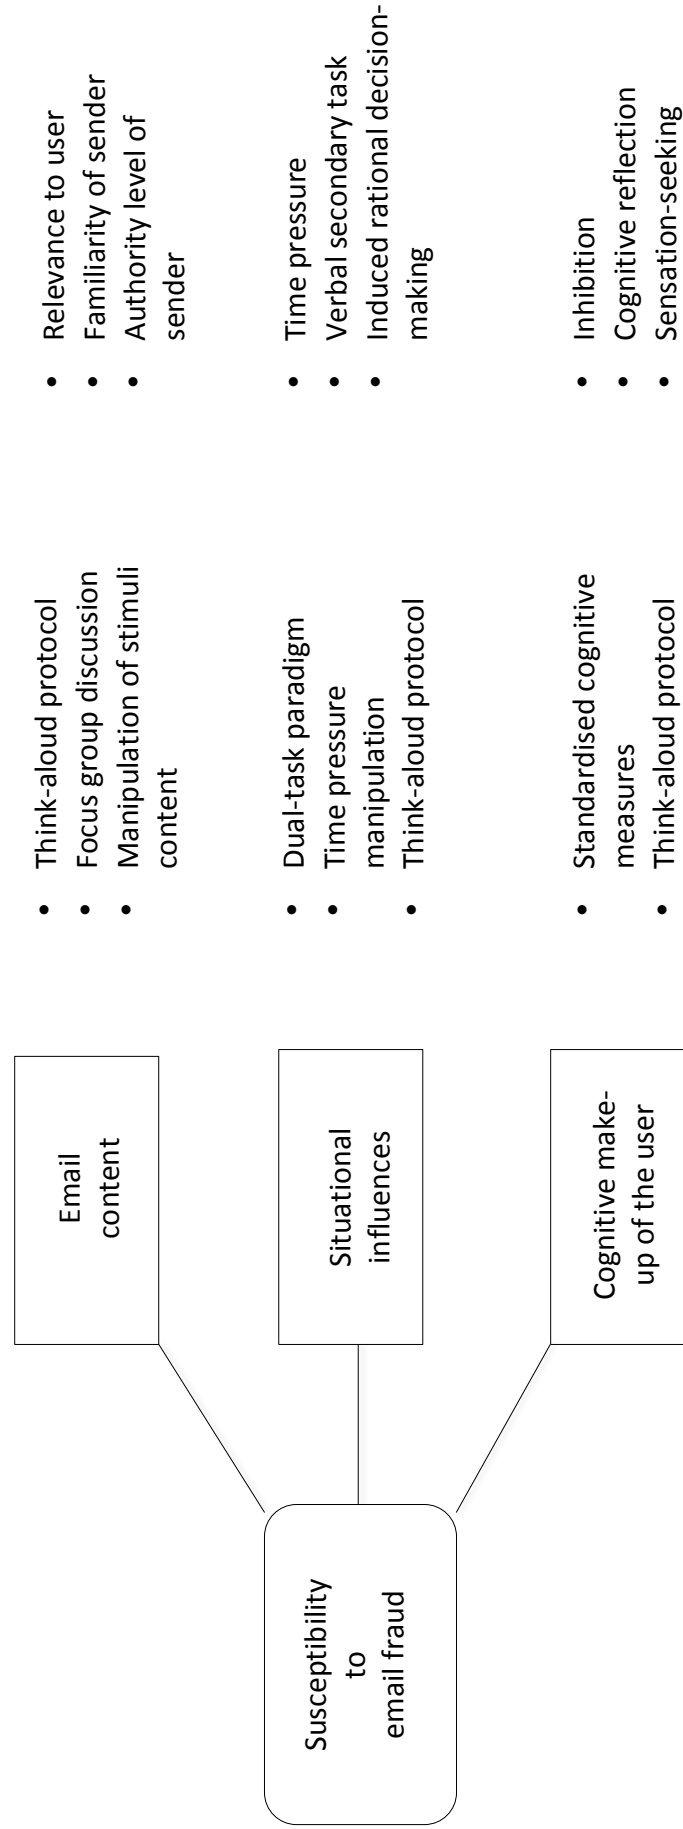


Figure 9.1 Diagram of evidence-based theoretical approaches to understanding individual differences in susceptibility to email fraud

Consistent findings across the studies in the thesis demonstrated that factors relating to an ability to suppress an impulsive response to a problem were found to predict higher accuracy on the email legitimacy task. In line with this, consideration of situational influences suggests that attempts to make decisions about email legitimacy whilst under increased cognitive load could be detrimental to accuracy. These differences both within participants in different situations, and between participants who demonstrate different cognitive abilities, may be explained by the cognitive strategies employed in decision-making.

One explanation for the differences demonstrated may come from dual-system theories of reasoning. These propose that some users are consistently better able to engage more rational, system 2 strategies for decision-making whilst others rely more heavily on impulsive, system 1 strategies, leading to more systematic errors (Tversky & Kahneman, 1975). Rational decision-making is known to relate to a higher working memory capacity (Kyllonen & Christal, 1990; Markovits, Doyon, & Simoneau, 2002), and this in turn has been linked to inhibitory capacity (Engle, 1996; Redick, Heitz, & Engle, 2007). The findings reported in this thesis would suggest that when a user's inhibitory capacity is limited - that is, they are unable to suppress an impulsive response - ability to see past an intuitive response regarding email legitimacy is also limited. Such users may consistently overlook cues that identify an email as suspicious as they do not process all information available to them or consider the negative consequences of a decision to respond in an impulsive manner. Study 5 provides support for this explanation, with qualitative data demonstrating that some users were consistently reliant on a single cue to

legitimacy, whilst others spent time contemplating numerous aspects of the information available to them when judging each email stimuli.

At the same time, the decision-making process can be influenced by situational factors that mean users, who might otherwise have engaged in more rational processing, are reliant on an impulsive response. The thesis showed evidence of impaired accuracy in email decision-making when participants were under induced time constraints, or their attention was divided between multiple tasks. This supports previous research demonstrating that participants demonstrate increased susceptibility to phishing emails when they are reliant on intuitive responses (Yan & Gozu, 2012; Harrison et al., 2016). Situational influences on susceptibility may also be related to ability to divide attention. In particular, the finding that participant accuracy is impaired whilst completing a secondary verbal task, may demonstrate the negative impact of trying to divide attention between multiple tasks at once. This has practical implications, discussed in more detail below, for the impact that multi-tasking behaviours such as this can have on susceptibility.

Although the level of variance in email task performance explained by these factors remains relatively low across experiments, the replication of findings with alternative methodologies demonstrates that they provide a solid grounding for establishing a more substantial explanation of email fraud victimisation in future research.

9.7.2 Methodological contributions. In addition to the insights into individual differences in email fraud susceptibility, the thesis employed a mixed

methods approach, demonstrating the value of a number of research paradigms to the assessment of email fraud susceptibility. Incorporating both quantitative and qualitative data collection methods allowed for alternative approaches to addressing the aims of the thesis and considering different aspects of susceptibility. Whilst the quantitative data provide evidence of relationships between cognitive factors and situational influences in email decision-making, the qualitative data provide a complementary source of evidence, which also considers the effect of different persuasive techniques employed in the emails themselves. As well as the different types of data collected, the methods used varied across the experiments. Methods that have not previously been used in relation to email fraud susceptibility were employed, such as focus group discussions, dual-task paradigms, reliability analysis, and an office simulation experiment. The use of different methodologies provided opportunity to assess susceptibility and factors influencing this under a number of different circumstances. Convergent findings across these different methodologies suggest more robust findings regarding individual differences in cognitive make-up between users, as well as situational factors influencing email decision-making.

In addition to the methods employed to assess factors influencing the decision-making process, this thesis has also contributed two novel measures of susceptibility. The email legitimacy task provides a practical, ethically sound methodology that can be administered either in the lab or online. Although there are a number of limitations to this methodology as a measure of susceptibility, it allows flexibility in terms of additional manipulations that can be administered, and variations in the exact stimuli included. Although the findings regarding test-retest reliability were not as strong as predicted, these may be explained by situational

factors that contribute to varying levels of susceptibility across different scenarios. Findings from Study 3 regarding the cognitive factors that are predictive of susceptibility were partially replicated in Study 7, where a more realistic measure of susceptibility was employed – thus supporting the validity of the email task, suggesting that susceptibility is being measured in the same way as response likelihood during the office simulation.

Secondly, the office simulation study designed for Study 7 provides an alternative measure of susceptibility. Using this task, participants remained naïve to the purpose of the research, allowing for a more naturalistic measure of response likelihood. Although less extensively tested through the thesis, this task provides a valuable complementary measure, alongside the email legitimacy task. The lab-based nature of this simulation means that the researcher still has experimental control over stimuli and situational factors that might affect susceptibility though. This provides advantages for experimental control, which are not available with methods such as a simulated phishing attack. With integration of new stimuli to keep these measures up to date with the latest techniques used in phishing attacks, it is hoped that they can be utilised, or provide a basis for methodologies employed in further research in this field.

9.7.3 Additional contributions. As well as addressing theoretical and methodological issues surrounding email fraud susceptibility, this thesis aimed to consider the applied implications of the findings to real world victimisation and possible ways to reduce this. Efforts to educate users to-date have been fairly restricted to awareness and identifying cues that might indicate the illegitimacy of an

email (e.g. Sheng et al., 2007; Kumaraguru et al., 2010). Whilst such programmes have been shown to induce some short-term improvements in cyber security behaviour (Abawajy, 2014), there is no evidence that these have long-term benefits for users. The findings from this thesis demonstrate the importance of incorporating human factors into efforts to reduce susceptibility, based on individual differences in the decision-making process surrounding email management. By incorporating psychological perspectives on cognitive differences, as well as the influence of situational factors, training programmes could be developed that are tailored to the user and inform them about when and why they might be more at risk to victimisation. These ideas will be discussed more in section 9.9.

9.8 Limitations of the research

Whilst there are a number of valuable insights generated from the studies in this thesis, there are also some methodological limitations that should be taken into consideration when contemplating these. Due to ethical restrictions, the email legitimacy task was developed as a lab-based measure of susceptibility that assessed how well participants could differentiate between phishing and legitimate email stimuli presented to them. With variations across the experiments, this methodology allowed for a controlled assessment of email management behaviour, without the need to obtain information from participants deceptively (for example, gathering email addresses for a simulated phishing attack) or requiring any violation of their personal email accounts. However, the nature of this forced-choice judgment task means that it is not necessarily representative of real world email management behaviour. When identifying the emails, participants are explicitly told that they will

differentiate between phishing and legitimate emails, thus resulting in them actively seeking out clues to legitimacy. Evidence from Parsons et al. (2013) demonstrated that explicit knowledge of the nature of this type of task led to improved accuracy, demonstrating that participants do respond differently, with more vigilance, when they are actively seeking to recognise phishing. This might be a beneficial approach to take in real life, but it is unlikely to be how most users approach day-to-day email management.

Although an effort was made to give participants enough context in the instructions for the email task to enable them to complete the task, findings from Study 5 (in which participants were asked to think aloud about their decisions during the task) highlighted the issue of not knowing whether each email was relevant to the recipient. For example, many participants noted that they would only believe an email from Amazon confirming an order if they knew they had just ordered the specified item. In future versions of the email task, it might be beneficial to provide context with each email, giving information about the scenario in which the email was received and the recipient's prior interactions with the company that the email purported to come from. This would go some way to addressing the issue of whether each email was relevant to the recipient.

A further limitation of this task is the lack of personal involvement with the email stimuli, meaning that participants may not take the decisions as seriously as they would if they were managing their own emails, given that there is no actual threat of loss for them. In order to address this, a cash prize for the best performance in the email task was offered to participants, to encourage participants to make an effort in recognising both phishing and legitimate emails. Whilst this

cannot replicate a monetary threat to participants' own personal accounts, it provides an experimental incentive that goes some way to mimic real world email management. Finally, the line-up nature of this task means that participants may compare cues between emails, giving them additional information to base their judgments on, which would not be available to them when they were making a decision about a personal email. The findings from Study 4 imply that this may be the case, as accuracy decreased when participants saw either more phishing or more legitimate email stimuli, indicating that any comparisons they were drawing on the basis that they were viewing equal amounts of each were impairing judgments.

In order to address some of the issues highlighted, the office simulation study was developed. During this task, participants were naïve to the purpose of the research, and so their judgments about emails they were managing were not affected by an objective to spot phishing emails. Therefore, this was intended to provide a more ecologically valid assessment of email management behaviour. The partial replication of findings between this and Study 3 suggests that both are measuring susceptibility in the same way, which supports the use of the email task as a less time consuming, and more easily administered measure in a lab-based environment.

However, there are still a number of methodological limitations to the office simulation. Although participants were encouraged to get into character as much as possible, and were set up in an office-like environment, there was still no monetary threat to them personally at the end of the day. This means that, as with the email legitimacy task, participants may not have felt that they were taking risks in choosing to respond to an email, as there is nothing at stake for them compared to if a similar

email had been received to their personal email account. Further to this, this study is still lab-based, and although this allows for control over the situational variables being measured in the study, it may inflict the same biases as other lab-based tasks. For example, the presence of a video camera in the room may induce the Hawthorne Effect, disrupting participants' natural behaviour through increased awareness of being observed (Blum & Naylor, 1968).

The most valid assessment of email management behaviour would be to observe response likelihood in users' actual inboxes, either through a simulated phishing attack or through monitoring of their computer, with permission. However, each of these methods raises ethical concerns relating to consent and the embarrassment caused to participants if they decided to respond to a phishing, or simulated phishing attack. Further to this, assessment of actual susceptibility would require participants entering personal information, to demonstrate that they would engage with the email to the extent of victimisation. Again, the ethical constraints surrounding this assessment make it difficult to use in practise (discussed in more detail by Jones, Towse, & Race, 2015). Therefore, the tasks used in this thesis provide valuable, ethically sound, alternatives that provide insight into how well users recognise phishing emails and the cues they rely on to make these decisions.

The exploratory nature of investigating cognitive influences on susceptibility meant that it was necessary to include a range of measures in the initial experiments, as there was little previous literature to suggest what factors might be important. However, the down side of this is that there is an increased chance of an over fitted regression model – with cognitive variables found to be predictive due to noise generated by the high number of predictor variables included (Harrell, 2001;

Babyak, 2004). In Study 3, the cognitive variables were split into two sets, with each participant completing only one of these. This means that the number of variables in each model was reduced, so it is hoped that this also reduced the likelihood of over fitting. Only the predictive variables from Study 3 were then included in Study 7, which meant the number of predictor variables was reduced to four, which again should have reduced the likelihood of an over fitting model.

It is important to note that although there was some level of consistency between studies in the cognitive variables found to be predictive, there was also consistency in the low predictive power of these findings. It is apparent that the cognitive variables discussed, as well as the situational factors outlined only provide a partial explanation of susceptibility. Nonetheless, these are valuable insights that contribute to a wider picture of why some users become victims of a specific phishing email whilst thousands do not. These findings will encourage further research to provide a more robust model of susceptibility, to which individual differences in cognitive make-up, and situational factors will contribute.

Finally, it should be acknowledged that the majority of data collected in this thesis came from undergraduate students, meaning the age range of the sample is limited to younger internet users. Studies 2 and 2a provide a valuable data set from students and their parents at a University open day, giving a wider demographic. However, the later studies rely on a sample of limited demographic. Study 2a demonstrated a higher likelihood of past response behaviour in the parent than the student age groups, but age was not found to predict accuracy when judging the small set of email stimuli presented to participants. This may be because the younger age sample in this study had less exposure to email communications, and thus had

less opportunity to fall victim, rather than lower susceptibility. The student age group does provide a relevant sample group though, as they are often subject to targeted attacks relating to issues such as student loans or University account security. Although further research would be necessary to establish how the findings reported in this thesis relate to a wider demographic, findings relating to this limited age group provide valuable insight into this at risk population.

Further research would be necessary to establish whether patterns found in Studies 3-7 using lab-based measures of susceptibility with a sample of younger internet users, are mirrored in other demographic groups. If this were the case, then any educational efforts developed from the findings in this research would be relevant to all internet users. It is possible though that the individual differences in cognitive ability may vary between age groups, with older subjects often finding it harder to divide attention (Kok, 2000) and demonstrating lower inhibition (Hasher & Zacks, 1988; Hillman et al., 2006). This would require alternative approaches to consider how educational efforts might be tailored to reduce susceptibility in these cases. Therefore, without further understanding of how our findings generalise to the wider internet population, it should be concluded that any outcomes reported from this thesis are relevant only to the student sample at this stage.

9.9 Future directions

9.9.1 Development of methodology. The methodological issues highlighted above outline the need to compare findings with a more valid measure of susceptibility to email fraud in order to demonstrate the relevance of these to real world email management behaviour. The ultimate measure of susceptibility is to

conduct a simulated phishing attack in which participants would be required to enter personal information – demonstrating their actual likelihood of becoming a victim of fraud. The ethical constraints of conducting such an attack mean that this is not always practical though, unless there is assurance that the personal information entered by participants would not be stored anywhere. As an interim measure, a simulated phishing attack where likelihood to click a link or download a file would provide a more valid measure than the email legitimacy task, but unfortunately this was not possible in the scope of this thesis.

In terms of addressing the ethical issues and measuring likelihood to actually enter personal details, one possible alternative is to conduct research with past victims of email fraud. Although there are issues with this, relating to the self-report of circumstances surrounding victimisation and obtaining a comparable control sample, as discussed in Chapter 1. Past victims are users who have previously demonstrated susceptibility, and so their status is clear, given past behaviour. As part of the thesis, a study working with past victims was designed but unfortunately this was not executed due to practical issues in obtaining the victim sample. This would be an avenue to consider in future research though. This would also provide the opportunity to assess the validity of the email legitimacy task – by assessing how well past victims perform on this. This would require comparison to a control group though, which might be difficult to obtain, as this would need to be a group of users who have been exposed to the same types of emails as the victims but not have fallen victim. Otherwise, it may be argued that the control group could be just as susceptible but have never been under circumstances that led them to become victims of fraud.

9.9.2 Theoretical implications. Development of the most valid measure of susceptibility that is practically and ethically feasible will establish whether findings related to the cognitive make-up of the user reported here are replicable. With consistent replication of findings regarding inhibition, and the effect of cognitive load on susceptibility, a more conclusive model of response likelihood could be generated. With a thorough understanding of individual differences - building upon initial findings reported in this thesis, considering the application of dual-system theories of decision-making and the effects of divided attention - efforts to reduce susceptibility can be tailored accordingly.

9.9.3 Applications of research findings. As well as developing the most reliable and valid measure of susceptibility to encourage further research into individual differences in susceptibility, it is also important that efforts are made to apply findings into educational efforts to reduce susceptibility. Research from computer science has demonstrated how successful different techniques for training users about cues to phishing emails are, with interactive games proving the most beneficial (Sheng et al., 2007; Abawajy, 2012). However, the effectiveness of these training techniques is often only measured on a short-term basis, meaning the long-term benefits of such programmes remain unknown.

Evidence from Studies 1 and 5 suggests that participants are reliant on cues that are often out-dated or do not address the more sophisticated nature of techniques employed by some fraudsters. Further to this, Study 7 demonstrated that simply warning participants about the presence of phishing emails in a computer system did not reduce their likelihood to respond. In order to provide long-term

protection from email fraud victimisation, it is important to do more than warn users of the existence of phishing and teach them about the cues to look out for, which are constantly changing with the advancement of technological approaches to phishing. Users have demonstrated an inability to apply knowledge to unfamiliar security scenarios (Downs et al., 2006). Therefore, it might be beneficial to tailor educational efforts towards teaching users about how phishing attacks are conducted, what the fraudsters are attempting to gain from them, and how situational factors might put them at higher risk. Further to this, including assessments of cognitive reflection in training programmes could allow personalised feedback to users about potentially heightened susceptibility, based on the findings from the thesis.

Additional research looking at the long-term effects of such educational programmes would allow for assessment of their success. By using lab-based measures of susceptibility, this could be assessed both before and after the implementation of the programme, in addition to periodical assessments over an extended period of time to ensure that participants are still aware of how to protect themselves, regardless of technological advances on the part of the fraudsters.

9.10 Conclusions

In a novel area of research to the field of psychology, this thesis provides an exploratory step towards understanding how users become victims of email fraud. By considering both individual differences and external influences, the studies reported outline key considerations for future efforts to protect users against victimisation. Incorporating mixed methods for assessing susceptibility also contributes to the field, with the development of two tasks that can hopefully be

used in future research. Though there is still vast scope for further research in this area, especially in the development of tools to protect users, the research from this thesis provides insight into the question of why thousands of users can receive the same email, and yet only a small proportion will respond to it.

References

- Abawajy, J. (2014). User preference of cyber security awareness delivery methods. *Behaviour and Information Technology*, 33 (3), 237-248.
- Abu-Nimeh, S., Nappa, D., Wang, X., and Nair, S. (2007). A comparison of machine learning techniques for phishing detection. In *Proceedings of The Anti-Phishing Working Group's Second Annual eCrime Researchers Summit*, 60-69. ACM.
- Action Fraud (2016). Phishing. Retrieved from:
<http://www.actionfraud.police.uk/fraud-az-phishing> [Accessed 1 August 2016].
- Babiyak, M. A. (2004). What you see may not be what you get: A brief, nontechnical introduction to overfitting in regression-type models. *Psychosomatic Medicine*, 66 (3), 411-421.
- Bailey, J., Mitchell, R., and Jensen, B. (2008). Analysis of student vulnerabilities to phishing. *Proceedings of the Fourteenth Americas Conference on Information Systems* (pp. 1-10). Toronto: Association of Information Systems.
- Barrett, L. F., and Tugade, M. M. (2004). Individual differences in working memory capacity and dual-process theories of the mind. *Psychological Bulletin*, 130 (4), 553-573.
- Beede, K.E., and Kass, S. J. (2006). Engrossed in conversation: The impact of cell phones on simulated driving performance. *Accident Analysis & Prevention*, 38, 415-421.
- Bergholz, A., De Beer, J., Glahn, S., Moens, M. F., Paaß, G., and Strobel, S. (2010). New filtering approaches for phishing email. *Journal of Computer Security*, 18, 7-35.
- Blum, M. L., and Naylor, J. (1968). *Industrial Psychology: It's Theoretical and Social Foundations*. New York: Harper and Row.
- Brand, M., Heinze, K., Labudda, K., and Markowitsch, H. J. (2008). *Cognitive Processing*, 9 (3), 159-173.

- Bullée, J. H., Montoya Morales, A. L., Junger, M., and Hartel, P. H. (2016). Telephone-based social engineering attacks: An study testing the success and time decay of an intervention. In *Proceedings of the Singapore Cyber-Security Conference (SG-CRC) 2016* (pp. 107-114). IOS Press.
- Carmines, E. G., and Zeller, R. A. (1979). *Reliability and Validity Assessment* (Vol. 17). Sage Publications.
- Cialdini, R. B. (1993). *Influence: The Psychology of Persuasion*. New York: Quill William Morrow.
- Cicchetti, D. V. (1994). Guidelines, criteria and rules of thumb for evaluating normed and standardized assessment instruments in psychology. *Psychological Assessment*, 6, 284-290.
- Citizen's Advice Bureau (2016). Phishing – spam emails and fake websites. Retrieved from: <https://www.citizensadvice.org.uk/consumer/scams/scams/common-scams/computer-and-online-scams/phishing-spam-emails-and-fake-websites/> [Accessed 1 August 2016].
- Cokely, E. T., and Kelley, C. M. (2009). Cognitive abilities and superior decision making under risk: A protocol analysis and process model evaluation. *Judgement and Decision Making*, 4 (1), 20-33.
- Colflesh, C. J., and Conway, A. R. (2007). Individual differences in working memory capacity and divided attention in dichotic listening. *Psychonomic Bulletin & Review*, 14(4), 699-703.
- Conway, A. R., Jarrold, C., Kane, M., Miyake, A., and Towse, J. N. (Eds.). (2007). *Variation in Working Memory*. New York: Oxford University Press.
- Copes, H., Kerley, K. R., Mason, K. A., and Van Wyk, J. (2001). Reporting behavior of fraud victims and Black's theory of law: An empirical assessment. *Justice Quarterly*, 18(2), 343-363.
- Daneman, M., and Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, 19 (4), 450-466.

- Daneman, M., and Newson, M. (1992). Assessing the importance of subvocalization during normal silent reading. *Reading and Writing*, 4 (1), 55-77.
- de Groot, A. D. (1946). *Het Denken van den Schaker* (Thought in Chess). Doctoral thesis. University of Amsterdam, Amsterdam (in Dutch).
- Dong, X., Clarke, J. A., and Jacob, J. (2008). Modelling user-phishing interaction. In *Proceedings of Human System Interaction*, 627-632. IEEE.
- Downs, J. S., Holbrook, M., and Cranor, L. F. (2006). Decision strategies and susceptibility to phishing. In *Proceedings of the Second Symposium on Usable Privacy and Security* (pp. 79-90). ACM.
- Downs, J. S., Holbrook, M., and Cranor, L. F. (2007). Behavioural response to phishing risk. In *Proceedings of the Anti-Phishing Working Groups Second Annual eCrime Researchers Summit*, 37-44. ACM.
- Duncker, K. (1945). *On Problem Solving*. Washington: The American Psychological Society.
- Engle, R. W. (1996). Working memory and retrieval: An inhibition-resource approach. In J. T. E. Richardson, R. W. Engle, L. Hasher, R. H. Logie, E. R. Stoltzfus, and R. T. Zacks (Eds.), *Working Memory and Human Cognition* (pp. 89-119). London: Oxford University Press.
- Ericsson, K. A., and Simon, H. A. (1993). *Protocol Analysis: Verbal Reports as Data* (revised edition). Cambridge, MA: MIT Press.
- Ericsson, K. A., and Simon, H. A. (1998). How to study thinking in everyday life: Contrasting think-aloud protocols with descriptions and explanations of thinking. *Mind, Culture, and Activity*, 5 (3), 178-186.
- Erickson, T. D., and Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behaviour*, 20 (5), 540-551.
- Eriksen, B. A., and Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & psychophysics*, 16(1), 143-149.

- Evans, J. B. T. (2003). In two minds: Dual-process accounts of reasoning. *Trends in Cognitive Science*, 7, 454-459.
- Fan, J., McCandliss, B. D., Sommer, T., Raz, A., and Posner, M. I. (2002). Testing the efficiency and independence of attentional networks. *Journal of Cognitive Neuroscience*, 14, 340-347.
- Fette, I., Sadeh, N., and Tomasic, A. (2007). Learning to detect phishing emails. In *Proceedings of the 16th International World Wide Web Conference*, 649-656. ACM.
- Finn, P., and Jakobsson, M. (2007). Designing ethical phishing experiments. *Technology and Society Magazine, IEEE*, 26(1), 46-58.
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, 19 (4), 25-42.
- Friedman, B., Hurley, D., Howe, D. C., Felten, E., and Nissenbaum, H. (2002). Users' conceptions of web security: A comparative study. In *CHI'02 Extended Abstracts on Human Factors in Computing Systems* (pp. 746-747). ACM.
- Garera, S., Provos, N., Chew, M., and Rubin, A. D. (2007). A framework for detection and measurement of phishing attacks. In *Proceedings of the 2007 ACM Workshop on Recurring Malcode*, 1-8. ACM.
- Germine, L., Nakayama, K., Duchaine, B. C., Chabris, C. F., Chatterjee, G., and Wilmer, J. B. (2012). Is the Web as good as the lab? Comparable performance from Web and lab in cognitive/perceptual experiments. *Psychonomic Bulletin and Review*, 19 (5), 847-857.
- Gilovich, T., Griffin, D., and Kahneman, D. (2002). *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press.
- Goldberg, L. R. (1999). A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. In I. Mervielde, I. Deary, F. De Fruyt, & F. Ostendorf (Eds.), *Personality Psychology in Europe*, 7, 7-28. Tilburg, The Netherlands: Tilburg University Press.

- Gottfredson, M. R., and Hirschi, T. (1990). *A General Theory of Crime*. Stanford: Stanford University Press.
- Guéguen, N., and Jacob, C. (2002). Solicitation by e-mail and solicitor's status: A field study of social influence on the web. *CyberPsychology & Behaviour*, 5 (4), 377-383.
- Hadlington, L. (2015). Cognitive factors in online behaviour. In A. Attrill (Ed.), *Cyberpsychology* (pp. 249-267). New York: Oxford University Press.
- Harrell, F. E. Jr. (2001). *Regression modelling strategies: With implications to linear models, logistic regression and survival analysis*. New York: Springer.
- Harrison, B., Vishwanath, A., and Rao, R. (2016). A user-centered approach to phishing susceptibility: The role of a suspicious personality in protecting against phishing. In *2016 49th Hawaii International Conference on System Sciences (HICSS)* (pp. 5628-5634). IEEE.
- Harrison, M. (2015). *Understanding the mechanisms underlying scam vulnerability* (Unpublished undergraduate dissertation). Lancaster University, Lancaster.
- Hasher, L., and Zacks, R. T. (1988). Working memory, comprehension, and aging: a review and a new view. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation* (vol. 22, pp. 193-225). San Diego: Academic Press.
- Hillman, C. H., Motl, R. W., Pontifex, M. B., Posthuma, D., Stubbe, J. H., Boomsma, D. I., and de Gaus, E. J. C. (2006). Physical activity and cognitive function in a cross-section of younger and older community-dwelling individuals. *Health Psychology*, 25 (6).
- Hinson, J. M., Jameson, T. L., and Whitney, P. (2003). Impulsive decision making and working memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 29 (2), 298-306.
- Hiscock, M. (1986). Lateral eye movements and dual-task performance. In H. J. Hanney (Ed.), *Experimental Techniques in Human Neuropsychology* (pp. 264-308). New York: Oxford University Press.

- Holtfreter, K., Reisig, M. D., Piquero, N. L., and Piquero, A. R. (2010). Low self-control and fraud offending, victimization, and their overlap. *Criminal Justice and Behaviour*, 37 (2), 188-203.
- Hong, K. W., Kelley, C. M., Tembe, R., Murphy-Hill, E., and Mayhorn, C. B. (2013). Keeping up with the Joneses: Assessing phishing susceptibility in an email task. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, September 2013*, 57 (1), 1012-1016.
- Hoyle, R. H., Stephenson, M. T., Palmgreen, P., Lorch, E. P., and Donohew, R. L. (2002). Reliability and validity of a brief measure of sensation seeking. *Personality and Individual Differences*, 32, 401-414.
- Islam, R., and Abawajy, J. (2013). A multi-tier phishing detection and filtering approach. *Journal of Network and Computer Applications*, 36 (1), 324-335.
- Jakobsson, M., Tsow, A., Shah, A., Blevins, E., and Lim, Y. K. (2007). What instills trust? A qualitative study of phishing. In *International Conference on Financial Cryptography and Data Security* (pp. 356-361). Berlin: Springer.
- Jagatic, T. N., Johnson, N. A., Jakobsson, M., and Menczer, F. (2005). Social phishing. *Communications of the ACM*, 50 (10), 94-100.
- Jones, H. S., Towse, J., and Race, N. (2015). Susceptibility to email fraud: A review of psychological perspectives, data-collection methods, and ethical considerations. *International Journal of Cyber Behavior, Psychology, and Learning*, 5 (3), 13-29.
- Kahneman, D. (2000). A psychological point of view: Violations of rational rules as a diagnostic of mental processes. *Behavioural and Brain Sciences*, 23, 681-683.
- Kane, M. J., Bleckley, M. K., Conway, A. R. A., and Engle, R. W. (2001). A controlled-attention view of working-memory capacity. *Journal of Experimental Psychology: General*, 130 (2), 169-183.
- Kane, M. J., and Engle, R. W. (2000). Working memory capacity, proactive interference, and divided attention: Limits on long-term memory retrieval.

Journal of Experimental Psychology: Learning, Memory, and Cognition, 26 (2), 336-358.

- Karat, J., Karat, C., and Brodie, C. (2009). Human-computer interaction viewed from the intersection of privacy, security, and trust. In A. Sears, and J. A. Jacko (Eds.), *Human-Computer Interaction: Design Issues, Solutions, and Applications*, 311-330. Boca Raton, FL: CRC Press.
- King, M. F., and Bruner, G. C. (2000). Social desirability bias: A neglected aspect of validity testing. *Psychology and Marketing*, 17 (2), 79-103.
- Klein, R. A., Ratcliff, K. A., Vianello, M., Adams Jr., R. B., Bahník, Š., Bernstein, M. J., ... and Cemalcilar, Z. (2014). Investigating variation in replicability. *Social Psychology*, 45, 142-152.
- Kok, A. (2000). Age-related changes in involuntary and voluntary attention as reflected in components of the event-related potential (ERP). *Biological Psychology*, 54 (1), 107-143.
- Kruglanski (1990). Motivations for judging and knowing: Implications for causal attribution. In E. T. Higgins & R. M. Sorrentino (Eds.), *The handbook of motivation and cognition: Foundation of social behaviour*, 2, 333–368. New York: Guilford Press.
- Kyllonen, P., and Christal, R. E. (1990). Reasoning ability is (little more than) working memory capacity? *Intelligence*, 14, 389-433.
- Langenderfer, J., and Shimp, T. A. (2001). Consumer vulnerability to scams, swindles, and fraud: A new theory of visceral influences on persuasion. *Psychology and Marketing*, 18 (7), 763-783.
- Leung, F., and Savithiri, R. (2009). Spotlight on focus groups. *Canadian Family Physician*, 55 (2), 218-219.
- Li, H., Edwards, S. M., and Lee, J. H. (2002). Measuring the intrusiveness of advertisements: Scale development and validation. *Journal of Advertising*, 31 (2), 37-47.

- Macmillan, N. A. (1993). Signal detection theory as data analysis method and psychological decision model. In G. Keren & C. Lewis (Eds.), *A Handbook for Data Analysis in the Behavioural Sciences: Methodological Issues* (pp. 21-57). Hillsdale, NJ: Erlbaum.
- Macmillan, N. A. (2002). Signal detection theory. In H. Pashler & J. Wixted (Eds.), *Stevens' Handbook of Experimental Psychology* (Third edition, vol. 4, pp. 43-90). New York: John Wiley & Sons, Inc.
- Mack, S. (2014). *Reasoning and judgements made in an online capacity. An exploration of how phishing emails influence decision making strategies* (Unpublished undergraduate dissertation). Lancaster University, Lancaster.
- Markovits, H., Doyon, C., and Simoneau, M. (2002). Individual differences in working memory and conditional reasoning with concrete and abstract content. *Thinking and Reasoning*, 8 (2), 97-107.
- McCormac, A., Calic, D., Parson, K., Zwaans, T., Butavicius, M., and Pattinson, M. (2016). *Test-retest reliability and internal consistency of the Human Aspects of Information Security Questionnaire (HAIS-Q)*. Paper to be presented at Australasian Conference on Information Systems, Wollongong, Australia.
- Mitnick, K. D., and Simon, W. L. (2002). *The Art of Deception*. Indianapolis: Wiley Publishing, Inc.
- Miyake, A., Just, M. A., and Carpenter, P. A. (1994). Working memory constraints on the resolution of lexical ambiguity: Maintaining multiple interpretations in neutral contexts. *Journal of Memory and Language*, 33, 175-202.
- Modic, D., and Anderson, R. J. (2014). We will make you like our research: The development of a susceptibility-to-persuasion scale. *Social Sciences Research Network*. Retrieved from <http://ssrn.com/abstract=2446971> [Accessed 16 July 2014].
- Modic, D., and Lea, S. E. G. (2011). *How neurotic are scam victims, really? The big five and Internet scams*. Paper presented at the 2011 Conference of the International Confederation for the Advancement of Behavioral Economics and Economic Psychology, Exeter: UK.

- Modic, D., and Lea, S. E. G. (2013). Scam compliance and the psychology of persuasion. *Social Sciences Research Network*. Retrieved at <http://ssrn.com/abstract=2364464> [Accessed 16 July 2014].
- Myers, S. (2007). Introduction to phishing. In M. Jakobsson, and S. Myers (Eds.), *Phishing and Countermeasures* (pp.1-29). New Jersey: John Wiley & Sons, Inc.
- National Fraud Authority (2011). A quantitative segmentation of the UK population. Helping to determine how, why and when citizens become victims of fraud. Accessed at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/118481/national-fraud-segmentation.pdf [Accessed 8 July 2013].
- Nederhof, A. J. (1985). Methods of coping with social desirability bias: A review. *European Journal of Social Psychology*, 15, 263-280.
- Norton (2014). Online fraud: Phishing. Retrieved from <http://uk.norton.com/cybercrime-phishing> [Accessed 12 July 2014].
- Onwuegbuzie, A. J., Dickinson, W. B., Leech, N. L., and Zoran, A. G. (2009). A qualitative framework for collecting and analyzing data in focus group research. *International Journal of Qualitative Methods*, 8 (3), 1-21.
- Office for National Statistics (2015). Crime Statistics, Year Ending June 2015. Available at: <http://www.ons.gov.uk/ons/rel/crime-stats/crime-statistics/year-ending-june-2015/sty-fraud.html> [Accessed: 1 August 2016].
- Pak, K. B. S., and Shadel, D. P. (2011). AARP Foundation National Fraud Victim Study. Retrieved from <http://assets.aarp.org/rgcenter/econ/fraud-victims-11.pdf> [Accessed 05 September 2013].
- Park, H., and Reder, L. M. (2004). Moses illusion: Implication for human cognition. *Cognitive Illusions*, 275-291.
- Parsons, K., McCormac, A., Pattinson, M., Butavicius, M., and Jerram, C. (2013). Phishing for the truth: A scenario-based study of users' behavioural response to emails. In *IFIP International Information Security Conference* (pp. 366-378). Berlin: Springer.

- Parsons, K., McCormac, A., Butavicius, M., Pattinson, M., and Jerram, C. (2014). Determining employee awareness using the human aspects of information security questionnaire (HAIS-Q). *Computers and Security*, 42, 165-176.
- Paulhus, D. L. (1991). Measurement and control of response bias. In J. P. Robinson et al. (Eds.), *Artifact in Behavioural Research* (pp. 17-59). New York: Academic Press.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2 (10).
- Porcelli, A. J., and Delgado, M. R. (2009). Acute stress modulates risk taking in financial decision-making. *Psychological Science*, 20 (3), 278-283.
- Pratt, T. C., Holtfreter, K., and Reisig, M. D. (2010). Routine online activity and internet fraud targeting: Extending the generality of routine activity theory. *Journal of Research in Crime and Delinquency*, 47 (3), 267-296.
- Reder, L. M., and Kusbit, G. W. (1991). Locus of the Moses illusion: Imperfect encoding, retrieval, or match? *Journal of Memory and Language*, 30 (4), 385-406.
- Redick, T. S., and Engle, R. W. (2006). Working memory capacity and attention network test performance. *Applied Cognitive Psychology*, 20, 713-721.
- Redick, T. S., Heitz, R. P., and Engle, R. W. (2007). Working memory capacity and inhibition: Cognitive and social consequences. In D. S. Gorfein, and C. M. MacLeod (Eds.), *Inhibition in Cognition* (pp. 125-142). Washington: American Psychological Association.
- Roda, C. (2011). Human attention and its implications for human-computer interaction. In C. Roda (Ed.), *Human Attention in Digital Environments* (pp. 11-62). Cambridge: University Press.
- Roets, A., and van Hiel, A. (2011). Item selection and validation of a brief, 15-item version of the Need for Closure scale. *Personality and Individual Differences*, 50, 90-94.

- Salah, K., Alcaraz Calero, J. M., Zeadally, S., Al-Mulla, S., and Alzaabi, M. (2013). Using cloud computing to implement a security overlay network. *IEEE Security & Privacy*, 11 (1), 44-53.
- Schneier, B. (2000a). *Secrets & Lies: Digital Security in a Networked World*. Indianapolis: Wiley Publishing Inc.
- Schneier, B. (2000b). *Semantic Attacks: The Third Wave of Network Attacks*. Schneier on Security blog. Retrieved from: <https://www.schneier.com/cryptogram/archives/2000/1015.html#1> [Accessed 2 August 2016].
- Schreck, C. J. (1999). Criminal victimization and low self-control: An extension and test of a general theory of crime. *Justice Quarterly*, 16.
- Shadel, D. P., and Pak, K. B. S. (2007). *The Psychology of Consumer Fraud*. (Unpublished doctoral thesis). Tilburg University, Netherlands.
- Sheng, S., Holbrook, M. B., Kumaraguru, P., Cranor, L. F., and Downs, J. S. (2010). Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, (pp. 373-382). ACM.
- Slavin, S. (2015). PsyScript 3 (Version 0.9.0) [Computer software]. Lancaster University. Retrieved July 3, 2014. Available from: <http://www.lancaster.ac.uk/psychology/research/research-software/psyscript3/>
- Slowiaczek, M. L., and Clifton Jr., C. (1980). Subvocalization and reading for meaning. *Journal of Verbal Learning and Verbal Behavior*, 19 (5), 573-582.
- Smagorinsky, P. (1998). Thinking and speech and protocol analysis. *Mind, Culture, and Activity*, 5 (3), 157-177.
- Smithson, J. (2000). Using and analysing focus groups: Limitations and possibilities. *International Journal of Social Research Methodology*, 3 (2), 103-119.
- Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Erlbaum.

- Stanovich, K. E., and West, R. F. (2002). Individual differences in reasoning: Implications for the rationality debate. In Gilovich, T., Griffin, D., and Kahneman, D. (Eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 421-440). New York: Cambridge University Press.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 18 (6), 643–662.
- Symantec (2014). Internet Security Threat Report 2014. Retrieved from: http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf [Accessed 4 August 2016].
- Tangney, J. P., Baumeister, R. F., and Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of personality*, 72 (2), 271-324.
- The Radicati Group (2015). Email Statistics Report, 2015-2019. Retrieved from: <http://www.radicati.com/wp/wp-content/uploads/2015/02/Email-Statistics-Report-2015-2019-Executive-Summary.pdf> [Accessed 23 November 2016].
- Tversky, A., and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5 (2), 207-232.
- Tversky, A., and Kahneman, D. (1975). Judgment under uncertainty: Heuristics and biases. In *Utility, Probability, and Human Decision Making* (pp. 141-162). Netherlands: Springer.
- van Someren, M., Barnard, Y. F., and Sandberg, J. A. (1994). *The Think Aloud Method: A Practical Approach to Modelling Cognitive Processes*. London: Academic Press.
- van Wilsem, J. (2013). ‘Bought it, but never got it’: Assessing risk factors for online consumer fraud victimization. *European Sociological Review*, 29 (2), 168-178.
- Vishwanath, A. (2015). Examining the distinct antecedents of e-mail habits and its influence on the outcomes of a phishing attack. *Journal of Computer-mediated Communication*, 20 (5), 570-584.

- Vishwanath, A., Herath, T., Chen, R., Wang, J., and Rao, H. R. (2011). Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model. *Decision Support Systems*, 51, 576-586.
- Webster, D. M., and Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67, 1049–1062.
- Whiteside, S. P., and Lynam, D. R. (2001). The five factor model and impulsivity: Using a structural model of personality to understand impulsivity. *Personality and Individual Differences*, 30, 669-689.
- Whiteside, S. P., Lynam, D. R., Miller, J. D., and Reynolds, S. K. (2005). Validation of the UPPS impulsive behaviour scale: A four-factor model of impulsivity. *European Journal of Personality*, 19, 559-574.
- Whitty, M. T., and Buchanan, T. (2012). The online romance scam: A serious cybercrime. *CyberPsychology, Behavior, and Social Networking*, 15(3), 181-183.
- Workman, M. (2007). Gaining access with social engineering: An empirical study of threat. *Information Systems Security*, 16 (6), 315-331.
- Wright, R. T., and Marett, K. (2010). The influence of experiential and dispositional factors in phishing: An empirical investigation of the deceived. *Journal of Management Information Systems*, 27 (1), 273-303.
- Yan, Z., and Gozu, H. Y. (2012). Online Decision-Making in Receiving Spam Emails Among College Students. *International Journal of Cyber Behavior, Psychology and Learning (IJCBL)*, 2(1), 1-12.
- Zuckerman, M., Eysenck, S., and Eysenck, H. J. (1978). Sensation seeking in England and America: cross-cultural age and sex comparisons. *Journal of Consulting and Clinical Psychology*, 46, 139–149.

Appendices

All appendices mentioned in the thesis can be viewed at

[<https://dx.doi.org/10.17635/lancaster/researchdata/117>].